

Customizing PRM Roadmaps at Query Time*

Guang Song Shawna Miller Nancy M. Amato
Department of Computer Science
Texas A&M University
College Station, TX 77843-3112
{gsong, slm5563, amato}@cs.tamu.edu

Abstract

In this paper, we propose a new approach for building and querying probabilistic roadmaps. In the roadmap construction stage, we build coarse roadmaps by performing only an approximate validation of the roadmap nodes and/or edges. In the query stage, the roadmap is validated and refined only in the area of interest for the query, and moreover, is customized in accordance with any specified query preferences. This approach, which postpones some of the validation checks (e.g., collision checks) to the query phase, yields more efficient solutions to many problems. An important benefit of our approach is that it gives one the ability to customize the same roadmap in accordance with multiple, variable, query preferences. For example, our approach enables one to find a path which maintains a particular clearance, or makes at most some specified number of sharp turns. Our preliminary results on problems drawn from diverse application domains show that this new approach dramatically improves performance, and shows remarkable flexibility when adapting to different query requirements.

1 Introduction

Automatic motion planning deals with finding a feasible sequence of motions to take some moving object (the ‘robot’) from a given initial configuration to some specified goal configuration. While complete motion planning algorithms do exist, they are rarely used in practice since they are computationally infeasible in all but the simplest cases [14]. For this reason, attention has focused on probabilistic methods, which sacrifice completeness in favor of computational feasibility and applicability.

In particular, a class of algorithms known as *probabilistic roadmap* (PRM) methods have been the subject of much recent work. These successes sparked a flurry of activity in which PRM motion planning techniques were applied to a number of challenging problems arising in a variety of fields including robotics (e.g., closed-chain systems

[11, 16]), CAD (e.g., maintainability [3, 8], deformable objects [2, 13]), and even computational Biology and Chemistry (e.g., ligand docking [4, 18], protein folding [19, 20]). Indeed, it can be argued that the PRM framework was instrumental in this broadening of the range of applicability of motion planning, as many of these problems had never before been considered candidates for automatic methods.

At the same time, these new applications served to point out some weaknesses in the PRM approach. In particular, it was observed that there were some classes of problems for which PRMs did not yield efficient solutions. For example, problems where the solution path must traverse *narrow passages* in C-space [12]. As a result, a number of PRM variants specifically targeted at this problem have been proposed, e.g., [1, 7, 12]. Also, some novel approaches for improving PRM efficiency have shown promising results [6, 17] (these methods do not address the narrow passage problem). Unfortunately, however, PRM solutions to many problems still take prohibitively long.

Another shortcoming of PRMs is that while they are very good at finding *a* path, they do not support applications which might impose particular, variable, requirements on the sought after path, e.g., maintaining a particular clearance from obstacles or minimizing the robot’s rotation as it moves. This issue has not received as much attention as, e.g., the narrow passage problem, because simply finding any path was considered a necessary first step in many cases. However, this issue must be addressed before PRMs can be used in many practical instances.

In this work, we propose a new approach for *building and querying* probabilistic roadmaps that addresses both the above shortcomings of current PRMs.

- In the roadmap construction stage, we build coarse roadmaps by performing only approximate validation of roadmap nodes and edges.
- In the query stage, the roadmap is validated and refined only in the area of interest for the query, and moreover, is *customized* in accordance with any specified query preferences.

This approach, which postpones some of the validation checks (e.g., collision checks) to the query phase, yields more efficient solutions to many problems. For example,

*This research supported in part by NSF CAREER Award CCR-9624315 (with REU Supplement), NSF Grants IIS-9619850 (with REU Supplement), ACI-9872126, EIA-9975018, EIA-9805823, and EIA-9810937, and by the Texas Higher Education Coordinating Board under grant ARP-036327-017.

consider a roadmap with n nodes, each of which is connected to its k nearest neighbors, who are at an average distance of d from it. Constructing this roadmap in its entirety would require $O(nkd)$ validity checks – the majority of which are not needed for any particular solution path. Our strategy of postponing validation leads to faster roadmap construction, while incurring only slightly higher query times. As discussed below, similar strategies have been used by other researchers [6, 17].

An important benefit of our approach is that it gives one the ability to *customize* the roadmap in accordance with any specified, variable, query requirements. For example, it might be specified that the solution path should maintain a certain clearance from the obstacles, that the robot’s rotation should be minimized, or that the path should contain at most some fixed number of bends. And moreover, these requirements might change for different queries in the same roadmap, or might actually be preferences (as opposed to requirements) with different priorities. Using the same roadmap, one might want several different paths: a path with a particular clearance and minimal rotation, a path with a desired clearance from obstacles, and another path with a different clearance threshold. Current PRMs do not support such customizable use of the *same* roadmap.

2 Related Work

In this section we briefly mention other approaches targeted at either improving the efficiency of PRMs or in supporting requirements on the query path other than the usual collision-free requirement. First, there has been some recent work aimed at improving the efficiency of PRMs by postponing most of the validation to the query stage [6, 17].

Like PRM, Lazy PRM [6] randomly generates nodes in the configuration space, and edges in the roadmap represent paths between roadmap nodes. Unlike PRM, all nodes and edges are initially assumed to be collision-free. This dramatically reduces the amount of time spent in the preprocessing phase. When presented with a query, a shortest path is found between the start and goal configurations. Then, the nodes (first) and edges (second) along the path are checked for collision. If a node or edge is found to be in collision, it is removed from the roadmap, and a new shortest path is extracted from the roadmap. This process repeats until a collision-free path is found or a path no longer exists between the start and goal.

A similar approach called Fuzzy PRM was independently proposed in [17]. Like Lazy PRM, this method does not check the validity of the roadmap edges during construction, but it does ensure that the roadmap nodes are collision-free. Fuzzy PRM assigns each edge a weight that corresponds to the probability that the edge represents a feasible path (1 for collision-free). During the query, the algorithm selects the path with highest probability. Each edge along the path is inserted into a priority queue. Then, the edge with the lowest probability is checked at a slightly higher

resolution, assigned a new probability, and reinserted in the priority queue. This process repeats until all edges on the path have probabilities of 1 or a collision is found.

We are aware of no PRM work which supports variable query preferences. However, there are a few probabilistic methods that enable one to specify, but not subsequently alter, query requirements [5, 15]. These methods are a bit different from PRMs in that they grow components rooted at the start (and possibly the goal) configuration. Since they incrementally construct paths, query requirements can be implemented as the roadmap is built. However, they do not provide support for later querying the same roadmap with different requirements.

3 Building and Querying Customizable PRM Roadmaps

In this section, we describe our new algorithm for the fast construction of PRM roadmaps that support multiple, and variable, query preferences. The general structure of our Customizable PRM (or C-PRM) follows the traditional PRM paradigm. The main differences are that in C-PRM, only a coarse, approximate roadmap is built during preprocessing, and detailed validation of the path, including customization for any particular query preferences, is performed during the query itself. That is, we postpone much of the time consuming validation until the query phase, and then only perform detailed validation on the portions of the roadmap deemed necessary for solving the particular query. In addition to yielding big performance gains, this strategy enables us to support additional, and variable, requirements or preferences on the properties of the solution path obtained. We note that some query preferences cannot be encoded in a roadmap, i.e., they must be enforced at query time. For example, if the same roadmap will be used to answer queries with different clearance thresholds, or if validity can only be tested after the start and goal are known (e.g., paths with at most k sharp turns).

We now describe in more detail how we construct coarse roadmaps during preprocessing (Section 3.1) and how these roadmaps are used to provide solution paths adhering to requirements or preferences specified during the query phase (Section 3.2).

3.1 Roadmap construction using fast, approximate evaluation

One of the strengths of the PRM method is that the global sampling (and validation) yields roadmaps that encode useful information about the connectivity of the free configuration space. Unfortunately, however, as previously noted, complete, exact evaluation of configurations in the roadmap construction phase leads to infeasible construction times for many important applications. Our strategy is to replace the exact validation, of nodes and/or edges, with faster, approximate methods. The goal is to perform ‘good enough’ eval-

uations so that the resulting roadmap provides some global view of the connectivity of the planning space, while postponing as much of the time consuming detailed validation until it is really needed.

While our strategy is similar in motivation to the recently proposed Lazy PRM [6] and Fuzzy PRM [17] methods, it is different in that we advocate approximate validation methods for the roadmap nodes *and* edges in the roadmap construction phase. In contrast, Lazy PRM performs no validation of either nodes or edges during roadmap construction, and Fuzzy PRM performs exact evaluation of nodes but no validation of edges (initial edge probabilities are set based on inter-node distances). We believe our approximate edge evaluation will produce roadmaps which better reflect the planning space, which is very important for difficult problems, such as those containing *narrow passages* in C-space [12]. In addition, as our approximate evaluation methods can be tuned to be more or less accurate (and consequently more or less time consuming), our methodology allows the user to determine the degree of exactness appropriate for the given application.

Below, we describe some methods for approximate node and approximate edge evaluation.

3.1.1 Approximate node evaluation

Using faster, coarser, collision detection. Suppose that paths with variable clearances will be desired when planning in a particular environment. Note that it may not be known *a priori* what the maximum possible clearance is, or what clearance requirements will be imposed in the query.

One approach is to determine the clearance of all nodes and edges during the roadmap construction phase, and then to use a standard graph search to query the roadmap for the solution path with maximal clearance. This requires a collision detection method which computes clearance information, such as the C-Space Toolkit (CSTk) [21].

Another strategy is to build a roadmap which does not contain clearance information. Such a roadmap would be faster to construct and could employ any collision detection method suitable for the environment (thus allowing the use of perhaps faster packages such as RAPID [10]). Then, at the query stage, the roadmap could be refined using more expensive clearance computations. This can be done by (iteratively) extracting a shortest path, and then checking it for compliance with the given clearance requirement – first the nodes, and then the edges. Performance is improved because clearance computations for many edges will not be performed, and in particular, we will only check edges that connect nodes with acceptable clearances. An intermediate approach might compute clearance information for the nodes, but not the edges, during roadmap construction.

Grid-based approximation. Another method is to compute an approximate value for each cell in a grid-based decomposition of C-space. During roadmap construction, the value for each configuration is determined by a simple, fast table lookup. Of course, due to the high-dimensionality

of most interesting C-spaces, this method is infeasible in many cases. Nevertheless, there exist some applications where useful approximations can be computed on a coarse grid. For example, as discussed in Section 4, this technique has been used for ligand binding, which arises in drug design [4, 18].

3.1.2 Approximate edge evaluation

As with approximate node evaluation, our goal here is to perform some fast, approximate validation of the edges. The motivation is that many invalid edges can be easily discovered, resulting in roadmaps that better reflect the connectivity of the free C-space.

Binary resolution approximation. This technique is very general and can be applied in any scenario. The strategy is to validate the intermediate configurations on an edge (a straight-line in C-space) according to a binary partitioning strategy. That is, we first validate the midpoint, then the midpoints of the two resulting subsegments, etc. Our experience shows this strategy discovers many invalid edges quickly. The resolution at which to terminate the process during roadmap construction can be some computed value related to the environment characteristics, or can be set by the user. A similar idea of increasing resolution checks is used in Fuzzy PRM [17] when attempting to increase edge probabilities during the query phase.

C-space clearance approximation. Another way to estimate whether an edge is collision-free is to use a test based on the C-space clearance of its two endpoints. The C-space clearance c , if known, of a node defines a sphere centered at that node of radius c which is entirely contained in the free C-space. If the clearance spheres of the two endpoints of an edge intersect, then there exists a collision-free path connecting them. (A path could still exist even if the spheres do not intersect.) Since we cannot compute clearances in C-space, we compute an approximation by selecting several directions at random, finding the clearance in those directions, and using the minimum as the C-space clearance approximation. Clearly, the accuracy of the approximation is very sensitive to the number k . In practice, the C-space clearance approximation seems to work well for values as small as $k = 3$.

3.2 Query Phase

In C-PRM queries are processed by adding the start and goal to the roadmap, and connecting them to their k closest neighbors. The shortest path between them is then searched using Dijkstra's algorithm. Since the path consists of edges and nodes that have only been approximately validated, finer resolution collision checking must now be performed on all path segments. In addition, any desired query specifications will be enforced at this stage. If any node or edge on the path fails to meet the query requirements, it is removed from the roadmap, a new shortest path is found in

Running time for C-PRM vs. PRM						
Application	# Rdmap Edges		Rdmp Const. Time (s)		Query Time (s)	
	C-PRM	PRM	C-PRM	PRM	C-PRM	PRM
Moving-Piano	2133	2019	24	104	1.73	1.38
CAD: alpha1.5	508	441	57	478	3.77	3.07
Ligand Binding	705	958	206	9523	41.16	38.54
Protein Folding	127296	96919	6.19hr	13.47hr	27.1min	11.4min

Table 1: Running times of C-PRM and a traditional PRM for various applications.

the refined roadmap, and the process iterates until a valid path is found or failure is reported.

3.2.1 Additional, multiple query preferences

Usually, there are other desirable properties for a path in addition to the basic requirement that it be collision-free. Some common preferences include large clearance, small rotation, low curvature (smoothness), few sharp corners, avoiding singularities for manipulators, or low potential energies for ligand binding and protein folding (see Table 2).

Common Query Preferences	
Application	Preference
CAD	clearance
Mobile robots	clearance, smoothness
Manipulators	singularity, clearance
Ligand binding	low potential
Protein folding	low potential

Table 2: Some common query preferences for different applications.

Note that it is not practical to store all information needed to verify such properties in the roadmap. Moreover, in some cases it is difficult, or even impossible, to associate such information with roadmap nodes or edges because it relates to global properties of the solution path. For example, one might ask for a path for a mobile robot that makes at most k sharp turns. Since this is not a local property, and since the start and goal are not known in advance, one cannot prune the roadmap so that it contains only valid solution paths. In this example, and indeed in many interesting queries, the only feasible way to enforce such preferences is during the query process.

In many cases, one would actually like to request a path meeting several query preferences but with different priorities. For example, one might ask for a path that: (1) is collision free, (2) restricts certain dofs to a particular range (e.g., to control tilt), (3) has some particular clearance, and (4) requires at most some specified amount of rotation. The numbers indicate priorities.

C-PRM can support such queries by iteratively refining, or customizing, the roadmap. First, a collision-free path is found. Next, the path is checked to see if all dofs are

within the specified limits. Then, clearances on the path are checked, and finally the rotation along the path is checked. At each stage, invalid nodes and edges are removed as they are discovered, and the process iterates. Finally, a path that meets the preferences to the degree possible is found, and in the process, the roadmap is customized with respect to these requirements.

4 Examples and Results for Various Applications

In this section, we consider examples drawn from a wide variety of application domains, including traditional motion planning applications, intelligent CAD, and computational Biology and Chemistry. We demonstrate how C-PRM can be applied to support different query preferences and that in all cases C-PRM shows significant performance gains over a traditional PRM, as shown in Table 1. All experiments were performed on a Pentium III 550 MHz PC using our C++ OBPRM library, which includes implementations of many PRM variants.




Figure 1: The piano mover's problem.

4.1 Robotics Applications

The Piano Mover's Problem. Figure 1 shows a house containing a piano and a table. The goal is to move the piano out of the house. For this problem, all roadmap nodes were checked for collision, but the edges were only approximately verified using the binary resolution approximation

scheme, where we checked the edges down to a resolution of 10 times the resolution required for exact validation. The results in Table 1 show that the C-PRM roadmap was constructed approximately five times faster than the PRM roadmap, while the cost paid in terms of a more expensive query was negligible (a fraction of a second). One can also see by comparing the number of edges in the C-PRM and PRM roadmaps that the binary approximation does fairly well in pruning invalid edges.

Iteratively Refined Roadmap				
Query	time (sec)	Roadmap Statistics		
		#N	#E	#CCs
makemap	468.00	2618	33914	3(2087:531:1)
C-free	14.64	2620	34081	3(2089:531:1)
tilt < 90°	67.36	2072	21562	4(1665:206:200:1)
rot < 15°/s	51.47	2072	21546	4(1665:206:200:1)
cl > 0.1	25.00	1059	7231	3(885:97:77)
cl > 0.2	23.00	767	4391	4(647:67:52:1)

Table 3: The evolution of the roadmap for the piano mover’s problem based on prioritized query preferences. The clearance requirement units are environment units. The rotation requirement units are degrees per step.

We also use this problem to illustrate the ability of C-PRM to support multiple, prioritized query preferences. The results of this process are summarized in Table 3. The first query, which enforces the collision-free requirement, results in the addition of the start and the goal to the roadmap, along with some edges. The tilting and clearance requirements are properties that can be tested locally for each node and edge, and as such, they result in the loss of both nodes and edges from the roadmap. For the tilting and clearance preferences, we first go through the roadmap and remove all nodes that do not satisfy the current requirement. This turned out to be useful since many of our nodes did not meet the requirements, and so pruning them initially greatly reduced the number of iterations required to find a feasible path. The rotation requirement does not result in the loss of any nodes, but does remove some edges from the roadmap – this is because it is a property related to the edges, not nodes.




Figure 2: A 3-R manipulator

Manipulators. Another practical application of our method involves multi-linked manipulators. Consider the




Figure 3: (a) Alpha Puzzle. (b) Ligand Binding.

example of an RRR planar manipulator in Figure 2, where the singularity regions are shown in the plot on the right. It is important to avoid singular configurations when planning because in those configurations the manipulator loses one or more degrees of freedom [9]. In the figure, path 1 is the shortest path from the start to the goal; it is undesirable because it passes very close to a singularity region. If the user were to add an additional query requirement specifying some minimum value on the absolute value of the determinant of the Jacobian, then path 2 or path 3 might be returned, depending on the particular minimum value specified.

4.2 Intelligent CAD

An important design criteria for complex mechanical systems is that certain parts should be easy to service and/or replace. Motion planning algorithms for testing such conditions automatically on CAD models would be a great aid to engineers [3, 8].

The CAD model we consider is the alpha puzzle, which consists of two identical twisted, intertwined tubes (one the obstacle and one the robot, 1008 triangles each); the objective of the puzzle is to separate the tubes (see Figure 3 (a)). An easier version of problem was obtained by scaling the obstacle tube by a factor of 1.5 in one dimension, which widened the gap between the prongs of the alpha shape.

For this problem, collision-free configurations were generated and edges were partially validated by the binary approximation scheme (again by checking at ten times the resolution required for exact validation). The results in Table 1 show that C-PRM solves this problem nearly 10 times faster than the PRM. The savings are due to the fact that C-PRM never completely validates many of the edges.

4.3 Computational Biology/Chemistry

We have also applied C-PRM to several problems arising in computational Biology and Chemistry such as ligand binding and protein folding. For these applications, different criteria are used to judge the validity of nodes and edges. In particular, instead of geometric collision detection tests, we need to determine if the *potential energy* of the configurations (called conformations) is low enough.

Ligand binding, also referred to as molecule docking, is an important problem in the design of new drugs [4, 18]. We applied C-PRM to the example shown in Figure 3 (b). The protein (M4 Lactate Dehydrogenase Ternary Complex, pdb file: 1LDM) is shown in wire-frame and the binding ligand, shown in space-fill, is an oxamate molecule. The roadmap was constructed using a grid-based potential approximation for the nodes (as in [4, 18]) and ‘exact’ validation for the edges, but again using the grid-based potential approximation for the internal nodes on the edge. During the query, exact potential computations were used. The results in Table 1 show C-PRM is nearly 50 times faster than the traditional PRM method. Moreover, even though C-PRM has postponed all exact potential computations to the query phase, the query times for the two approaches are still quite comparable. This example shows the grid-based approximation for node validation is a very powerful technique for this application.

We have also applied C-PRM to compute folding pathways for several small proteins (see [19, 20] for more detail). The roadmap for protein A (120 dof) was constructed using exact potential computations for the nodes and the binary approximation scheme for the edges which used a resolution six times that required for exact evaluation. Here, C-PRM is roughly twice as fast as the PRM approach. Less improvement is obtained here than was observed for the ligand binding because C-PRM mainly optimizes the roadmap connection phase, and this application spends more of its time in the node generation phase than the other applications studied do.

5 Conclusion

In this paper, we describe a new approach for *building and querying* probabilistic roadmaps. This approach, which postpones some of the validation checks (e.g., collision checks) to the query phase, was seen to yield more efficient solutions to a wide variety of problems. An important benefit of our approach is that it gives one the ability to *customize* the roadmap in accordance with multiple, variable, query preferences. This feature addresses an aspect that had been absent in previous PRMs, and potentially makes the PRM framework more attractive for use in many practical instances.

Acknowledgements

We would like to thank the robotics group at Texas A&M. The alpha puzzle was designed by Boris Yamrom of the Computer Graphics & Systems Group at GE’s Corporate Research & Development Center.

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 *Proc. Int. Workshop on Algo-*

Algorithmic Foundations of Robotics (WAFR), pages 155–168, 1998.

[2] E. Anshelevich, S. Owens, F. Lamiroux, and L. Kavraki. Deformable volumes in path planning applications. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2000.

[3] 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.

[4] 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)*, 2001. To appear. This work will be presented as a poster at RECOMB’01.

[5] 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.

[6] R. Bohlin and L. E. Kavraki. Path planning using lazy prm. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 521–528, 2000.

[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)*, pages 1018–1023, 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] John J. Craig. *Introduction to Robotics: Mechanics and Control, 2nd Edition*. Addison-Wesley Publishing Company, Reading, MA, 1989.

[10] S. Gottschalk, M.C. Lin, and D. Manocha. Obb-tree: A hierarchical structure for rapid interference detection. Technical Report TR96-013, University of N. Carolina, Chapel Hill, CA, 1996.

[11] L. Han and N. M. Amato. A kinematics-based probabilistic roadmap method for closed chain systems. In *Proc. Int. Workshop on Algorithmic Foundations of Robotics (WAFR)*, 2000.

[12] D. Hsu, L. Kavraki, J.-C. Latombe, R. Motwani, and S. Sorkin. On finding narrow passages with probabilistic roadmap planners. In *Proc. Int. Workshop on Algorithmic Foundations of Robotics (WAFR)*, 1998.

[13] L. Kavraki, F. Lamiroux, and C. Holleman. Towards planning for elastic objects. In *Proc. Int. Workshop on Algorithmic Foundations of Robotics (WAFR)*, 1998.

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

[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, J.H. Yakey, and L.E. Kavraki. A probabilistic roadmap approach for systems with closed kinematic chains. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 1999.

[17] C. L. Nielsen and L. E. Kavraki. A two level fuzzy prm for manipulation planning. Technical Report TR2000-365, Computer Science, Rice University, Houston, TX, 2000.

[18] A.P. Singh, J.C. Latombe, and D.L. Brutlag. A motion planning approach to flexible ligand binding. In *7th Int. Conf. on Intelligent Systems for Molecular Biology (ISMB)*, pages 252–261, 1999.

[19] G. Song and N. M. Amato. A motion planning approach to folding: From paper craft to protein folding. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2001. To appear.

[20] G. Song and N. M. Amato. Using motion planning to study protein folding pathways. In *Proc. Int. Conf. Comput. Molecular Biology (RECOMB)*, pages 278–287, 2001.

[21] P. G. Xavier and R. A. LaFarge. A configuration space toolkit for automated spatial reasoning: Technical results and ldrd project final report. Technical Report SAND97-0366, Sandia National Laboratories, 1997.