Path Planning for Spatial Closed Kinematic Chains with Spherical Joints *

J.C. Trinkle
Intelligent Systems Principles Department
Sandia National Laboratory
Albuquerque, NM 87185-1004 U.S.A.

R. James Milgram
Department of Mathematics
Stanford University
Stanford, CA 94305 U.S.A.

Given a robot in a workspace with obstacles and start and goal configurations, q_S and q_G, the "generalized movers' problem" is to construct a continuous collision-free path for the robot connecting q_S and q_G. In it's full generality, this problem is extremely challenging. The most efficient, complete algorithm for solving the general problem is Canny's roadmap algorithm which operates in the space of all configurations of the system, C-space, denoted by $\mathcal{C}$. In $\mathcal{C}$, points represent specific robot poses and (continuous) curves represent robot motions. Canny's roadmap algorithm produces a one-dimensional skeleton, $\mathcal{R}$ (a roadmap), of $\mathcal{C}$ with two properties that lead to algorithmic completeness: (1) for each component of $\mathcal{C}$, $\mathcal{R}$ has exactly one component and (2) a path connecting any point in a component of $\mathcal{C}$ to the corresponding component of $\mathcal{R}$ can always be generated. The algorithm runs in single exponential time, with the exponent equal to the number of degrees of freedom. Because the potential number of components in $\mathcal{C}$ is at least exponential in the dimension of $\mathcal{C}$, the complexity of the roadmap algorithm is worst-case optimal.

The complexity of complete motion planning algorithms for the generalized movers' problem have fueled two major thrusts in motion planning research over the past 20 years: the search for sub-classes of the general movers' problems for which complete polynomial-time algorithms exist, and approximate methods which trade completeness for average-case efficiency. Other than roadmap methods, the main class of complete planning methods are known as cell decomposition methods. In these methods, $\mathcal{C}$ is decomposed into a set of non-overlapping cells and their connectivity is represented in a graph. After identifying the cells containing q_S and q_G, one performs a graph search to obtain a sequence of cells connecting q_S and q_G; the union of the cells in this sequence forms a "channel." The last step is to extract a continuous path through the channel. The problem with cell decomposition methods is that for the generalized movers' problem, there are no polynomial-time decomposition algorithms. Polynomial-time algorithms for simplified problems have been developed, but work in this area dwindled as successful probabilistic methods began to appear.

Probabilistic roadmap methods (PRMs) have taken center stage for about the last 10 years. Roughly speaking, PRMs use random search to construct a probabilistic roadmap, $\mathcal{R}$, as a graph whose nodes and arcs represent free configurations and collision-free paths linking them. If enough random samples are generated, the components of the graph will be in one-to-one correspondence with the components of $\mathcal{C}$ and it will be "easy" to connect arbitrary start and goal configurations to the graph. Despite their incompleteness, PRMs have become popular because they have successfully solved many problems in very high-dimensional C-spaces. They have also been adapted to problems with various types of physical constraints. For example, PRMs have been applied to problems with continuously deformable bodies, problems with significant dynamic effects, problems involving mechanisms with loops, and problems of dexterous manipulation, where the kinematic loop topology varies and further complications arise from contact mechanics.

While the successes of PRMs are clear, one should expect that as these methods are pushed into domains with complicated constraints, such as those with closed kinematic loops, it will become increasingly difficult to construct randomized roadmaps. In the case of dexterous manipulation, plans are composed of path segments restricted to strata of differing dimensions. The application of PRMs would require one to generate roadmaps in each stratum and then to connect them. In general, the number of strata will be exponential in the dimension of the ambient C-space. Given this, one should expect the amount of random sampling required for building a good probabilistic roadmap to be exponential in the dimension of C-space.

Our approach represents a return to the first research thrust discussed above; the search for complete, polynomial-time algorithms. We study the problem of planning reconfigurations of spatial kinematic closed chains with spherical joints. Our problem is further restricted by neglecting collisions with obstacles and with other links. While our class of problems is too simple to be directly applicable to most tasks of practical interest in robotics, our approach sheds new light on exact approaches to path planning. Our approach is based on a complete understanding of the singular sets of certain maps, which is achieved via techniques of modern mathematics not previously exploited. What makes our class of planning problems difficult even though collisions are ignored is the complexity of the valid portion of C-space, which is a real algebraic variety of co-dimension three that is not necessarily parameterizable (see http://www.cs.sandia.gov/∼jctrink/closed_chain_mp.html for animations of plans).

DOD Applications and Future Work: Path planning is relevant in many DOD application areas including urban warfare, facilities security, bio-warfare, and intelligence gathering. In future work, we plan to push our mathematical methods toward more realistic problems that demand collision avoidance. We are also applying our methods to protein folding and structure prediction problems. Solutions to these problems will facilitate the design of anti-toxins for thwarting biological attacks.

*This work has been supported by Sandia National Laboratories and the National Science Foundation.