

# Analysis of the Evolution of C-Space Models built through Incremental Exploration

Marco Morales, Roger Pearce, Nancy M. Amato  
{marcom,rap2317,amato}@cs.tamu.edu

Technical Report TR06-013  
Parasol Lab  
Department of Computer Science  
Texas A&M University  
College Station, TX 77843-3112

September 18, 2006

## Abstract

Many sampling methods for motion planning explore the robot's configuration space (C-space) starting from a set of configuration(s) and incrementally explore surrounding areas to produce a growing model of the space. Although there is a common understanding of the strengths and weaknesses of these techniques, metrics for analyzing the incremental exploration process and for evaluating the performance of incremental samplers have been lacking. We propose the use of local metrics that provide insight into the complexity of the different regions in the model and global metrics that describe the process as a whole. These metrics only require local information and can be efficiently computed. We illustrate the use of our proposed metrics to analyze representative incremental strategies including the Rapidly-exploring Random Trees, Expansive Space Trees, and the original Randomized Path Planner. We show how these metrics model the efficiency of C-space exploration and help to identify different modeling stages. In addition, these metrics are ideal for adapting space exploration to improve performance.

## 1 Introduction

Probabilistic sampling in C-space allows the solution of many interesting motion planning problems that were previously impractical [1, 3, 6, 7, 10, 12, 14, 15, 23]. This approach avoids computing an exact representation of the planning space by sampling the configuration space (C-space): the space of feasible robot configurations and transitions between configurations. The result is an approximate model that encodes motions that the robot can perform. This general methodology has found diverse applications, such as the study of protein folding in Biology and Chemistry [2, 5, 20–22], the development of virtual prototypes in manufacturing and mechanical design [4, 8], and the simulation of characters for animation and games [17, 18].

Much work has been done to improve probabilistic sampling, especially on mechanisms to bias sampling towards regions of the space that model highly constrained robot motions. In this paper we are particularly interested in incremental planners that, starting from a set of configurations, incrementally expand the model by exploring surrounding C-space regions [3, 6, 10, 16]. A general set of metrics to evaluate the performance and efficiency of these methods has been lacking.

We propose a set of metrics that measure local and global aspects of sampling and we apply them to analyze several incremental samplers. Previously [19], we proposed a set of metrics that characterizes samples as they are added to a model and we applied them to analyze different Probabilistic Roadmap Methods (PRMs) [12] that model the space through a *roadmap* sampled globally. In [24] we applied these metrics to decide when to stop roadmap construction. Here, we propose to extract local metrics that provide insight into the complexity of the different regions in the model and global metrics that monitor the modeling process.

Sampling metrics can be applied to analyze, to compare, and to improve planners. We can gain insight into the strengths and weaknesses of a sampling strategy by analyzing features of the model produced as it progresses. We can compare different planners by studying their ability to discover new regions over time. We can improve planners by adapting sampling to the features of the space discovered or by decid-

ing what method to apply in each case by evaluating their performance as in [24] and [11].

Without loss of generality, we study the case where there is a single model component undergoing expansion. After introducing the general methodology, we apply these metrics to study the growth of Rapidly-exploring Random Trees (RRTs) in its non-biased (RRT-Expand) [16] and biased (RRT-Connect) [13] versions, Expansive Space Trees (ESTs) [10], and the original Randomized Path Planner (RPP) [3].

## 2 C-space Modeling through Incremental Exploration

Probabilistic planners build an approximate model of the valid C-space (C-free) by selecting random samples of configurations and transitions between them according to some strategy. The resulting model is usually a graph or tree with vertices representing feasible configurations and edges representing feasible transitions between configurations. Some planners explore the space to find new samples following some strategy to expand the model in an incremental fashion [3, 6, 10, 16].

Methods that model C-space through incremental exploration usually start from some set of valid configurations where the root trees that are expanded in increments as they explore the space surrounding areas already in the tree. In each increment, a node is added to the tree by a local exploration around a selected node. RRTs [16] are expanded by selecting a random configuration  $x_{rand}$  from the C-space, identifying the closest node  $x_{near}$  in the tree, and adding to the tree the last valid configuration  $x_{new}$  on the line of length  $\delta$  that goes from  $x_{near}$  to  $x_{rand}$ . ESTs [10] are expanded by selecting a node  $x_g$  in the tree based on a probability biased towards unexplored areas, and adding the last valid configuration on the line of length  $\delta$  that goes from  $x_g$  in a random direction also biased towards unexplored areas.

Some planners bias the exploration of C-space to solve specific queries. RRT-Connect [13] roots one tree at the start configuration and another at the goal configuration. Expanding one tree by a step, it attempts to connect the closest node from the other tree to the newly added node of the first tree. This process is repeated, alternating trees for expansion and

connection, until the trees are connected or a maximum number of iterations is reached. RPP [3] roots a tree at the start configuration, each node is expanded by searching for a closeby node with a smaller value of a given potential function that attracts it to the goal configuration. When no such node can be found, the latest configuration is a local minimum, and a small number of escape attempts are tried through Brownian motion. If every escape attempt fails the expansion restarts from one previous node picked at random from the parents of the local minimum.

### 3 Evolution of C-space Models built through Incremental Exploration

The goal of a C-space model is to reflect the coverage and connectivity of the underlying C-space. However, it is usually difficult to sample highly constrained areas of the C-free, affecting the chances of the planner to correctly model its coverage and connectivity.

Although finding the exact connectivity and coverage achieved with a C-space model is only practical in some cases, generally we can extract information from the construction process to evaluate the progress achieved by the planner. On one hand, studies have been made to exhaustively compare the coverage of the C-space of a given problem with those achieved by different PRM-based sampling strategies [9]. This kind of study requires a discretization of the space that is not always feasible to obtain. On the other hand, a set of metrics proposed in [19] characterizes the samples as they are added to the model to identify the contribution of the sample to the quality of the model. Efficient approximations of these metrics can be obtained during model construction. Here we discuss these metrics and propose new metrics in the context of planners that incrementally explore the C-space from a given configuration.

#### 3.1 Effects on coverage and connectivity during incremental exploration

Following [19], when a planner adds a valid configuration  $v$  and a selected subset of all its valid connections

to a model  $M$ , the connectivity and coverage of the original model changes in exactly one of the following ways:

1. **cc-create** —  $v$  lies outside the coverage region of all the components in  $M$ . A new component is then created as seen in Fig. 1(b,c).
2. **cc-merge** —  $v$  lies inside the coverage region of more than one component of  $M$ . These components can be merged as seen in Fig. 2(a).
3. **cc-expand** —  $v$  lies inside the coverage of exactly one component of  $M$  and it increases the coverage of the component as seen in Fig. 2(b).
4. **cc-oversample** —  $v$  lies inside the coverage of exactly one component of  $M$  without increasing the coverage of the component, this is an oversample that does not improve the model as seen in Fig. 2(c).




Figure 1: (a) A two-dimensional environment with a polygonal obstacle. (b) One sample is added. (c) One more sample is added. Note that the coverage of the two samples overlaps above the obstacle.

Here, we pay more attention to *cc-expand* and *cc-oversample* nodes because they carry most of the information in incremental exploration. *cc-create* nodes are produced only when a tree is started and *cc-merge* nodes when building more than one tree or when connecting the tree to a goal configuration. Otherwise, the nodes are *cc-expand* or *cc-oversample*.

#### 3.2 Local and global characterization of C-space sampling

We propose to compile a history of local metrics and global metrics to gain insight into the planning process and the C-space discovered. Local metrics help




Figure 2: New samples (hollow dots) are added. (a) A *cc-merge* sample falls in the intersecting coverage of the two existing samples connecting the two regions. (b) A *cc-expand* sample increases the coverage of the component to the left, but does not change the connectivity. (c) *Oversampling* does not increase coverage or connectivity.

us to understand the complexity of the C-space regions around sampled areas. Global metrics help us to understand the overall progress made by the planner over time. These metrics are general, independent of the dimensionality of the problem, and can be applied to any probabilistic sampler that explores C-space incrementally.

### 3.3 Local metrics

Local metrics characterize the complexity of the C-space region around a particular sample. We define as *growth sites* the nodes that have been selected for expansion, regardless of whether the expansion attempt is successful. By keeping track of growth sites we can analyze the distribution of the sampling. By keeping track of the history of the success of the growth sites, we can estimate the complexity of their surrounding regions. Also, we can compute more expensive tests for approximating whether new nodes are *cc-expand* or *cc-oversample*. Thus, the following metrics are maintained for each growth site:

- **growth attempts** – Number of times the node has been selected for incremental expansion. Node 1 in Fig. 3 (left) has four growth attempts.
- **successful growths** – Number of successful growth attempts. Node 1 in Fig. 3 (left) has two successful growths.
- **obstruction ratio** – Complement of the ratio of successful growths to growth attempts. The ob-

struction ratio is undefined for nodes that have no growth attempts (nodes in the fringe that have not been selected). Node 1 in Fig. 3 (left) has an obstruction ratio of 0.5.

In addition, we compute an **expansion ratio** for new nodes – the expansion ratio for a node is the complement of the percentage of neighbors of the parent that the node can connect to. This test may require additional validity tests than specified by the incremental planner. Node 10 in Fig. 3 (right) has an expansion ratio of 0.5 whereas node 9 has an expansion ratio of 0.




Figure 3: Incremental model of the problem in Figure 1. Numbers indicate vertices and their placement order. Solid arrows indicate successful connections and their direction (a) Dotted lines indicate failure growth attempts from the source node. (b) Hollow arrows indicate connection attempts for the expansion ratio test when nodes 9 and 10 were added to the model.

These metrics model the complexity of the areas sampled in the past as well as the distribution of the sampling at the node level. If the model is expanding uniformly, most growth sites should have a similar number of *growth attempts*. In RRTs, sites that have more than average growth attempts may point to areas of the tree that are repeatedly selected for growth, but whose Voronoi regions fail to shrink. Sites frequently selected with low *successful growths* have a high *obstruction ratio* indicating a highly constrained region. Ideally, the sites with low obstruction ratio should have smaller probabilities of selection for growth.

We propose to group nodes in order to gain insight into the complexity of their neighborhoods. We

group growth sites in *bubbles*, n-dimensional spheres of radius  $r$ . We create a bubble centered on every new growth site that does not lie inside any existing bubble. When a new growth site lies inside an existing bubble, a link is made between the growth site and the bubble. This way, the bubbles represent the coverage of the C-space achieved by the planner approximated at a resolution defined by  $r$ . By combining the node statistics of the growth sites in each bubble, we obtain a characterization of the space it represents that provides information that can be used to adjust the sampling distributions. For example, bubbles with many growth sites that have low obstruction ratio may indicate oversampling.

Instead of computing the volume covered by the bubbles, we compute the ratio of bubbles to total growth sites as an indication of useful expansions of the model.

### 3.4 Global metrics

Global metrics keep track of the evolution of the incremental model, they aggregate local metrics and global information over slices of time during the planning. We divide the incremental expansion attempts into sets (bins) of equal size and aggregate the following statistics:

- **percentage of successful growths** – Indicates how difficult it is to expand the model.
- **average obstruction ratio of growth sites** – A high value indicates a highly constrained region of the space.
- **ratio of bubbles to growth sites** – Indicates the efficiency of the sampler in discovering unexplored regions. A high value indicates that most growth sites produce new bubbles, a low value indicates oversampling.
- **ratio of bubbles to growth sites from successful expansions** – Similar to the previous, but only includes bubbles and growth sites created during successful expansions.

The evolution of global metrics is correlated with the ability of the planner to increase its knowledge of the planning space over time. A high *percentage of successful growths* and low *average obstruction ratio*

*of growth sites* over some time indicates that the areas being covered are mostly free. A low *ratio of bubbles to growth sites* indicates oversampling.

## 4 Application: Analysis of the evolution of incremental planners

We demonstrate the application of local and global metrics in an experimental analysis of the evolution of models produced with different incremental planners applied to different environments as described below:

- **Problem domain** – The problems studied involve rigid robots moving in three-dimensional environments.
- **Incremental strategies** – RRT-Expand [16], EST [10], and the biased RRT-Connect [13] and RPP [3].
- **Expansion parameters** – In each iteration, we apply the expansion rule of the method studied. We define a maximum number of samples and bin them into even intervals. Since each planner may produce more than one sample per iteration, bin size had some variations. For each planner and environment we used the parameters ( $\delta_{t,q}$ ) that showed the best performance.
- **Study parameters** – The bubble radius  $r$  was chosen to be about 15% more than the expansion length. Smaller bubbles overestimate the coverage increase, larger bubbles underestimate it.

### 4.1 Environments

We selected environments that have both open spaces and narrow passages in order to evaluate the metrics. We evaluate the way in which the metrics characterize the planner’s ability to cover the space finding their way through the passages. All the robots studied here were six-DOF rigid bodies.

The *walls* environment (Fig. 4-walls) has a cubic robot that must pass through four short narrow passages from right to left.




Figure 4: (*walls*) 6-DOF cubic robot, four short passages; (*maze*) solid view, 6-DOF robot, and wire view; (*hook*) 6-DOF hook robot, two medium passages.

The *maze* (Fig. 4-maze) environment has a 6-DOF rigid body that must move through a long narrow maze to go from the top to the bottom. The entrances of the passages are small.

The *hook* (Fig. 4-hook) environment has a 6-DOF rigid body resembling a hook that must pass through two walls using translations and rotations. The narrow passages in this environment are long due to the shape of the robot.

## 4.2 Experiments

In each environment we defined one initial configuration for the non-biased methods RRT-Expand and EST, and two initial configurations on the extremes of the environments for the biased methods RRT-Connect and RPP. In each independent run, we ran the expansion strategy for a maximum number of iterations and for a number of desired nodes. We collected global and local statistics over all the samples and over the latest bin.

We extracted aggregated metrics for all methods when mapping the walls (Fig. 5) and the maze environments (Fig. 6). We present the ratio of bubbles to growth sites (total and per bin) and the ratio of bubbles to growth sites created during successful expansion (per bin). The total ratio of bubbles to growth sites shows the efficiency of the sampler in discovering new regions of the C-space. The ratios per bin give us a picture of the recent progress made. These metrics show the features of sampling and allow us to compare them among planners.

We show two global metrics for the unbiased sam-

plers RRT-Expand and EST when mapping the hook environment (Fig. 7). We show the percentage of successful growths per bin and the accumulated expansion ratio.

### 4.2.1 Walls Environment

Each of the planners evaluated has a steady decline in their ratio of bubbles to growth sites but at different rates. This corresponds to the speed at which the majority of the C-space was mapped by each planner. RRT-Connect, which covers at 500 nodes, has the highest ratio and the slowest decline. RRT-Expand, which covers at 1000 nodes, starts with a ratio of bubbles to growth sites just below .3 and has a slow decline. EST, which covers at 1400 nodes, starts with a ratio of .2 that quickly declines to about .05 after which it has a slower decline. RPP, which covers at 2300 nodes, starts with a high ratio that quickly declines to 0.2 where it continues at a slower rate.

At the bin level we see that the ratio of bubbles to growth sites and the ratio of bubbles to growth sites from successful expansions follow a similar behavior with spikes that mark drastic changes in coverage. Some of these spikes correlate to the iterations where the planner makes its way through to the following chamber. Other spikes may be due to the quick coverage of non-apparent tightly constrained regions of the C-space. Also note that the non-biased RRT-Expand and EST show the biggest difference between the ratio of bubbles to growth sites and the ratio of bubbles to growth sites from successful expansions which may indicate more time spent on free areas of the space.

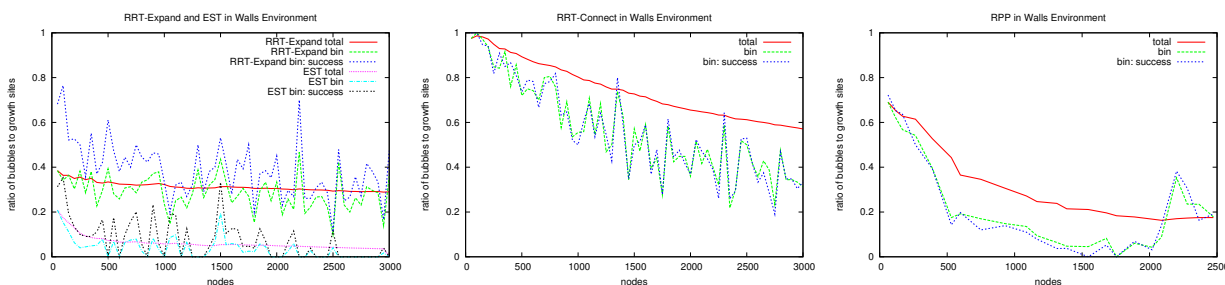


Figure 5: Global metrics for *walls* environment mapped with RRT-Expand (left), EST (left), RRT-Connect (center), and RPP (right). For each sampler we show the total ratio of bubbles to growth sites (total), the bin ratio of bubbles to growth sites (bin), and the bin ratio of bubbles to growth sites from successful expansions (bin: success) with bins of either 50 or 200 nodes

#### 4.2.2 Maze Environment

Similar to the walls environment, in the maze each planner shows a steady decline in their ratio of bubbles to growth sites. This trend held for many other environments we used during our experiments. The maze environment has two free chambers on either side of the actual maze. The mapping of the two free chambers can be seen in the bin ratios of each planner. RRT-Connect, RRT-Expand, and RPP start with a high ratio of bubbles to growth sites computed over bins as they map the first chamber. After negotiating the maze, RRT-Connect and RRT-Expand show a large increase in their ratio of bubbles to growth sites computed over bins as they map the second chamber. RRT-Connect does not show as pronounced of a spike as it maps the second chamber because its ratio of bubbles to growth sites is already high due to its quick progress. RRT-Connect maps the free space in about 700 nodes, RPP maps in about 1800 nodes and RRT-Expand maps in about 5800 nodes. EST failed to make significant progress through the maze.

#### 4.2.3 Hook Environment

For the hook environment, we show the percentage of successful growths per bin and the accumulated expansion ratio for both RRT-Expand and EST. The percentage of successful growths per bin shows how successful the planner is in creating new expansions at different slices of time. This metric alone is not enough to show how “good” the newly created sam-

ples are. The accumulated expansion ratio helps to show the planners progress in creating high quality samples. An increase in the accumulated expansion ratio means that new samples have been added with high expansion ratios.

RRT-Expand shows a steady increase of successful growths per bin as the planner progresses. Also, the accumulated expansion ratio for RRT-Expand increases steadily showing that the planner is making progress mapping new areas. On the other hand, EST’s percentage of successful growths per bin continuously stays between 70% and 100% and when examining the accumulated expansion ratio we can see the reason. EST’s expansion ratio grows at a much slower rate than that of RRT-Expand, it is mostly creating new samples in previously mapped areas and it is not successful in finding samples with high expansion ratios.

## 5 Conclusions

In this paper we propose several qualitative metrics for use in the analysis, comparison, and improvement of incremental expansion planners. Building on previous work [19], we show that most samples from an incremental planner will be either a cc-oversample or cc-expand, and propose aggregating statistics both local and global with every expansion attempt. At the local level we collect node statistics and group them in bubbles rooted at growth sites to shed light on the complexity of a region. Tightly constrained

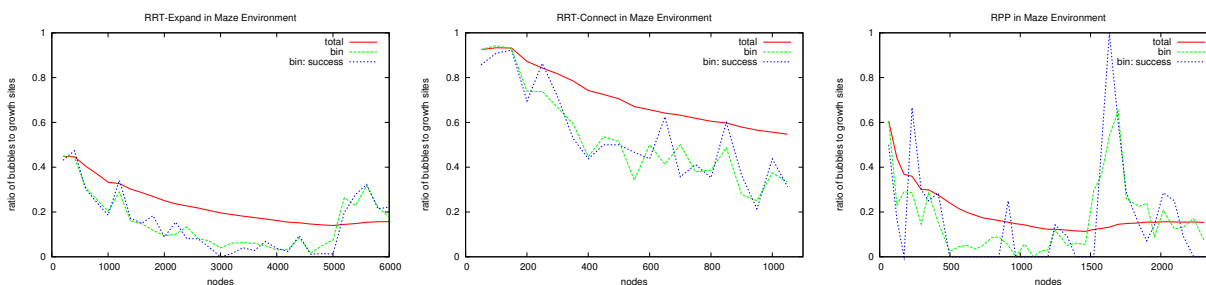


Figure 6: Global metrics for *maze* environment mapped with RRT-Expand (left), RRT-Connect (center), and RPP (right). For each sampler we show the total ratio of bubbles to growth sites (total), the bin ratio of bubbles to growth sites (bin), and the bin ratio of bubbles to growth sites from successful expansions (bin: success) with bins of either 200 nodes

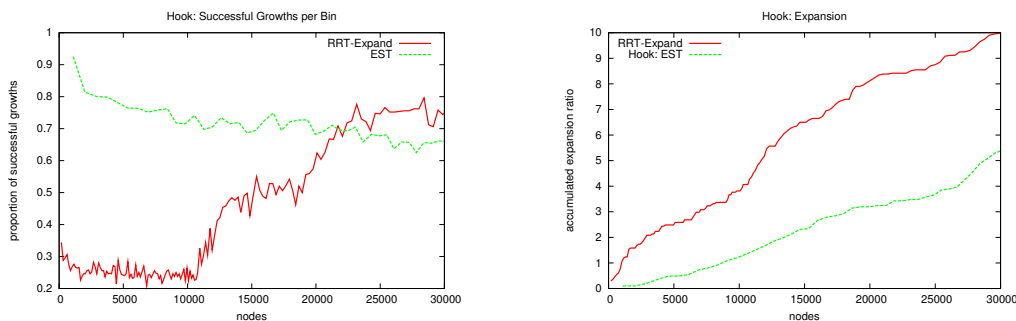


Figure 7: Results for *hook* environment mapped with RRT-Expand and EST. We show the proportion of successful growths per bin and the accumulated expansion ratio with bins of 50 nodes

regions can be identified for their obstruction in the sampling process. At the global level we collect statistics from nodes and bubbles that monitor the modeling process to provide insight into how well the planner is progressing – if the planner is making poor progress it can also be seen with these metrics. We used the metrics proposed to evaluate the performance of four incremental planners in three different environments with narrow passages. Our metrics helped to explain the relative performance of the methods studied.

## References

[1] N. M. Amato, O. B. Bayazit, L. K. Dale, C. V. Jones, and D. Vallejo. OBPRM: An obstacle-based PRM for 3D workspaces. In *Robotics: The*

*Algorithmic Perspective*, pages 155–168, Natick, MA, 1998. A.K. Peters. Proc. Third Workshop on Algorithmic Foundations of Robotics (WAFR), Houston, TX, 1998.

- [2] N. M. Amato and G. Song. Using motion planning to study protein folding pathways. *J. Comput. Biol.*, 9(2):149–168, 2002. Special issue of Int. Conf. Comput. Molecular Biology (RECOMB) 2001.
- [3] J. Barraquand and J. C. Latombe. Robot motion planning: A distributed representation approach. *Int. J. Robot. Res.*, 10(6):628–649, 1991.
- [4] O. B. Bayazit, G. Song, and N. M. Amato. Enhancing randomized motion planners: Exploring with haptic hints. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 529–536, 2000.

- [5] O. B. Bayazit, G. Song, and N. M. Amato. Ligand binding with OBPRM and haptic user input: Enhancing automatic motion planning with virtual touch. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 954–959, 2001. This work was also presented as a poster at *RECOMB 2001*.
- [6] P. Bessiere, J. M. Ahuactzin, E. G. Talbi, and E. Mazer. The Ariadne’s clew algorithm: Global planning with local methods. In *Proc. IEEE Int. Conf. Intel. Rob. Syst. (IROS)*, volume 2, pages 1373–1380, 1993.
- [7] V. Boor, M. H. Overmars, and A. F. van der Stappen. The Gaussian sampling strategy for probabilistic roadmap planners. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, volume 2, pages 1018–1023, May 1999.
- [8] H. Chang and T. Y. Li. Assembly maintainability study with motion planning. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 1012–1019, 1995.
- [9] R. Geraerts and M. H. Overmars. Reachability analysis of sampling based planners. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 406–412, 2005.
- [10] D. Hsu, R. Kindel, J. C. Latombe, and S. Rock. Randomized kinodynamic motion planning with moving obstacles. In *Proc. Int. Workshop on Algorithmic Foundations of Robotics (WAFR)*, pages SA1–SA18, 2000.
- [11] D. Hsu, G. Sánchez-Ante, and Z. Sun. Hybrid PRM sampling with a cost-sensitive adaptive strategy. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 3885–3891, 2005.
- [12] L. E. Kavraki, P. Svestka, J. C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robot. Automat.*, 12(4):566–580, August 1996.
- [13] J. J. Kuffner and S. M. LaValle. RRT-Connect: An Efficient Approach to Single-Query Path Planning. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 995–1001, 2000.
- [14] J.-P. Laumond and T. Siméon. Notes on visibility roadmaps and path planning. In *Proc. Int. Workshop on Algorithmic Foundations of Robotics (WAFR)*, 2000.
- [15] S. M. LaValle and J. J. Kuffner. Randomized kinodynamic planning. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 473–479, 1999.
- [16] S. M. LaValle and J. J. Kuffner. Rapidly-Exploring Random Trees: Progress and Prospects. In *Proc. Int. Workshop on Algorithmic Foundations of Robotics (WAFR)*, pages SA45–SA59, 2000.
- [17] J.-M. Lien, O. B. Bayazit, R.-T. Sowell, S. Rodriguez, and N. M. Amato. Shepherding behaviors. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 4159–4164, April 2004.
- [18] J.-M. Lien, S. Rodriguez, J.-P. Malric, and N. M. Amato. Shepherding behaviors with multiple shepherds. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 3413–3418, April 2005.
- [19] M. A. Morales A., R. Pearce, and N. M. Amato. Metrics for analyzing the evolution of C-Space models. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 1268–1273, May 2006.
- [20] A. Singh, J. Latombe, and D. Brutlag. A motion planning approach to flexible ligand binding. In *7th Int. Conf. on Intelligent Systems for Molecular Biology (ISMB)*, pages 252–261, 1999.
- [21] X. Tang, B. Kirkpatrick, S. Thomas, G. Song, and N. M. Amato. Using motion planning to study RNA folding kinetics. In *Proc. Int. Conf. Comput. Molecular Biology (RECOMB)*, pages 252–261, 2004.
- [22] S. Thomas, X. Tang, L. Tapia, and N. M. Amato. Simulating protein motions with rigidity analysis. In *Proc. Int. Conf. Comput. Molecular Biology (RECOMB)*, pages 394–409, 2006.
- [23] S. A. Wilmarth, N. M. Amato, and P. F. Stiller. MAPRM: A probabilistic roadmap planner with sampling on the medial axis of the free space. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, volume 2, pages 1024–1031, 1999.

- [24] D. Xie, M. Morales, R. Pearce, S. Thomas, J.-M. Lien, and N. M. Amato. Incremental map generation (IMG). In *Proc. Int. Workshop on Algorithmic Foundations of Robotics (WAFR)*, July 2006.