J.C. Trinkle

Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180, USA
trink@cs.rpi.edu

R. James Milgram

Department of Mathematics
Stanford University
Stanford, CA 94305, USA
milgram@math.stanford.edu

Complete Path Planning for Closed Kinematic Chains with Spherical Joints

Abstract

We study the path planning problem, without obstacles, for closed kinematic chains with n links connected by spherical joints in space or revolute joints in the plane. The configuration space of such systems is a real algebraic variety whose structure is fully determined using techniques from algebraic geometry and differential topology. This structure is then exploited to design a complete path planning algorithm that produces a sequence of compliant moves, each of which monotonically increases the number of links in their goal configurations. The average running time of this algorithm is proportional to n^3. *While less efficient than the* O(n) *algorithm of Lenhart and Whitesides, our algorithm produces paths that are considerably smoother. More importantly, our analysis serves as a demonstration of how to apply advanced mathematical techniques to path planning problems.*

Theoretically, our results can be extended to produce collision-free paths, paths avoiding both link–obstacle and link–link collisions. An approach to such an extension is sketched in Section 4.5, but the details are beyond the scope of this paper. Practically, link–obstacle collision avoidance will impact the complexity of our algorithm, forcing us to allow only small numbers of obstacles with "nice" geometry, such as spheres. Link–link collision avoidance appears to be considerably more complex. Despite these concerns, the global structural information obtained in this paper is fundamental to closed kinematic chains with spherical joints and can easily be incorporated into probabilistic planning algorithms that plan collision-free motions. This is also described in Section 4.5.

KEY WORDS—exact path planning, complete algorithm, critical points, differential topology, algebraic geometry, gradient flow

The International Journal of Robotics Research
Vol. 21, No. 9, September 2002, pp. 773-789,
©2002 Sage Publications

1. Introduction

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 its full generality with multiple robots and arbitrary link and obstacle geometries, this problem is extremely challenging. The most efficient, complete algorithm for solving the general problem is the roadmap algorithm of Canny (1988) 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 assumes polyhedral bodies and takes as input a formula describing the collision-free portion of configuration space, $\mathcal{C}_{free}$. It produces a one-dimensional skeleton, $\mathcal{R}_{free}$ (a roadmap), with two properties that lead to algorithmic completeness: (1) for each component of $\mathcal{C}_{free}$, $\mathcal{R}_{free}$ has exactly one component; and (2) a path connecting any point in a component of $\mathcal{C}_{free}$ to the corresponding component of $\mathcal{R}_{free}$ 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 $\mathcal{C}_{free}$ is exponential in the dimension of $\mathcal{C}$, the complexity of the roadmap algorithm is worst-case optimal (Canny 2002).

The complexity and implementation difficulties of complete motion planning algorithms for the generalized movers' problem (Canny 1988; Schwartz, Hopcroft, and Sharir 1987) have fueled two major thrusts in motion planning research over the past 20 years: the search for sub-classes of the generalized movers' problem for which complete polynomial-time algorithms exist, and approximate methods which trade completeness for average-case efficiency; discussions of many exact and approximate methods can be found in Latombe's text (Latombe 1991). Other than roadmap methods, the main

class of complete planning methods are known as cell decomposition methods. In these methods, $\mathcal{C}_{free}$ is decomposed into a set of non-overlapping cells and their connectivity is represented by a graph. After identifying the cells containing q_S and q_G, a graph search is performed 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. To the best of our knowledge, Collins' algorithm is the best option available, but it is doubly exponential in the dimension of the ambient C-space (Schwartz, Hopcroft, and Sharir 1987). Polynomial-time algorithms for simplified problems have been developed—for example, motion planning for a "ladder" (a line segment) moving in space among polyhedral obstacles (Schwartz, Hopcroft, and Sharir 1987)—but the work in this area dwindled as successful probabilistic methods began to appear.

Probabilistic roadmap methods (PRMs), beginning primarily with the work of Overmars and Svestka (1994), Kavraki and Latombe (1994) and Kavraki et al. (1996), have taken center stage for about the last 10 years. Roughly speaking, PRMs use random search to construct a probabilistic roadmap, $\mathcal{R}_{free}$, 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 $\mathcal{C}_{free}$ and it will be "easy" to connect arbitrary start and goal configurations to the graph. Despite their incompleteness, PRMs have become popular because they are easy to implement and 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 (Kavraki, Lamiraux, and Holleman 1998), problems with significant dynamic effects (LaValle and Kuffner 2001),[1] problems involving mechanisms with loops (Han and Amato 2000; LaValle, Yakey, and Kavraki 1999), and problems of dexterous manipulation, where the kinematic loop topology varies and further complications arise from contact and friction constraints (Farahat, Stiller, and Trinkle 1995; Trinkle, Farahat, and Stiller 1995; Trinkle and Hunter 1991).

While the successes of PRM methods are clear, we should expect that, as these methods are pushed into domains with complicated constraints, such as those with kinematic loops, it will become increasingly difficult to construct randomized roadmaps. For example, in the case of dexterous manipulation, plans are composed of path segments restricted to strata of differing dimensions (caused by varying numbers of contact constraints). The application of PRMs would require us to generate roadmaps in each stratum and then to connect them (Cherif and Gupta 1999; Son 1996; Trinkle et al. 1993). In general, the number of strata will be exponential in the dimension of the ambient C-space. Given this, we should expect the amount of random sampling required for building a good probabilistic roadmap to be exponential in the dimension of C-space.

The approach taken in this paper represents a return to the first research thrust discussed above: the search for complete, polynomial-time algorithms for a sub-class of the generalized movers' problem. We study the problem of planning reconfigurations of spatial kinematic closed chains with spherical joints and their planar analogs. Our problem is further restricted by neglecting collisions with obstacles and with other links (although we discuss how these could be handled in Section 4.5). While admittedly, our class of problems is too simple to be directly applicable to most tasks of practical interest in robotics, our approach sheds some new light on exact approaches to path planning problems. The approach is based on a complete understanding of the singular sets of certain maps, which is achieved via techniques of modern mathematics, some of which were not previously available. What makes this class of path planning problems difficult even though collisions are ignored is the complexity of the valid portion of C-space, $\mathcal{C}_{kin}$,[2] which is a real algebraic variety of co-dimension three (for spatial closed chains) or two (for planar closed chains) that is not necessarily parametrizable.

Beyond the scope of this paper, we have found that our methods can be applied to more complex problems, including the collision-free motion planning problem for closed chains (in the presence of obstacles and joint limits), and to spatial kinematic chains with revolute joints that represent proteins. In the latter case, we have found that stable secondary structures in proteins, known as "beta-sheets" and "alpha-helices," correspond to singular sets of degrees 1 and 2, respectively. Moreover, we cannot fully understand how proteins fold without the analysis done for the special case of planar closed chains that we discuss later. This is partly because many proteins develop planar substructures and pass through their critical points as they fold, but also because some of the singular sets, while not planar, have the same structure as these spaces of planar configurations.

1.1. Related Work on Closed Kinematic Chains

In this section, we discuss in more detail the connections our work has to existing research for systems with closed kinematic chains. There are three primary categories: exact path planning, probabilistic motion planning, and mechanism singularity analysis. Note that the terminology "motion planning" generally implies that collisions are avoided. We use the term "path planning" to imply that collisions are ignored.

1. Note that C-space is replaced by state space for dynamic problems.

2. The subscript "*kin*" indicates that the valid configurations must satisfy the kinematic constraint of loop closure.

1.1.1. Exact Motion Planning

In the early 1990s Lenhart and Whitesides (1994) and Iarocci (1994) studied the problem addressed in this paper with the same assumptions. They proved several global results for the valid configuration space $\mathcal{C}_{kin}$ of closed chains, which allowed them to develop a complete path planning algorithm that used $O(n)$ "line-tracking" moves. Given q_S, q_G, and the link lengths, $O(n)$ line-tracking moves were used to construct a path to transform the closed chain from q_S into a triangle. The same was done for q_G. Then $O(n)$ line-tracking moves were used to transform one triangle into the other. The main problems with this algorithm are that the paths produced contain a large number of slope discontinuities (one for each "basic" move; line-tracking moves can be composed of several basic moves) and that every plan contains an intermediate triangular configuration even when q_S and q_G are quite "near" each other.

In this paper, we use additional global properties of $\mathcal{C}_{kin}$ to design a planning algorithm that also uses at most $n-2$ *accordion moves* with one velocity discontinuity per move. In addition, we exploit the global structure of $\mathcal{C}_{kin}$ in a special planar case that allows us to complete a path in progress with a single linear interpolation. As a result, our paths are smoother that those produced by the algorithm of Lenhart and Whitesides (1994). On the other hand, their algorithm is more efficient, because each line-tracking move can be computed in constant time, whereas in our implementation the computation time of *accordion moves* is proportional to m^2, where m is the number of links in the current accordion.

1.1.2. Probabilistic Motion Planning

Recently, there have been two efforts to plan collision-free motions for systems with kinematic loops and fixed topology, both adopting the PRM paradigm (Han and Amato 2000; LaValle, Yakey, and Kavraki 1999). The approach taken by LaValle, Yakey, and Kavraki (1999) targeted more general systems: systems with constraints expressed as algebraic equations (not limited to kinematic constraints). Because of its generality, the method was somewhat inefficient. Points in a probabilistic roadmap $\mathcal{R}_{kin}$ were found by first choosing random points in the ambient C-space and then applying an iterative scheme to reduce the residuals of the constraints (hopefully to zero). The configuration found by each successful minimization was tested for collisions before inserting it into the roadmap. Planning queries were processed in the usual way. Planar test problems with up to eight links and two loops were computed in several hours.

Han and Amato (2000) specialized the method of LaValle, Yakey, and Kavraki for problems in which the constraint equations enforced kinematic loop closure, and they implemented it for floating spatial mechanisms. Considerable improvement was obtained by breaking the kinematic loops to allow closed-form inverse kinematic solutions to replace the iterative solution approach used by LaValle, Yakey, and Kavraki. An additional improvement was obtained by applying the PRM idea hierarchically. With one link fixed in space, a probabilistic roadmap, $\mathcal{R}_{kin}$, of the chain's valid configurations was generated without concern for collisions. Then, since the link fixed during the generation of $\mathcal{R}_{kin}$ was really allowed to be positioned and oriented freely in the motion planning problem of interest, copies of $\mathcal{R}_{kin}$ were randomly placed in $SE(3)$ (the special Euclidean group). Nodes and arcs of the copies of $\mathcal{R}_{kin}$ generating collisions were eliminated. Finally, a randomized planner for rigid bodies was used to try to connect corresponding nodes in the copies of $\mathcal{R}_{kin}$ to complete the full roadmap. Han and Amato solved planning problems for planar and spatial linkages with 7–9 links and workspace obstacles. Roadmap construction required less than 1 min of cpu time.

As a final comment on PRMs, because of the complexity of the valid subset of C-space, we should expect roadmap construction times for generalized movers' problems to increase exponentially with the dimension of the ambient C-space and the degrees of the polynomials representing the constraints. As suggested by Han and Amato, advantages gained by using "more deliberate" methods, such as incorporating closed-form inverse kinematic solutions and the approach presented here, will become increasingly important.

1.1.3. Mechanism Singularity Analysis

As indicated above, our approach relies on an understanding of the global structure of the singular sets of certain maps. Approximately 20 years ago, researchers studying the design of spatial mechanisms recognized the importance of various singular sets and developed methods for finding points in them; see, for example, Baker (1978), Derby (1980) and Sugimoto, Duffy, and Hunt (1982). Their basic approach was to obtain the Jacobian of the loop closure equations as a function of the kinematic parameters and the joint variables. Then various design questions were answered by the numerical solution of appropriate equations derived from the Jacobian. More specifically, Baker (1978) found limiting positions of links, and Derby (1980) found the maximum reach of an open spatial chain. Sugimoto, Duffy, and Hunt (1982) found "stationary" configurations (a generalization of Baker's limiting positions) and "uncertainty" configurations by monitoring all seven (6×6) minors of the Jacobian as a driving link was incrementally moved.

The distinguishing common thread in this previous work—and this line of research continues today; see, for example, Carricato and Parenti-Castelli (2002) and Downing, Samuel, and Hunt (2002)—is the use of numerical methods to find individual singular points of co-rank one.[3] This approach is in

3. Given a C^∞ map $f : M \longrightarrow N$ of differentiable manifolds, a point $m \in M$ is singular if and only if the Jacobian of f at m does not have full rank. In this case, the *co-rank* of the singular point is the number $dim(N) - dim(Im(\boldsymbol{J}))$, where $\boldsymbol{J}$ is the Jacobian, $dim(N)$ is the dimension of the manifold N, and $Im(J)$ is the image of J.

striking contrast to ours, in which we obtain the entire singular set and characterize it globally. The closest the previous work comes to this is to iteratively obtain many singular points for a one-degree-of-freedom mechanism in an attempt to approximate the singular set. However, such an approach cannot be used to determine the properties and full global structure that is needed in our approach.

The mechanism work most closely related to our work was carried out by Ting (1989), Ting and Liu (1991) and Alizade and Sandor (1985). They focused on the problem of discovering "fully rotatable" links within a mechanism. While the connection to singularity analysis may not be immediately clear, the existence of "fully rotatable" links implies important global structure and is determined by proving that conditions used to define "special" and "uncertainty" configurations cannot be satisfied. In their work, Alizade and Sandor focused on two particular classes of one-degree-of-freedom spatial mechanisms. Their method ultimately relied on the numerical solution of equations derived from singularity conditions, and thus could not be used to provide the complete global topological properties of C-space that we seek. However, by restricting his attention to planar closed chains with revolute joints, Ting was able to derive closed-form rules for rotatability for three classes of closed chains (one degenerate). As will be seen later, we use some of Ting's results, in addition to our newer results, to design our path planning algorithm for planar closed chains.

1.2. Summary of Our Approach

We begin with a closed kinematic chain with spherical joints (revolute for planar closed chains). One link, the base link, is fixed in the workspace. We wish to determine the detailed structure of the valid portion of C-space, $\mathcal{C}_{kin}$. This problem has resisted analysis for a very long time. However, if we break the kinematic loop by disconnecting a joint at one end of the base link, we form an open chain based at the other end of the base link. The C-space of this open chain is a product of special orthogonal groups $SO(3)^{n-1}$ and the end-point of the open chain is free to map anywhere it can (inside the annular workspace of the open chain). However, because we do not consider collisions in our analysis, the shapes of the links (and also their orientations about the line segments connecting the joints) is not relevant. Thus the essence of the mechanism's kinematics is represented more concisely as a product of spheres $(S^2)^{n-1}$, and this then serves as the basis of our analysis.[4]

Since the relevant C-space of an open chain is the product of spheres, we can focus on the end-point map, $f : (S^2)^{n-1} \longrightarrow \Re^3$, ($f$ sends each configuration of the open chain to the image of the end-point of the final link) of this C-space to construct $\mathcal{C}_{kin}$ of the original closed chain. This is obtained as $f^{-1}(e)$, with e being the point at the end of the base link where the joint was disconnected to break the chain. In order to determine $f^{-1}(e)$ we use techniques from algebraic geometry and differential topology to determine the inverse images of *all* the points p in the image, $Im(f)$ (not just e), $f^{-1}(p)$, as p runs over all points $f(q), q \in (S^2)^{n-1}$. Thus we find the desired C-space $f^{-1}(e)$ as a special case. We denote the image space as $\mathcal{V}$, and note that it is either a closed ball or a closed annulus.[5]

In detail, first we determine the singular points of f, which are not thought of as isolated points, but rather as points making up an intrinsic algebraic sub-variety of $(S^2)^{n-1}$. In the special case considered here, these singular sets turn out to be 2^{n-2} disjoint spheres $S^2 \subset (S^2)^{n-1}$. Moreover, f maps these singular spheres onto concentric spheres in $Im(f)$ centered at the base of the open chain, the specific radii of these image spheres depending on the explicit lengths of the links. This is actually a special case of a very general and important result in mathematics, the Morse–Sard theorem that states, for C^∞ maps, that the image of the singular set always has measure 0 in the image manifold; see Guillemin and Pollack (1974). In our situation, where we are dealing with an algebraic map of algebraic varieties, the singular set is an algebraic sub-variety and the image is, consequently, an algebraic sub-variety of $\mathcal{V}$ of dimension at most one less than the dimension of $\mathcal{V}$.

The image of the singular set decomposes $\mathcal{V}$ into connected regions, in the case here, open annuli centered at the base of the open chain. Since $(S^2)^{n-1}$ is compact, the map f, restricted to the inverse image of each open region is a fibration with fiber a closed, compact manifold of dimension $2n-5$; see Guillemin and Pollack (1974). Consequently, the inverse images for any two points in one of these regions are diffeomorphic. Thus the problem reduces to two basic steps: first, find the inverse image in one of these regions; and, second, understand how the image changes when we pass through the image of a singular point.

When the singular set is generic, the map f restricted to the singular set is a local diffeomorphism and standard methods of differential topology and algebraic geometry allow the determination of the way in which the global structure of the space $f^{-1}(p)$ changes as we pass through the image of the singular points. Among the critical things we learn is how the number and structure of the connected components in the inverse image changes. Indeed the component set provides the solution to the existence problem: a path from configuration q_S to configuration q_G exists if and only if q_S and q_G are in the same path component in $f^{-1}(p)$. Further information about the global structure of $f^{-1}(p)$ gives information on how to construct efficient paths consisting of long geodesic segments.

4. In Section 4.5 we indicate some of the steps that are needed to go from the model used here to models that would be needed to consider obstacles and self-intersections of the links. The extension of the results is direct, involving lifting paths in a product of S^2 to paths in a product of $SO(3)$.

5. The analysis can be extended to include prismatic joints (with finite travel), but when we do this, the C-space of the open chain is a product of spheres and intervals.

It is worth noting that, while in this paper we only consider the end-point map into $\Re^3$, $\mathcal{V}$ could be chosen to be a subset of $\Re^3 \times S^2$ if the orientation of the last link is constrained. Also, the techniques described above are still available to handle the analysis in this case.

2. Problem Formulation

Let $\mathcal{M}$ denote a closed kinematic chain with n links connected by spherical joints in Euclidean space, $\Re^N$, with $N = 2$ or $N = 3$. When $N = 3$, $\mathcal{M}$ represents a spatial n-bar mechanism with spherical joints. In Figure 1, the joints, in positions $\vec{j}_i;\ 1, \ldots, n;$, are represented by spheres and the links by cylinders connecting them. The vector along link i from joint i to joint $i+1$ will be denoted by $\vec{l}_i;\ i = 2, \ldots, n;\ \vec{l}_1 = \vec{j}_1 - \vec{j}_n$ with lengths, $||\vec{l}_i|| = l_i;$, are the distances between the centers of the joints and are assumed to be strictly positive and constant. Without loss of generality, we also assume that the first (or base) link is fixed to lie on the x-axis with its initial end at the origin. Each link's orientation is specified by two angles, ϕ_i measured in the xy-plane and ψ_i measured from the z-axis, as shown in Figure 2. A configuration q of $\mathcal{M}$ is specified by the orientations of all the unfixed links (i.e., $q = [\psi_2, \phi_2, \ldots, \psi_n, \phi_n]$). Thus the space $\mathcal{C}_{kin}$ of all valid configurations of $\mathcal{M}$ is contained in the ambient C-space, $\mathcal{C}$, the product of $n-1$ spheres, $(S^2)^{(n-1)}$.

When $N = 2$, $\mathcal{M}$ represents a planar n-bar mechanism with revolute joints (see Figure 3). In this case, $\psi_i = \frac{\pi}{2}$ for all i, and so a configuration is represented as an $(n-1)$-tuple, $q = [\phi_2, \ldots, \phi_n]$ and the C-space of $\mathcal{M}$ is contained in the product of $n-1$ circles, $(S^1)^{(n-1)}$.

For $\mathcal{M}$ to form a closed loop as required, the lengths and angles must satisfy the following constraint:

$$F(q) = \sum_{i=1}^{n} l_i \begin{bmatrix} \sin(\psi_i)\cos(\phi_i) \\ \sin(\psi_i)\sin(\phi_i) \\ \cos(\psi_i) \end{bmatrix} = 0. \qquad (1)$$

Equation (1) restricts the C-space of $\mathcal{M}$ to a real algebraic variety, $\mathcal{C}_{kin}$ (not necessarily a manifold[6]), embedded in $\mathcal{C}$. Our problem can now be stated as follows.

Path Planning Problem: Given start and goal configurations, $q_S \in \mathcal{C}_{kin}$ and $q_G \in \mathcal{C}_{kin}$, find a continuous path $\tau : [0, 1] \longrightarrow \mathcal{C}_{kin}$ such that $\tau(0) = q_S$ and $\tau(1) = q_G$.

As a preview of our results, Figure 4 shows start and goal configurations for two six-link planar closed chains with the same base link lengths (bold and horizontal), but with slightly different lengths for the other links. While not intuitively obvious, only one of these two problems has a solution, which our algorithm (implemented in Matlab on a 650 MHz Pentium III PC) found in less than 0.1 cpu seconds. This same algorithm has been successful in solving motion planning problems for planar closed chains with up to 10,000 links (in a few seconds) for one class of planar closed chains (Ting's "Class I" (Ting 1989; Ting and Liu 1991)), but only up to 500 links (in 16 h) for the other class (Ting's "Class II").

6. By m-dimensional manifold, we mean a topological space of dimension m that is locally Euclidean.

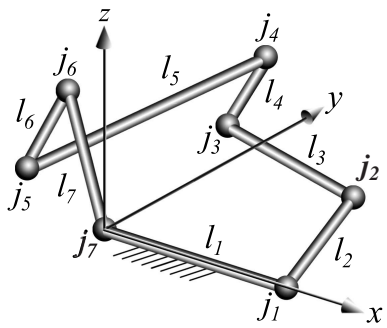

Fig. 1. Spatial seven-link closed chain with spherical joints. Link 1 is fixed on the x-axis.

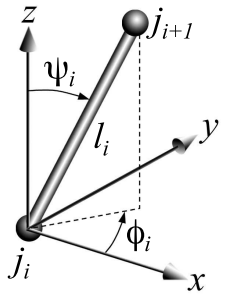

Fig. 2. Link orientation measured in spherical coordinates.

3. Properties of the Closure Variety, $\mathcal{C}_{kin}$

In our discussion of configuration spaces, it is convenient to introduce the following notation:

- $\emptyset$ is the empty set;
- L is the sum of the link lengths (i.e., $L = \sum_{i=1}^{n} l_i$); and
- λ_i is the length of the ith longest link.

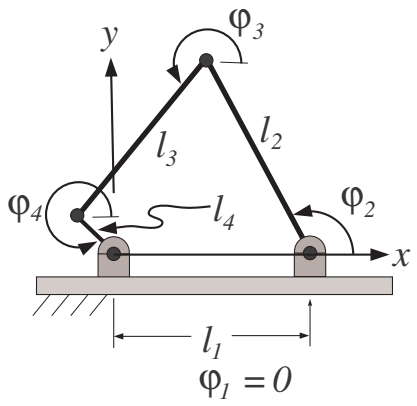

Fig. 3. Four-link planar closed chain.

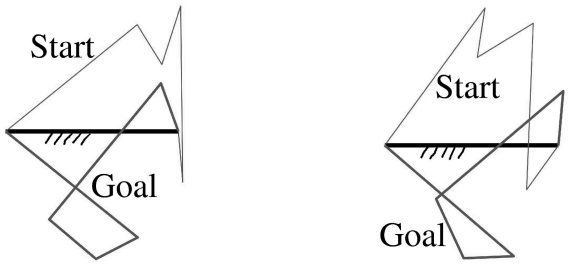

Fig. 4. Start and goal configurations of two six-link closed chains with slightly different link lengths. For one, no connecting path exists.

3.1. $\mathcal{C}_{kin}$ for $\mathcal{M}$ with Two or Three Links

A closed chain cannot be formed with fewer than two links. For the numbers of links, $n = 2$ and $n = 3$, we can determine $\mathcal{C}_{kin}$ by inspection:

- for $n = 2$, $\mathcal{C}_{kin}$
$$= \begin{cases} \text{a single point} & \text{if} \quad \lambda_1 = \lambda_2 = \frac{L}{2} \\ \emptyset & \text{if} \quad \lambda_1 \neq \lambda_2 \end{cases}$$

- for $n = 3$, $\mathcal{C}_{kin}$
$$= \begin{cases} S^1 & \text{if} \quad \lambda_1 < \frac{L}{2} \text{ for } \mathcal{M} \text{ in } \Re^3 \\ \text{two isolated points} & \text{if} \quad \lambda_1 < \frac{L}{2} \text{ for } \mathcal{M} \text{ in } \Re^2 \\ \text{a single point} & \text{if} \quad \lambda_1 = \frac{L}{2} \\ \emptyset & \text{if} \quad \lambda_1 > \frac{L}{2} \end{cases}.$$

For the case with $n = 3$, a distinction was made between the spatial and planar cases. When λ_1 is less than $\frac{L}{2}$, the links can form a triangle. In the spatial case, l_1 is fixed on the x-axis, but the other two links can rotate about it. Thus $\mathcal{C}_{kin}$ is a circle embedded in the ambient C-space, $\mathcal{C}$. In the planar case, one can assemble the triangle in two ways; "elbow-up" and "elbow-down," and once assembled, no motion is possible. Thus $\mathcal{C}_{kin}$ is two isolated points in $\mathcal{C}$.

3.2. $\mathcal{C}_{kin}$ for $\mathcal{M}$ with More than Three Links

When the number of links, n, is greater than three, it has been shown independently by Ting (1989), Ting and Liu (1991) and Lenhart and Whitesides (1994) that the properties of $\mathcal{C}_{kin}$ are strongly dependent on the number of "long links." We use the refined definition of "long links" found in Kapovich and Millson (1995):

DEFINITION 1. A subset $\mathcal{L}_L$ of the links is referred to as the "long links" if and only if the sum of the lengths of every pair of distinct links in $\mathcal{L}_L$ are strictly greater than $\frac{L}{2}$.

Note that $\mathcal{L}_L$ is not unique, so the number of long links is defined as the maximum cardinality of all possible $\mathcal{L}_L$. For example, given lengths [0.255, 0.252, 0.247, 0.246], there are three possible sets of long links, each of cardinality 2, so this set has two long links. However, when there are three long links, $\mathcal{L}_L$ is unique and consists of the three longest links. We can draw the following conclusions from Definition 1:

1. *There can be no more than three long links.* Given the definition of long links, four long links would have lengths summing to a value strictly greater than L, which is a contradiction.

2. $\mathcal{C}_{kin}$ *is empty if and only if there is one long link.* If $\lambda_1 > \frac{L}{2}$, then the other links would not be able to reach from one end of the longest link to the other. Thus, loop closure equation (1) could not be satisfied. The "only if" direction is proved in Kapovich and Millson (1995) by a relatively direct (but non-obvious) induction, starting from the case where there are only two links.

The following less obvious theorem was proved in different ways by Lenhart and Whitesides (1994) and Kapovich and Millson (1996):

THEOREM 1. (Lenhart and Whitesides (1994) and Kapovich and Millson (1996)). For $\mathcal{M}$ in $\Re^3$, if $\mathcal{C}_{kin}$ is non-empty, then $\mathcal{C}_{kin}$ has only one component.

This result is important, because it answers the path existence question. If the longest link length λ_1 is less than or equal to $\frac{L}{2}$, then a continuous path exists between any two points in $\mathcal{C}_{kin}$. Interestingly (as indicated by the example shown in Figure 4) this result does not apply to $\mathcal{M}$ in $\Re^2$. This issue is discussed in the next section.

Our main theoretical results for closed chains, Theorem 2, give more detail about the structure of $\mathcal{C}_{kin}$. These results were derived using techniques from differential topology and algebraic geometry, and were first presented and proved in Milgram and Trinkle (2002). Denoting the interval by I, we have:

THEOREM 2. (Milgram and Trinkle (2002)). For $\mathcal{M}$ in $\Re^2$ or $\Re^3$ with given lengths, $l_1, \ldots, l_n$, then:

(a) Except for a finite number of base lengths, l_1, $\mathcal{C}_{kin}$ is a closed compact manifold:

$$\begin{array}{lll} \text{of dimension} & n-3 & \text{for } \mathcal{M} \text{ in } \Re^2 \\ \text{of dimension} & 2n-5 & \text{for } \mathcal{M} \text{ in } \Re^3. \end{array}$$

(b) Whenever $\mathcal{C}_{kin}$ is a manifold, it is the boundary of a manifold $\mathcal{W}$ with a boundary,[7] which is given as the union of sub-manifolds of the form:

$$\begin{array}{lll} (S^1)^s \times I^{(n-s-2)} & \text{for} & \mathcal{M} \text{ in } \Re^2 \\ (S^2)^s \times (I^2)^{(n-s-2)} & \text{for} & \mathcal{M} \text{ in } \Re^3, \end{array}$$

where the values that s may take are determined by the link lengths.

(c) The intersection of any two such products is again of the form $(S^2)^r \times (I^2)^{n-r-2}$. In detail, these various $(S^1)^s \times I^{(n-s-2)}$ glue together in completely explicit ways that depend on the lengths of the links; technically they glue via "standard plumbing" (Milnor 1967).

The "finite number of base lengths" referred to in Theorem 3a are critical base lengths of the closed chain. At these lengths, the topology of $\mathcal{C}_{kin}$ undergoes significant structural changes. These lengths correspond to the singular configurations of the end-point map of the open chain of links 2 through n (used as described in Section 1.2 to construct $\mathcal{C}_{kin}$). The $2^{(n-2)}$ critical lengths are given as follows:

$$l^*_{1,j} = |\sum_{i=2}^{n} \sigma_i l_i| \tag{2}$$

where each $\sigma_i = \pm 1$. When the length of the base link is equal to one of the critical lengths, then all the links of $\mathcal{M}$ can be colinear. For example, if lengths $\{l_2, l_3, l_4\}$ are given as $\{1, 2, 3\}$, then the critical lengths of the base link would be $l^*_1 = \{0, 2, 4, 6\}$. Choosing $l_1 = 4$ would yield a critical four-link closed chain admitting the singular configuration, $\phi_1 = \phi_2 = 0$ and $\phi_3 = \phi_4 = \pi$, with all links lying on the x-axis. $\mathcal{C}_{kin}$ would be a figure eight, which is not representable as a union of products of circles and intervals glued via standard plumbing.

7. An m-dimensional manifold with a boundary is smooth everywhere except along a boundary of dimension $m-1$ which is an $(m-1)$-dimensional manifold.

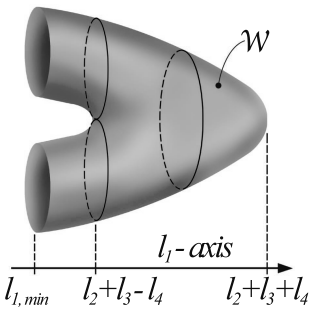

Fig. 5. Manifold $\mathcal{W}$ with boundary representing $\mathcal{C}_{kin}$ parametrized by l_1.

To illustrate Theorem 3b, consider the four-bar shown in Figure 8. As drawn, with l_1 less than $l_2 + l_3 - l_4$, s is equal to 1 and $\mathcal{C}_{kin}$ is two circles[8] and $\mathcal{W}$ is homeomorphic to $I \times S^1$ and can be drawn as the two-dimensional surface shown in Figure 5. If l_1 were lengthened to satisfy $l_2 + l_3 - l_4 < l_1 < l_2 + l_3 + l_4$, then s would equal zero and $\mathcal{W}$ would become homeomorphic to a disk, I^2.

3.3. Special Results for $\mathcal{M}$ in the Plane

For $\mathcal{M}$ in $\Re^2$, $\mathcal{C}_{kin}$ may have more than one component. This fact can be viewed as a result of the additional restriction placed on the motions of the links. However, what may be surprising is that when $\mathcal{C}_{kin}$ is not empty, it may have one or two components; no other number is possible.

The possibility of an extra component makes planning more difficult in the planar case, but only slightly so, because of the following two theorems. The first theorem can be used to determine the number of components of $\mathcal{C}_{kin}$.

THEOREM 3. (Ting and Liu (1991) and Lenhart and Whitesides (1994)). For $\mathcal{M}$ in $\Re^2$, $\mathcal{C}_{kin}$ has two components if and only if $|\mathcal{L}_L| = 3$, otherwise $\mathcal{C}_{kin}$ has only one component.

Theorem 3 implies that the number of components can be computed in $O(n)$ time. If there are two components, then we must determine if q_S and q_G are in the same component. This can be done in $O(1)$ time, as we need only to compute the sign of the cross product of the vectors along the two longest links for both q_S and q_G (Trinkle and Milgram 2001). If and only if the signs are the same, q_S and q_G are in the same component.

The second theorem gives the structure of $\mathcal{C}_{kin}$ when there are two components.

8. Why $\mathcal{C}_{kin}$ is two circles and how we compute s is discussed in Section 3.4.

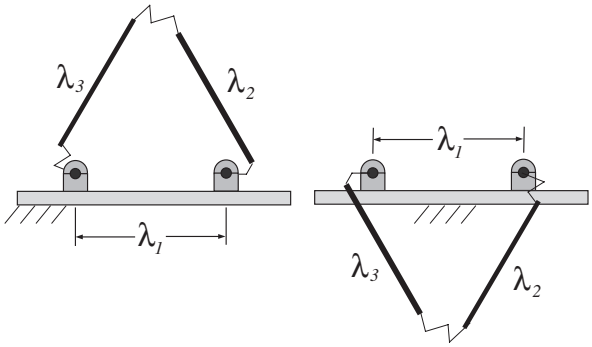

Fig. 6. A pair of configurations in different components of $\mathcal{C}_{kin}$.

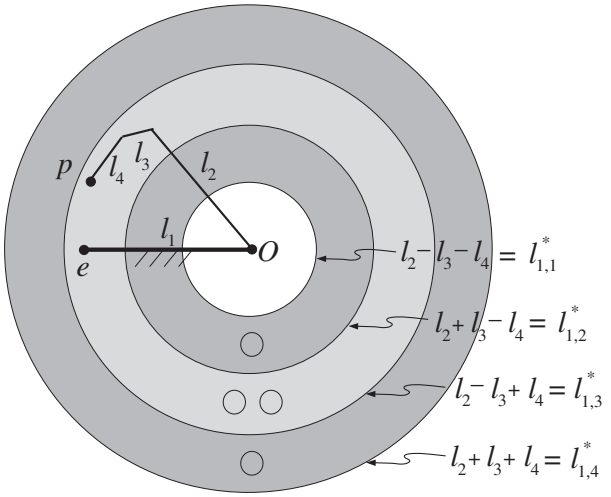

Fig. 7. Three-link open chain with its three open annular regions defined by the four singular image circles. The structure of $\mathcal{C}_{kin}$ of the four-link closed chain formed by fixing p in any of the annular regions is shown by the small circles at "six o'clock" in the annuli.

THEOREM 4. (Kapovich and Millson (1995)). For $\mathcal{M}$ in $\Re^2$, when $\mathcal{C}_{kin}$ has two components, each component is a torus $(S^1)^{n-3}$.

The implication of the toroidal structure is that when $\mathcal{M}$ has three long links, each of the other links, the *short links*, can rotate through a full circle regardless of the orientations of the other short links, while the long links maintain loop closure (Figure 6 shows a typical closed chain with two components). This result also means that each component of $\mathcal{C}_{kin}$ can be covered by a single chart parametrized by the angles of the short links. Moreover, if q_S and q_G are in the same component, then any path connecting the start and goal angles of the short links is valid. In our path planning algorithm (discussed in Section 4), the procedure TwoComponentMove() follows a geodesic on $(S^1)^{n-3}$ between the start and goal angles of the short links. This amounts to linear interpolation, which requires $(O(n-3))$ time, and is the reason we were able to generate paths for 10,000-link closed chains in several seconds in Matlab.

3.4. Construction of $\mathcal{C}_{kin}$

The discussion of the construction of $\mathcal{C}_{kin}$ in this section follows the general approach described in Section 1.2. While the method is valid in $\Re^3$, it is most clearly described through planar examples. We work with planar four-link and five-link closed chains to allow graphical illustration.

To construct $\mathcal{C}_{kin}$, we use the end-point map $f : \mathcal{C} \longrightarrow \Re^2$ for an open chain of m links based at the point (x_b, y_b). Explicitly:

$$f(q) = \begin{bmatrix} x_b \\ y_b \end{bmatrix} + \sum_{i=1}^{m} l_i \begin{bmatrix} \cos(\phi_i) \\ \sin(\phi_i) \end{bmatrix}; \quad q \in (S^1)^m. \quad (3)$$

3.4.1. $\mathcal{C}_{kin}$ for a Four-Link Closed Chain in the Plane

Consider the four-link closed chain whose fourth joint has been removed, thus breaking the chain as shown in Figure 7. The image of the end-point of the resulting open chain is the annulus with inner and outer radii $l^*_{1,1}$ and $l^*_{1,4}$. It is easy to see that the end-point map of the resulting open three-chain is singular when its links are parallel to each other and that the image of the singular configurations is four concentric circles. Recall that our approach requires that we obtain $\mathcal{C}_{kin}$ when the end-point of the open chain is fixed at a point p in any one of the three open annular regions formed by removing the four singular image circles (not necessarily the one containing the original end-point location, e) and that we determine how $\mathcal{C}_{kin}$ changes as p moves across the singular image circles.

We first show how to determine $\mathcal{C}_{kin}$ for a planar four-bar as illustrated by Figure 8. The process is then a kind of induction, i.e., we use knowledge of $\mathcal{C}_{kin}$ for three-link closed chains (given in Section 3.1) to obtain $\mathcal{C}_{kin}$ for four-link closed chains. Disconnecting the third joint of the four-link chain shown in Figure 8 yields two open chains attached to either end of the base link with end-point maps, f_L and f_R. The images of these maps are the circle E_L and the gray annulus E_R. Clearly, the loop closure equation (1) can only be satisfied in the intersection of the end-point maps, $E = E_1 \cup E_2$ (drawn as bold arcs terminated by small filled circles). Thus we obtain a "parametrized" family of C-spaces for three-link closed chains as describing the C-space of a four-link closed chain.

In detail, to construct $\mathcal{C}_{kin}$, we use the inverse images of the maps f_L and f_R. First consider f_R. There are two ways ("elbow-up" and "elbow-down") for the open two-chain to reach a given point $\vec{j}_3$ in $Int(E_R)$ (the interior of E_R), so $f_R^{-1}(\vec{j}_3)$ is two distinct points in $\mathcal{C}$. The singular points of f_R

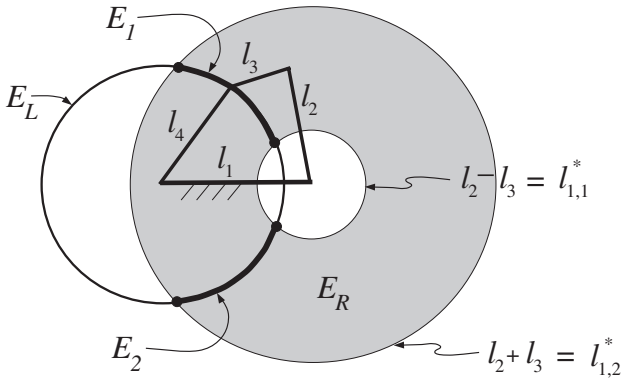

Fig. 8. Four-bar with $\mathcal{C}_{kin}$ homeomorphic to two circles.

map to the boundary of the annulus and the inverse image of each of those points is a single point in $\mathcal{C}$. The map f_L has no singular points and each point in its image can be reached in exactly one way. Now consider the reassembled closed chain with joint 3 fixed at a point $\vec{j}_3$ in E_1 or E_2. At such a point, the contribution to $\mathcal{C}_{kin}$ for the closed chain is the product of the inverse images $f_R^{-1}(\vec{j}_3)$ and $f_L^{-1}(\vec{j}_3)$. When $\vec{j}_3$ is in $Int(E_1)$ or $Int(E_2)$, the inverse image of the closed chain is two points (the product of two points with one point) in $\mathcal{C}$. For $\vec{j}_3$ fixed on the boundary of E_1 or E_2, the product of inverse images is one point in $\mathcal{C}$.

To determine $\mathcal{C}_{kin}$ for the closed chain with $\vec{j}_3$ allowed to move in E, we use Theorem 5.2 in Milgram and Trinkle (2002). The theorem states that the inverse image of any C^∞ rectifiable curve γ in the interior of a region that is disjoint from the image of the singular points (in our case $Int(E_R)$), is the Cartesian product $I \times f_R^{-1}(\gamma)$ (thus the parametrized family now can be recognized as a space). Applying this to the closed chain, the curve $Int(E_1)$ is rectifiable, so its inverse image is the product of two points and an interval, i.e., two arcs (one for "elbow-up" and one for "elbow-down") in $\mathcal{C}$. To complete $\mathcal{C}_{kin}$, note that the inverse images at the boundary points are single points. These points identify the ends of the two curves, creating a closed loop. Thus the inverse image of E_1 is homeomorphic to S^1. Since the same arguments apply to E_2, $\mathcal{C}_{kin}$ is homeomorphic to two copies of S^1.

The preceding discussion allows us to determine $\mathcal{C}_{kin}$ for all planar four-link closed chains graphically. For example, if the length l_1 were in the range $(|l_2 - l_3| + l_4, l_2 + l_3 + l_4)$, E_L would intersect the annular region in Figure 8 in a single arc indicating that $\mathcal{C}_{kin}$ would be a single circle. Likewise, if E_L were contained in the annulus then the inverse image would again be two circles. The critical situations for four-link closed chains are those for which the circle E_L is tangent to one or both of the bounding circles of the annulus. Relating these ideas back to Figure 7, l_1 in the interval $(l^*_{1,3}, l^*_{1,4})$ corresponds to fixing the end-point of the open three-chain in the outer annular region and thus $\mathcal{C}_{kin}$ is a single circle. Decreasing l_1 to the value as shown in the original four-link closed chain causes p to cross the second largest singular image circle, and $\mathcal{C}_{kin}$ changes from a single circle (through a "figure eight" when $l_1 = |l_2 - l_3| + l_4$) to two disjoint circles.

In summary, we have seen that if we take the based open three-link chain with lengths l_2, l_3, l_4, then the image of the end-point map will be a closed disk if $l_i < \frac{1}{2}(l_2 + l_3 + l_4)$ for $i = 2, 3, 4$, and otherwise it will be an annulus. Moreover, the image of the singular set will be the concentric circles with radii given by equation (2).[9] Additionally, in the outermost region, the region between $l_2 + l_3 + l_4$ and the next largest singular image circle, the inverse image of any point under the end-point map will be a single circle, and this is the configuration space of a closed four-link chain where $l_2 + l_3 + l_4 > l_1 > (l_2 + l_3 + l_4) - 2l_{min}$ and l_{min} is the minimum of l_2, l_3, and l_4.

3.4.2. $\mathcal{C}_{kin}$ for an n-Link Closed Chain in the Plane

We now see an easy induction breaking up an $(n + 1)$-link closed chain into an $(n - 1)$-link open chain based at one end of the base link, and a one-link chain based at the other end. Using the observation that the inverse image of any point on the largest singular image circle of radius $\sum_1^{n-1} l_i$ consists of a single point for each n and $(n - 1)$-tuple of link lengths $(l_1, l_2, \ldots, l_{n-1})$, it now shows that the inverse image of any point in the outermost annular region defined by:

$$\sum_1^{n-1} l_i > l_n > \sum_1^{n-1} l_i - 2l_{min}$$

is S^{n-2}, where l_{min} is the minimum of l_i over $i = 1, \ldots, n-1$. (Similar arguments work in $\Re^3$ after replacing circles with spheres and S^{n-2} with S^{2n-5}.)

Thus, we have achieved the initial steps of the method sketched in Section 1.2 for determining the explicit structure of $\mathcal{C}_{kin}$. The remainder of the construction is given, as indicated there, via standard techniques from differential topology and algebraic geometry. If the end-point of link $n + 1$ can pass through a critical circle in $\Re^2$ or a critical sphere in $\Re^3$, the inverse image changes via the attachment of a single "handle"[10] for each critical point in the inverse image of that sphere. Physically, the change in the specification of $\mathcal{C}_{kin}$ is simply that the number of fully rotatable links increases by

9. It is not hard to show this analytically. In general, for any based, open, serial chain in $\Re^3$ with spherical joints or in $\Re^2$ with revolute joints, the image of the singular set will be concentric spheres in $\Re^3$ and circles in $\Re^2$ of radii given by equation (2) and the singular sets themselves will be configurations where all the links are parallel.

10. A "handle" can be thought of as a thickened line segment or interval, or more precisely, the Cartesian product of a disk and another disk. Handle attachment is the process of identifying two points on the existing set with the ends of the thickened line segment. For example, a basket could be formed by attaching a handle to a disk.

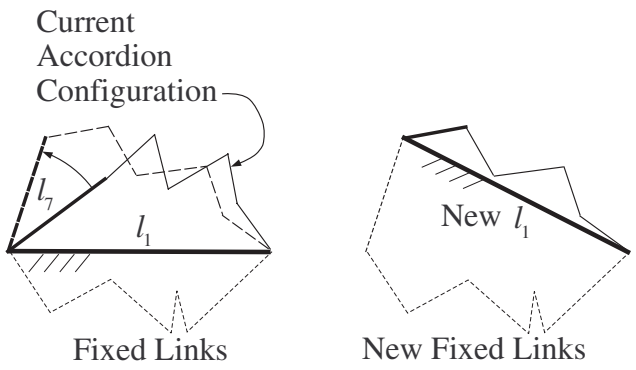

Fig. 9. An *accordion move* (left) and the resulting reduced closed chain (right).

one, and using these two observations, we achieve the main result, Theorem 2 in Section 3.2.

4. A Complete Path Planning Algorithm

Given the start and goal configurations, q_S and q_G, and the link lengths, $l_1, \ldots, l_n$, our algorithm works as follows. The subset of links that are in their goal orientations is fixed (initially this set is empty). A subset of the moveable links is chosen as controlled links. The controlled links are then moved into their goal orientations, while the other movable links comply to maintain the loop closure constraint (1). Figure 9 shows a situation in which there are six fixed links (drawn with short dashes) treated as one (labeled l_1 and drawn bold), one controlled link (labeled l_7), and five compliant links. As the controlled link is driven to its goal orientation (drawn with long bold dashes), the compliant links behave much like an accordion, and so we call such a move an *accordion move*. After the *accordion move*, the controlled links are added to the set of fixed links, thereby strictly increasing the number of links that have achieved their goal orientations and strictly reducing the number of movable links. Our planning algorithm is simply a sequence of *accordion moves* that ends when no movable links remain. We emphasize that this algorithm is complete, because an *accordion move* always exists and can be computed whenever one or more links are not in their goal orientations.

4.1. Accordion Moves

Given current and goal configurations, q and q_G, of a closed chain, the construction of an *accordion move* involves three basic steps:

(1) the identification of a suitable set of controlled links;

(2) the generation of a path of controlled links that transforms them into their goal orientations; and

(3) the generation of a path of the accordion that complies with the motion of the controlled links such that the loop closure constraint is satisfied throughout the motion.

Of course, these paths should be as smooth and efficient as possible.

4.1.1. Selecting Controlled Links and Generating Their Paths

Part (*c*) of Theorem 2 provides an excellent way to choose the controlled subset of movable links and to generate a path for it. Recall that $\mathcal{C}_{kin}$ is built out of regions, each having the form $(S^2)^s \times I^{2n-2s-3}$. This structure directly suggests a partitioning of the movable links into the controlled and accordion subsets. The s links responsible for creating $(S^2)^s$ can rotate freely with the assurance that the loop closure constraint can be satisfied by the remaining movable links, which are responsible for creating $I^{2n-2s-3}$. Therefore, we may choose an arbitrary continuous path of the controlled links on $(S^2)^s$, such as a geodesic. The advantage of this approach is that each *accordion move* can transform a large number of links into their goal orientations simultaneously. However, in general, the computational cost of obtaining the full description of $\mathcal{C}_{kin}$ (and thus determining which s links are fully rotatable) for each closed chain in the sequence of *accordion moves* is exponential in the number of movable links.

If we cannot afford to compute the full description of $\mathcal{C}_{kin}$, we can choose controlled links based upon the minimum and maximum length functions of an open chain. Let $\mathcal{L}$ denote a subset of the movable links and let D denote the length of this subset considered as an open chain (i.e., the distance between the end-points of the chain). Given the configuration q and the subset, the length of the open chain is

$$D(\mathcal{L}, q_{\mathcal{L}}) = || \sum_{i \in \mathcal{L}} \vec{l}_i ||, \tag{4}$$

where we recall that $\vec{l}_i$ is the vector along link i whose length is denoted by $l_i = ||\vec{l}_i||$. The maximum length of the chain represented by this subset is simply the sum of the link lengths

$$M(\mathcal{L}) = \max_{q_{\mathcal{L}}}(D(\mathcal{L}, q_{\mathcal{L}})) = \sum_{i \in \mathcal{L}} l_i, \tag{5}$$

while the minimum is

$$m(\mathcal{L}) = \min_{q_{\mathcal{L}}}(D(\mathcal{L}, q_{\mathcal{L}})) = \max(0, \{2l_i - M(\mathcal{L}) \mid i \in \mathcal{L}\}), \tag{6}$$

which is 0 unless there is a link in $\mathcal{L}$ having length greater than $\frac{1}{2}M(\mathcal{L})$.

Let us denote three index sets of the links, $\mathcal{L}_A$, $\mathcal{L}_F$, and $\mathcal{L}_C$, which define the accordion, fixed, and controlled links,

respectively, and their corresponding partitions of the configuration q, denoted by q_A, q_F, and q_C. We now view the closed chain as two open chains: one composed of the fixed and controlled links, and the other composed of the compliant links (the accordion). Treating the set of fixed links as a single link, we can search the space of possible index sets to find $\mathcal{L}_C$ of maximal cardinality such that:

$$\begin{aligned} m(\mathcal{L}_A) \;&\leq\; |m(\mathcal{L}_C) - D(\mathcal{L}_F, q_G)| \\ &\leq M(\mathcal{L}_C) + D(\mathcal{L}_F, q_G) \;\leq\; M(\mathcal{L}_A). \end{aligned} \tag{7}$$

If a nonempty $\mathcal{L}_C$ is found satisfying equation (7), then all links in $\mathcal{L}_C$ can be rotated freely while the other links comply to maintain loop closure. However, due to the combinatorial nature of the problem of finding $\mathcal{L}_C$ of maximal cardinality, it may be more practical to use heuristics to identify sets of large, but suboptimal cardinality.

Finally, it is possible that the maximum cardinality of $\mathcal{L}_C$ found as described above is zero, therefore we suggest a simple alternative for selecting the set of controlled links. Because the current and goal configurations are always in the same component of $\mathcal{C}_{kin}$, it must be possible to move any single link into its goal orientation. In $\Re^3$, the current and goal orientations of the controlled link define a great circle on the sphere representing the C-space of the controlled link (when disconnected from the accordion). The structure of $\mathcal{C}_{kin}$ guarantees that the controlled link can follow at least one of the two arcs of the great circle connecting the current and goal configurations of the controlled link. We simply compare the length requirements dictated by the two great circles with the extreme lengths of the accordion to decide which arc to traverse.

4.1.2. Generation of Compliant Accordion Motions

Given a path transforming the controlled links into their goal orientations, we must compute a smooth complying motion of the accordion. The problem of finding such a motion will be decomposed into two simpler problems. Let $q_{\overline{A}}$ denote the orientations of the fixed and controlled links (i.e., $q_{\overline{A}} = (q_C,\ q_F)$). The two problems are as follows. First, given known configurations q_C and q_F of the controlled and fixed links, determine a configuration q_A of the accordion such that $D(\mathcal{L}_A, q_A) = D(\mathcal{L}_{\overline{A}}, q_{\overline{A}})$. Second, given q_A such that $D(\mathcal{L}_A, q_A) = D(\mathcal{L}_{\overline{A}}, q_{\overline{A}})$ is satisfied, determine a rigid rotation, $\boldsymbol{R}$, of the accordion such that the loop closure constraint (1) is satisfied.

Let us begin with the observation that if we are given a subset, $\mathcal{L}$, of the links, their configurations, $q_{\mathcal{L}}$, and an arbitrary orthogonal transformation $\boldsymbol{R}$, then the following holds:

$$D(\mathcal{L}, q_{\mathcal{L}}) = ||\boldsymbol{R} \sum_{j \in \mathcal{L}} \vec{l}_j|| = || \sum_{j \in \mathcal{L}} \boldsymbol{R}\vec{l}_j||. \tag{8}$$

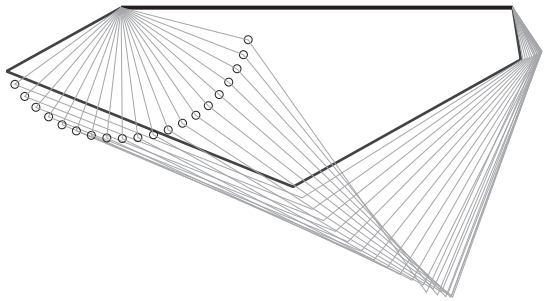

Fig. 10. *Accordion move* with one controlled link and three accordion links.

As an example, if we are in the plane and the orientation of $\vec{l}_j$ is ϕ_j and $\boldsymbol{R}$ represents a rotation through an angle θ, then the orientation of $\boldsymbol{R}\vec{l}_j$ is $\theta + \phi_j$. The equation simply indicates that the entire open chain defined by the subset $\mathcal{L}$ can be rotated freely as a rigid unit without changing the length of the chain.

Assume now that the subset of interest, $\mathcal{L}_A$, identifies a given accordion and that the configuration $q_{\overline{A}}$ of the fixed and controlled links is given. Furthermore, assume that the length, $D(\mathcal{L}_{\overline{A}}, q_{\overline{A}})$, of the chain of fixed and controlled links is within the range of possible accordion lengths (i.e., there exists some q_A such that $D(\mathcal{L}_A, q_A) = D(\mathcal{L}_{\overline{A}}, q_{\overline{A}})$) and that such a value of q_A has been obtained. As indicated above, the loop closure constraint will generally not be satisfied by q_A, but a rotation $\boldsymbol{R} \in SO(3)$ always exists such that loop closure is achieved.

These length-matching and rotation steps can be performed so that if the chain of controlled and fixed link moves along a C^∞ path, $q_{\overline{A}}(t) : t \longrightarrow [0, 1]$ coordinated by t and the accordion moves along another C^∞ path $q_A(t) : t \longrightarrow [0, 1]$ such that for each t the respective lengths are equal, then we can find a C^∞ map of the unit interval $\boldsymbol{R}(t) : [0, 1] \longrightarrow SO(3)$ so that the closure condition is satisfied everywhere along the path. Figure 10 shows the result of our planar implementation of a smooth *accordion move* (broken into 20 steps) with one controlled link and three accordion links.

The remaining unexplained piece of our planning algorithm is a method for generating a path of the accordion that changes its length from its current value $D(\mathcal{L}_{\overline{A}}, q_{\overline{A}})$. To do this we use the gradient vector field associated with the accordion's length function $D(\mathcal{L}_A, q_A)$ and we follow the flow lines from the current configuration to the desired length. This approach has the desirable property that the relative motions of the links in the accordion are minimized (as illustrated in Figure 10). However, gradient following can fail only if the flow line converges to an intermediate critical point. Fortunately there are only two types of critical points and both can be handled (as required for completeness).

The first type is a nondegenerate saddle point. At these points the links of the accordion are colinear and the length of

the accordion is nonzero. Since the length of the accordion is independent of the orientation of the first link, we can fix this link's orientation during gradient following and then notice that there is a finite number of nondegenerate saddle points on a known lattice, which must be avoided. This is done by checking proximity to such points while stepping along the flow line. If the current configuration comes too close to a saddle, first a second-order correction process is used to modify the vector field in this neighborhood of the saddle, forcing the path away from the saddle without changing the length of the accordion. This adjustment can fail in one very special, but detectable, case. In this case, we use a linear move of the type used by Lenhart and Whitesides (1994) to get beyond the saddle point. In either case, once past the saddle point, our algorithm returns to following the gradient vector field. Note that, since we pass through critical points, our algorithm cannot get stuck at them.

The second type of critical point to avoid is degenerate in the sense that the gradient is not well defined. Fortunately, such points only occur when the length of the accordion is zero (which again is easy to detect). If the motion of the controlled links require that the length of the accordion passes through zero, then the accordion motion is broken into intervals that are monotonic in length and the accordion motion for each interval is computed separately. Any zero length points will occur on the boundary of one or more of these intervals. Thus a path containing zero length points will be computed from both sides of those points, yielding a continuous path. It is worth mentioning that this non-differentiable motion part of our algorithm never had to be activated in any of our tests. However, it must be handled to achieve algorithmic completeness. (But in the program, as it currently exists, we have only put in the hooks for this part, not the actual code.)

The final step, once the accordion move has been specified, is to modify it by multiplying by a continuous family of rotations so as to obtain the actual complying motion of the accordion. In the two-dimensional case, where the program has been implemented, this is easy. It involves simply multiplying by a rotation through an angle θ, where θ is easily determined and is C^∞ along the part of the path where the length is non-zero, and continuous where it is zero, after perhaps reparametrizing the path (see Figure 10 again). The process in $\Re^3$ is similar, but somewhat more complex involving path lifting in the fibration $S^1 \rightarrow SO(3) \longrightarrow S^2$ and a smoothing reparametrization of the path near the points where the length is zero.

4.2. Special Considerations for $\mathcal{M}$ in $\Re^2$

The basic algorithm is the same in both the planar and spatial cases: accordion moves and chain reductions (accomplished by assigning the set of controlled links from the most recent *accordion move* to the set of fixed links $\mathcal{L}_F$) are executed. However, in the planar case, the algorithm must be modified to handle the situation in which a reduction causes the current and goal configurations to be in different components of $\mathcal{C}_{kin}$ of the current closed chain. The flow chart shown in Figure 11 shows our algorithm for the planar case, where i is the current number of links in the chain (with all fixed links counted as one) and C is the number of controlled links used in the most recent *accordion move*. The procedures InversionMove() and TwoComponentMove() execute special *accordion moves* that are discussed next.

A *two-component move* is used whenever $\mathcal{C}_{kin}$ has two components and the current and goal configurations are in the same one. The results of Theorem 4 provide all the detailed structure of $\mathcal{C}_{kin}$ needed to compute a final move to the goal. In this case, one of the three long links is treated as fixed, while the other two constitute the accordion. The remaining links are the controlled links, which are moved along a geodesic in $(S^1)^{i-3}$ to their goal locations. It is guaranteed that the long links can comply to maintain loop closure. Recall that in general, determining the set of fully rotatable links requires time exponential in the number of movable links. However, in the planar case, all that is required is to find the three longest links in the current closed chain, which can be done in $O(i)$ time. Choosing and constructing a geodesic can also be done in $O(i)$ time, since we need only to find the direction of shortest rotation to the goal for each link and then to perform linear interpolation from the current to goal angles of these links.

The other special *accordion move* used in the planar case is called an *inversion move*. This is needed when the linkage reduction following an *accordion move* leads to $\mathcal{C}_{kin}$ of the reduced closed chain with two components and with the current and goal configurations in different components. The picture on the left-hand side of Figure 12 shows the situation where links 6 and 9 have just been the fixed and controlled links in an *accordion move*. After replacing them with a single link (labeled λ_3), the current and goal configurations are in different components. The objective of an *inversion move* is to move the closed chain from an "elbow-up" to an "elbow-down" configuration. This is accomplished by temporarily reactivating the joint between the links just fused (links 6 and 9 in the figure) while controlling the relative angle of the other two long links (those labeled λ_1 and λ_2 in the figure). After an *inversion move*, the reactivated joint is again fixed, and now the current and goal configurations are in the same component of $\mathcal{C}_{kin}$. A *two-component move* completes the plan. Note that during every move leading to this sort of problem, the links with lengths λ_1 and λ_2 exist as individual links, and so an *inversion move* is always possible to correct the problem.

4.3. Average Complexity

The most computationally intensive part of our planning algorithm is the standard *accordion move*. This procedure uses an iterative Newton-type algorithm for tracking flow lines of the accordion length function. Each iteration computes the gradient, which requires $O(|\mathcal{L}_A|^2)$ time. In order to provide a

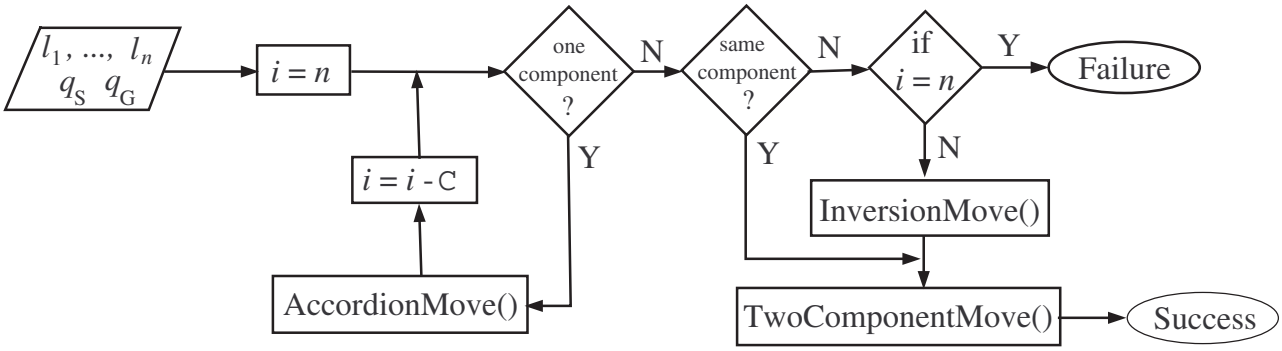

Fig. 11. Path planning flow chart for $\mathcal{M}$ in $\Re^2$, where i is the number of links in the current chain with all fixed links counted as a single link. The variable C in the box containing $i = i - C$ is the cardinality of the set of controlled links (i.e., $C = |\mathcal{L}_C|$).

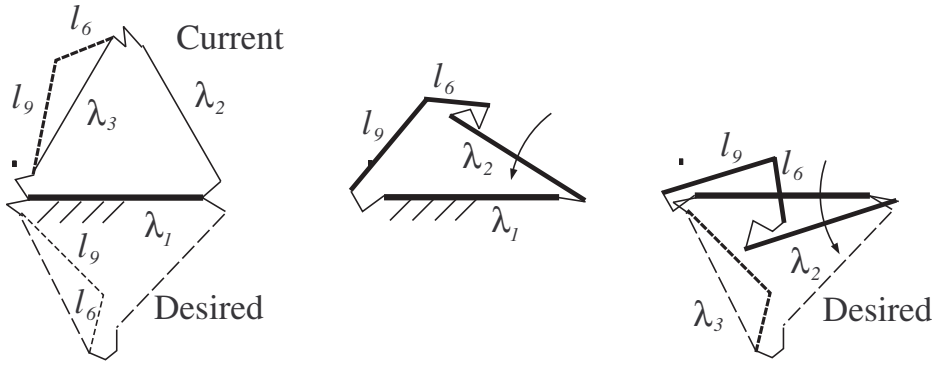

Fig. 12. An *inversion move* (left to right) reconfiguring a closed chain from "elbow-up" to "elbow-down."

good guess for Newton iterations, our implementation breaks each *accordion move* into $O(1)$ small steps and uses the configuration obtained at the end of each sub-move as the starting point for the next. As a result, the number of Newton iterations is small and, in our experience, has been independent of the number of links in the accordion.

Since, in our implementation, the main loop is executed $O(n)$ times, and *accordion moves* require $O(|\mathcal{L}_A|^2)$ time, we expect to see complexity proportional to n^3 or worse (depending on the convergence behavior of the gradient flow algorithm). Figure 13 shows the average planning times (circles) for a closed chain with various numbers of links. Each of the 27 circles represents 100 randomly generated problems for which solutions existed. The link lengths and $n-3$ of their orientations were chosen at random. The base link angle was fixed at zero and the other two unspecified link angles were determined by solving the loop closure equation. An average of $n-5$ *accordion moves* were used for each problem. Notice that the average computation time appears to be strongly cubic, which indicates that, on average, the number of Newton iterations per step within each *accordion move* was constant.

4.4. Comparison with a Local Planner

Our algorithm was compared to a simple algorithm that used only local geometric information. The local algorithm can be thought of as a simple potential field method or proportional controller applied to an open kinematic chain. Imagine breaking the linkage at the left end of the base link, yielding an open chain with $n-1$ revolute joints based on the right side of the fixed base link. Corrections to the current configuration are generated by comparing it with the goal configuration as follows:

$$\begin{bmatrix} \boldsymbol{J} \\ \beta \boldsymbol{I} \end{bmatrix} \mathrm{d}\phi = \begin{bmatrix} \boldsymbol{x}_{desired} - \boldsymbol{x}_{current} \\ \phi_{goal} - \phi_{current} \end{bmatrix}. \tag{9}$$

Here $\boldsymbol{J}$ is the Jacobian of the end-point map of the $(n-1)$-joint open chain, $\boldsymbol{I}$ is the identity matrix, $\boldsymbol{x}_{desired}$ is the location of the left end of the base link, $\boldsymbol{x}_{current}$ is the current location

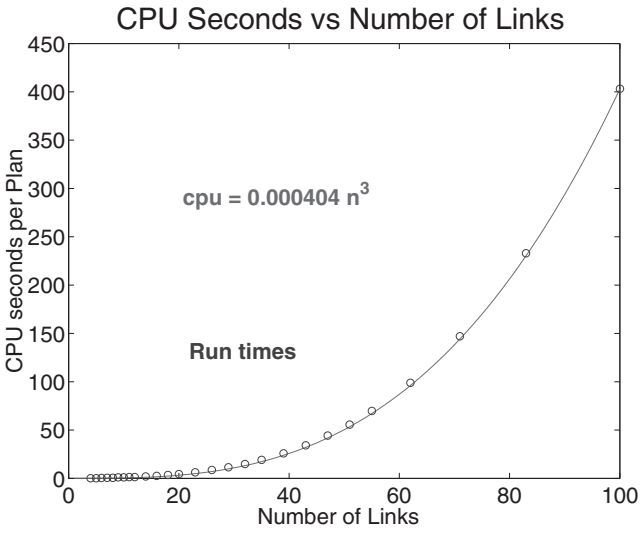

Fig. 13. CPU times to compute motion plans for linkages with 4 to 100 links.

(numerically) of the left end of the base link, ϕ_{goal} is goal configuration, of the linkage, $\phi_{current}$ is the current configuration of the linkage, and β is a scalar weight that determines the relative importance of maintaining the loop closure constraint and attracting the mechanism to its goal configuration.

A number of tests were run to verify the correctness of our algorithm and to compare its performance with the local algorithm just described. For linkages with from 4 to 500 links, the link lengths and the start and goal configurations were chosen at random. Therefore, planning problems with connected and disconnected C-spaces were attempted. Since our algorithm is complete it has solved all problems, so Table 1 only contains the results of the local algorithm.

In Table 1, for each number of links considered (leftmost column), 100 random problems were generated. The second column shows the number of problems for which solutions existed and the last column shows the number of problems solved by the local method. Notice that as the number of links increased, the success rate of the local path planner increased. For problems with over 20 links, the local method solved every problem we generated, but it is guaranteed that the local method (and all other purely local methods) will fail for certain problems even when solutions exists for large numbers of links.

Table 1

numLinks	numExist	numLocal
4	48	31
5	83	36
6	96	44
8	100	55
10	100	81
12	100	85
16	100	95
20	100	98
> 20	100	100

Figure 14 shows a problem drawn from our test set. While the local method generated a smooth motion, it failed to connect the start and goal configurations. Figure 15 shows the individual joint angle trajectories *NOT* achieving their goals (the circles on the right-hand side of the plot).

On the other hand, the complete algorithm found a solution with three *accordion moves* (see Figure 16). The slope discontinuities at the 20- and 40-step points indicate the ends of the *accordion moves*. (The jump discontinuity at the seventh step is a plotting artifact caused by wrap-around from $-\pi$ to π.) We can also see which links were used to drive the linkage, since driving trajectories appear as linear segments. The link whose angle started at about 1.2 and moved to about 3.1 by

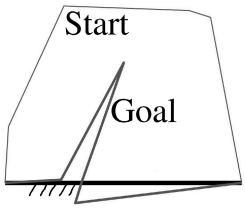

Fig. 14. Planning problem with solution not found by the local algorithm.

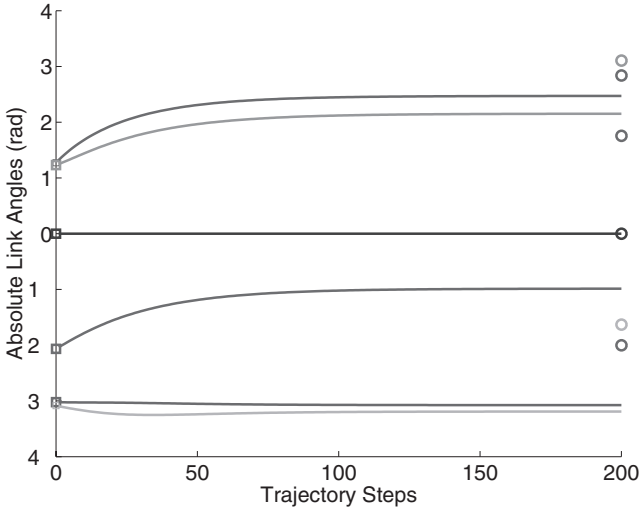

Fig. 15. Faulty trajectory found by the local algorithm.

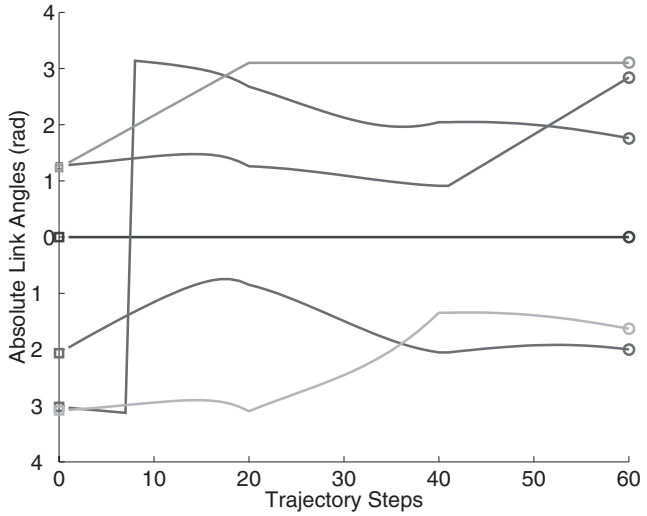

Fig. 16. Trajectory found by our algorithm.

step 20, drove the first *accordion move*. Since it was fused to the base link from then on, the remainder of its trajectory is horizontal.

4.5. Avoiding Obstacle–Link and Link–Link Collisions

The final step in developing a motion planner for closed chains would be to include modifications to ensure collision-free motions; both link–obstacle and link–link collisions (which include joint limits as a special case) should be avoided. Two possible ways to achieve this functionality are: (1) to extend our method to determine the structure of $\mathcal{C}_{kin,free}$ (i.e., the collision-free portion of $\mathcal{C}_{kin}$); and (2) to combine our method with existing probabilistic roadmap methods.

The second option is easiest. Recall that when applying the PRM methodology, we choose points at random in $\mathcal{C}_{kin}$ and then attempt to connect these points with a "local" planner to build the randomized roadmap, $\mathcal{R}_{kin}$. In this case, the "local" planner is any planner that is believed to be capable of finding a path between path-connected configurations; "local" does *not* imply that the planner should use only local geometric information such as that discussed in Section 4.4. Instead, we could use the planner developed in this paper, or that developed by Lenhart and Whitesides (1994). Using one of these planners would reduce the amount of sampling of $\mathcal{C}_{kin}$, since the planner will always succeed in connecting candidate pairs of configurations that are in fact path-connected. Once a connecting path has been obtained, then we must test for collisions along the path, and only accept the path if it is deemed collision-free. If a collision occurs along the path, then the pair of configurations under consideration could be labeled as not connected, or we could attempt to modify the path to eliminate the collision(s).

Now consider the first option: the extension of our global technique. For various reasons, the extension to avoid link–link collisions appears to be much more difficult than that of avoiding link–obstacle collisions. In both cases, however, there are common features that should be understood. The basic idea is that we begin with the well-understood C-space, $\mathcal{C}_{kin}$, studied here and remove from it the C-space of all collisions, $\mathcal{C}_{kin,bad}$. What remains is the collision-free portion, $\mathcal{C}_{kin,free}$. Fortunately, $\mathcal{C}_{kin,bad}$ can be obtained using an approach similar to that used to obtain $\mathcal{C}_{kin}$. For example, suppose there is a point obstacle in the workspace of a closed chain with links modeled as line segments (note, however, that the approach is not limited to point obstacles and one-dimensional links). We would obtain the closed set of invalid configurations as a union of C-spaces corresponding to each link constrained to be in contact with the obstacle. The C-space for a given constrained link can be found using the techniques employed in this paper, because requiring contact with the point obstacle effectively partitions the closed chain into two smaller closed chains whose C-spaces we know how to construct. The C-space of the chain constrained to contact the obstacle is the product of the C-spaces of the two smaller chains. The invalid configurations corresponding to link–link collisions could be found similarly, but with two links constrained to intersect rather than a link and an obstacle. Then the key issue is to identify the path components in

$$\mathcal{C}_{kin,free} = \mathcal{C}_{kin} - \mathcal{C}_{kin,bad}.$$

In general this is extremely difficult. However, when the original $\mathcal{C}_{kin}$ is non-singular we can take advantage of homology theory to provide an effective method for determining the path components. We assume $\mathcal{C}_{kin}$ is non-singular in what follows.

For any space with a finite number of path components we have that the zero-dimensional singular homology group, $H_0(X; \Re) = \Re^n$ where n is the number of components. Moreover, if Y is a closed manifold without boundary of dimension m and $Z \subset Y$ is a closed subspace in Y then we can apply Alexander–Poincaré duality which gives the result:

$$H_0(Y - Z; \Re) \cong H^m(Y, Z; \Re)$$

where $H^m(Y, Z; \Re)$ is the ordinary singular coholonomy of a pair. On the other hand, there is an exact sequence:

$$H^{m-1}(Y; \Re) \to H^{m-1}(Z; \Re) \to H^m(Y, Z; \Re) \to \Re \to 0$$

that allows us to calculate $H^m(Y, Z, \Re)$ once we understand the group $H^{m-1}(Z; \Re)$, since we already understand the group $H^{m-1}(Y; \Re)$ from Theorem 2. Moreover, when Z is a reasonable space—a real algebraic variety, for example—then it has the homotopy type of a complex of dimension $m - 1$, so we need to determine the top-dimensional cohomology group of the complex Z.

There are algorithms for doing this that build inductively on Theorem 2 (see the work of Vassiliev (1992) for some information on how this works) but they are computationally expensive and have not yet been entirely worked out for the case at hand where we believe that a significant amount of simplification of Vassiliev's general approach may be possible. In this case, Vassiliev's approach works as follows. When dealing with simple obstructions, for example a single point $\vec{X}$ in the workspace, then the space $\mathcal{C}_{kin,bad}$ consists of the union of varieties $\mathcal{V}_i$, where $\mathcal{V}_i$ is the set of configurations so that a point somewhere on link i touches $\vec{X}$. The variety $\mathcal{V}_i$ can be determined using the same techniques that gave Theorem 2. But to determine $\mathcal{C}_{kin,bad}$ and the relevant homology group, we also need to know the intersection $\mathcal{V}_{i_1} \cap \ldots \cap \mathcal{V}_{i_r}$.

5. Conclusion

The planning problem discussed here appears at the outset to be a simple path planning problem with holonomic constraints (Latombe 1991). However, such planning problems are not easily solved without knowledge of the global structure of the C-space of the system. In this paper, we used techniques from algebraic geometry and differential topology to determine the global structure of the C-spaces of all n-link closed chains with spherical joints (and n-link closed chains with revolute joints in the plane). Then we demonstrated the application of this knowledge to the design of a complete path planning algorithm. An interesting point is that, in many motion planning problems, the path existence and construction problems have the same difficulty. However, for the problem studied here, path existence is easier to determine than path construction.

The extension of our method to avoid collision, while theoretically possible, will face practical limitations as the number of links and obstacles increases, but these numbers are not yet known. On the other hand, our algorithm can serve quite well as a "local" planner in PRMs, in which candidate motions are generated and then tested for collisions. In such a scheme, there are actually two other motion generation methods that we might include: a local method such as that discussed in Section 4.4 and the global method of Lenhart and Whitesides (1994). The local method could be called first (since it produces the smoothest paths), and then if it fails, either the method presented here or the algorithm of Lenhart and Whitesides (1994) could be called.

In future work, we plan to extend the methods used here to problems in which it is important to avoid collisions, such as many occurring manufacturing and protein-folding applications.

Acknowledgments

The authors would like to thank Tony Maciejewski for motivating the comparison of our global algorithm to a local algorithm and Patrick G. Xavier for his contributions to the proof of algorithmic completeness. The research described in this paper was supported by the Research Foundations Program at Sandia National Laboratories. Sandia is a multi-program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000. The work of RJM was also partially supported by the Nation Science Foundation.

References

Alizade, R. I., and Sandor, G. N. 1985. Determination of the conditions of existence of complete crank rotation and of the instantaneous efficiency of spatial four-bar mechanisms. *Mechanism and Machine Theory* 20:155–163.

Baker, J. E. 1978. On the investigation of extrema in linkage analysis using screw system algebra. *Mechanism and Machine Theory* 13:333–343.

Canny, J. F. 1988. *The Complexity of Robot Motion Planning*. MIT Press.

Canny, J. F. 2002. Personal communication.

Carricato, M., and Parenti-Castelli, V. 2002. Singularity-free fully-isotropic translational parallel mechanisms. *International Journal of Robotics Research* 21(2):161–174.

Cherif, M., and Gupta, K. K. 1999. Planning quasi-static fingertip manipulation for reconfiguring objects. *IEEE Transactions on Robotics and Automation* 15(5):837–848.

Derby, S. 1980. The maximum reach of revolute jointed manipulators. *Mechanism and Machine Theory* 16:255–261.

Downing, D. M., Samuel, A. E., and Hunt, K. H. 2002. Identification of special configurations of the octahedral manipulator using the pure condition. *International Journal of Robotics Research* 21(2):147–160.

Farahat, A. O., Stiller, P. F., and Trinkle, J. C. August 1995. On the geometry of contact formation cells for systems of polygons. *IEEE Transactions on Robotics and Automation* 11(4):522–536.

Guillemin, V., and Pollack, A. 1974. *Differential Topology*. Prentice Hall.

Han, L., and Amato, N.M. 2000. A kinematics-based probilistic roadmap method for closed-chain systems. In *Proc. Workshop on Algorithmic Foundations of Robotics*, pp. 233–245.

Iarocci, M. October 1994. Three-dimensional graphics tool for motion design. Master's thesis, School of Computer Science, McGill University.

Kapovich, M. and Millson, J. 1995. On the moduli spaces of polygons in the euclidean plane. *Journal of Differential Geometry* 42:133–164.

Kapovich, M. and Millson, J. 1996. The simplectic geometry of polygons in euclidean space. *Journal of Differential Geometry* 44:479–513.

Kavraki, L. E., Lamiraux, F., and Holleman, C. 1998. Towards planning for elastic objects. In *Proceedings, Workshop on Algorithmic Foundations of Robotics*, pp. 313–327.

Kavraki, L. E., and Latombe, J.-C. May 1994. Randomized preprocessing for configuration space for fast path planning. In *Proc. IEEE International Conference on Robotics and Automation*, Vol. 3, pp. 2138–2145.

Kavraki, L. E., Švestka, P., Latombe, J.-C., and Overmars, M. H. August 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580.

Latombe, J.-C. 1991. *Robot Motion Planning*. Kluwer Academic Publishers.

LaValle, S. M., and Kuffner, J. J. 2001. Randomized kinodynamic planning. *International Journal of Robotics Research* 20(5):378–400.

LaValle, S. M., Yakey, J. H., and Kavraki, L. E. 1999. A probabilistic roadmap approach for systems with closed kinematic chains. In *Proc. IEEE International Conference on Robotics and Automation*, pp. 1671–1676.

Lenhart, W. J., and Whitesides, S. H. 1994. Reconfiguring simple polygons. *Discrete and Computational Geometry*.

Milgram, R. J., and Trinkle, J. C. 2002. The geometry of configuration spaces for closed chains in two and three dimensions. *Homology, Homotopy, and Applications*. Submitted.

Milnor, J. 1967. *Singular Points of Complex Hypersurfaces (Annals of Mathematics Studies No 61)*. Princeton University Press.

Overmars, M., and Svestka, P. 1994. A probabilistic learning approach to motion planning. In *Proc. Workshop on Algorithmic Foundations of Robotics*. A. K. Peters, Boston, MA.

Schwartz, J., Hopcroft, J., and Sharir, M. 1987. *Planning, Geometry, and Complexity of Robot Motion*. Ablex.

Son, W. May 1996. Dexterous manipulation planning for a planar whole-arm manipulator. Master's thesis, Texas A&M University Department of Computer Science.

Sugimoto, K., Duffy, J., and Hunt, K. H. 1982. Special configurations of spatial mechanisms and robot arms. *Mechanism and Machine Theory* 17:119–132.

Ting, K. L. December 1989. Mobility criteria of single loop n-bar linkages. *Trans. ASME, Journal of Mechanisms, Transmissions and Automation in Design* 111(4):504–507.

Ting, K. L., and Liu, Y. W. 1991. Rotatability laws for n-bar kinematic chains and their proof. *ASME Journal of Mechanical Design* 113(1):32–39.

Trinkle, J. C., Farahat, A. O., and Stiller, P. F. August 1995. First-order stability cells of active multi-rigid-body systems. *IEEE Transactions on Robotics and Automation* 11(4):545–557.

Trinkle, J. C., and Hunter, J. J. April 1991. A framework for planning dexterous manipulation. In *Proc. IEEE International Conference on Robotics and Automation*, pp. 1245–1251.

Trinkle, J. C., and Milgram, R. J. 2001. Motion planning for planar *n*-bar mechanisms with revolute joints. In *Proc. IEEE International Conference on Intelligent Robots and Systems*, pp. 1602–1608.

Trinkle, J. C., Ram, R. C., Farahat, A. O., and Stiller, P. F. May 1993. Dexterous manipulation planning and execution of an enveloped slippery workpiece. In *Proc. IEEE International Conference on Robotics and Automation*, Vol. 2, pp. 442–448.

Vassiliev, V. A. 1992. Complements of discriminants of smooth maps, topology, and applications. In *Transactions of Mathematics Monographs*, Vol. 98. American Mathematical Society.