Flat-State Connectivity of Linkages under Dihedral Motions*

Greg Aloupis, Erik D. Demaine, Vida Dujmović, Jeff Erickson
Stefan Langerman, Henk Meijer, Ileana Streinu, Joseph O'Rourke
Mark Overmars, Michael Soss, Godfried T. Toussaint

Abstract

We explore which classes of linkages have the property that each pair of their flat states—that is, their embeddings in $\mathbb{R}^2$ without self-intersection—can be connected by a continuous dihedral motion that avoids self-intersection throughout. Dihedral motions preserve all angles between pairs of incident edges, which is most natural for protein models. Our positive results include proofs that open chains with nonacute angles are flat-state connected, as are closed orthogonal unit-length chains. Among our negative results is an example of an orthogonal graph linkage that is flat-state disconnected. Several other results are obtained for other restricted classes of linkages. Many open problems are posed.

1 Introduction

Motivation: Locked Chains. There has been considerable research on reconfiguration of polygonal chains in 2D and 3D while preserving edge lengths and avoiding self-intersection. Much of this work is on the problem of which classes of chains can *lock* in the sense that they cannot be reconfigured to straight or convex configurations. In 3D, it is known that some chains can lock [CJ98], but the exact class of chains that lock has not been delimited [BDD$^+$01]. In 2D, no chains can lock [CDR00, Str00]. All of these results consider chains with *universal joints*.

Motivation: Protein Folding. The backbone of a protein can be modeled as a polygonal chain, but the joints are not universal; rather the bonds between residues form a nearly fixed angle in space. The study of such *fixed-angle* chains was initiated in [ST00], and this paper can be viewed as a continuation of that study. Although most protein molecules are linear polymers, modeled by open polygonal chains, others are rings (closed polygons) or star and dendritic polymers (trees) [SW88, FK97].

The polymer physics community has studied the statistics of "self-avoiding walks" [MOS90, MS93, Whi83], i.e., non-self-intersecting configurations, often constrained to the integer lattice. To generate these walks, they consider transformations of one configuration to another, such as "pivots" [Lal69] or "wiggling" [MJHS85]. Usually these transformations are not considered true molecular movements, often permitting self-intersection during the motion, and perhaps are better viewed as string edits.

*Research initiated at the 17th Winter Workshop on Computational Geometry, Bellairs Research Institute of McGill University, Feb. 1–8, 2002. Email addresses: athens@cs.mcgill.ca, edemaine@mit.edu, vida@jeff.cs.mcgill.ca, jeffe@cs.uiuc.edu, sl@jeff.cs.mcgill.ca, henk@cs.queensu.ca, streinu@cs.smith.edu, orourke@cs.smith.edu, markov@cs.uu.nl, soss@chemcomp.com, godfried@cs.mcgill.ca.

In contrast, this paper maintains the geometric integrity of the chain throughout the transformation, to more closely model the protein folding process. We focus primarily on transformations between planar configurations.

Fixed-angle linkages. Before describing our results, we introduce some definitions. A *(general) linkage* is a graph with fixed lengths assigned to each edge. The focus of this paper is *fixed-angle linkages*, which are linkages with, in addition, a fixed angle assigned between each pair of incident edges. We use the term *linkage* to include both general and fixed-angle linkages.

A *configuration* or *realization* of a general linkage is a positioning of the linkage in $\mathbb{R}^3$ (an assignment of a point in $\mathbb{R}^3$ to each vertex) achieving the specified edge lengths. The *configuration space* of a linkage is the set of all its configurations. To match physical reality, of special interest are *non-self-intersecting configurations* or *embeddings* in which no two nonincident edges share a common point. The *free space* of a linkage is the set of all its embeddings, i.e., the subset of configuration space for which the linkage does not "collide" with itself.

A configuration of a fixed-angle linkage must additionally respect the specified angles. The definitions of configuration space, embedding, and free space are the same. A *reconfiguration* or *motion* or *folding* of a linkage is a continuum of configurations. Motions of fixed-angle linkages are distinguished as *dihedral motions*.

Dihedral motions. A dihedral motion can be "factored" into *local dihedral motions* or *edge spins* [ST00] about individual edges of the linkage. Let $e = (v_1, v_2)$ be an edge for which there is another edge e_i incident to each endpoint v_i. Let Π_i be the plane through e and e_i. A *(local) dihedral motion about* e changes the dihedral angle between the planes Π_1 and Π_2 while preserving the angles between each pair of edges incident to the same endpoint of e. See Fig. 1. The edges incident to a common vertex in a fixed-angle linkage are moved rigidly by a dihedral motion. In particular, if the edges are coplanar, they remain coplanar.[1] If we view e and $e_1 \in \Pi_1$ as fixed, then a dihedral motion spins e_2 about e.

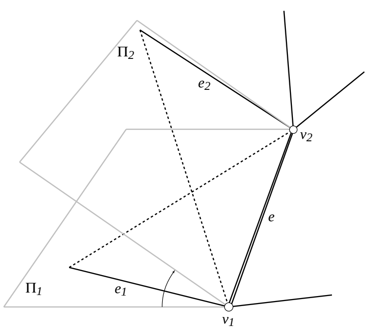

Figure 1: A local dihedral motion (spin) about edge e.

Flat-state connectivity. A *flat state* of a linkage is an embedding of the linkage into $\mathbb{R}^2$ without self-intersection. A linkage X is *flat-state connected* if, for each pair of its (distinct, i.e., incongruent) flat states X_1 and X_2, there is a dihedral motion from X_1 to X_2 that stays within the free space throughout. In general this dihedral motion alters the linkage to nonflat embeddings in $\mathbb{R}^3$ intermediate between the two flat states. If a linkage X is not flat-state connected, we say it is *flat-state disconnected.*

Flat-state disconnection could occur for two reasons. It could be that there are two flat states X_1 and X_2 which are in different components of free space but the same component of configuration space: they are connected by dihedral motions, but not by dihedral motions that avoid self-intersection. Or it could be that the two flat states are in different components of configuration space: there is no dihedral motion between them even permitting self-intersection. The former rea-

[1] Our definition of "dihedral motion" includes rigid motions of the entire linkage, which could be considered unnatural because a rigid motion has no local dihedral motions. However, including rigid motions among dihedral motions does not change our results. For a linkage of a single connected component, we can modulo out rigid motions; and for multiple connected components, we always pin vertices to prevent rigid motions.

son is the more interesting situation for our investigation; currently we have no nontrivial examples of the latter possibility.[2]

Results. The main goal of this paper is to delimit the class of linkages that are flat-state connected. Our results apply to various restricted classes of linkages, which are specified by a number of constraints, both topological and geometric. The topological classes of linkages we explore include general graphs, trees, chains (paths), both open and closed, and sets of chains. We sometimes restrict all link lengths to be the same; we call these *unit-length* linkages. We consider a variety of restrictions on the angles of a fixed-angle linkage, where the angle between two incident links is the smaller of the two angles between them within their common plane:

- An *obtuse* angle is one strictly greater than $90°$.
- An *orthogonal* angle is exactly $90°$.
- An *acute* angle is one strictly less than $90°$.
- A *nonacute* angle is either orthogonal or obtuse.

A chain *has a monotone state* if it has a flat state in which it forms a monotone chain in the plane. For sets of chains in a flat state, we *pin* each chain at one of the end links, keeping its position fixed in the plane.

In some cases we restrict the motions of a linkage in one of two ways. First, we may enforce that only certain edges permit local dihedral motion, in which case we call the linkage *partially rigid*. (Such a restriction also constrains the flat states that we can hope to connect, slightly modifying the definition of flat-state connected.) Second, we may restrict the motion to consist of a sequence of $180°$ edge spins, so that each move returns the linkage to the plane. Most of our examples of flat-state disconnected linkages are either partially rigid or restricted to $180°$ edge spins.

With the above definitions, we can present our results succinctly in Table 1.

2 Flat-State Disconnection

It may help to start with negative results, as it is not immediately clear how a linkage could be flat-state disconnected. Several of our examples revolve around the same idea, which can be achieved under several models. We start with partially rigid orthogonal trees, and then modify the example for other classes of linkages.

2.1 Partially Rigid Orthogonal Tree

An *orthogonal tree* is a tree linkage such that every pair of incident links meet at a multiple of $90°$. *Partial rigidity* specifies that only certain edges permit dihedral motions. Note that the focus of a dihedral motion is an edge, not the joint vertex.

Fig. 2(a–b) shows two incongruent flat states of the same orthogonal tree; we'll call the flat states $X_{(a)}$ and $X_{(b)}$. All but four edges of the tree are frozen, the four incident to the central degree-4 root vertex x. Call the 4-link branch of the tree containing a the *a-branch*, and similarly for the others. Label the vertices of the a-branch (a, a_1, a_2, a_3), and similarly for the other branches.

We observe three properties of the example. First, as mentioned previously, fixed-angle linkages have the property that all links incident to a particular vertex remain coplanar throughout all

[2] Our only examples involve restricting motion so much that the two flat states are completely rigid.

Constraints on Fixed-Angle Linkage				Flat-state connectivity
Connectivity	Angles	Lengths	Motions	
Open chain	—	—	—	?
	has a monotone state		—	?
	nonacute	—	—	Connected
	equal acute	—	—	Connected [ADM+02]
	each in $(60°, 90°]$	unit	—	Connected [ADM+02]
	—	—	180° edge spins	Disconnected
	orthogonal	—	180° edge spins	Connected
Set of chains, each	orthogonal	—	—	Connected
pinned at one end	orthogonal	—	partially rigid	Disconnected
Closed chain	—	—	—	?
	nonacute	—	—	?
	orthogonal	—	—	?
	orthogonal	unit	—	Connected
Tree	—	—	—	?
	orthogonal	—	—	?
	orthogonal	—	partially rigid	Disconnected
Graph	orthogonal	—	—	Disconnected

Table 1: Summary of results. The '—' means no restriction of the type indicated in the column heading. Entries marked '?' are open problems.

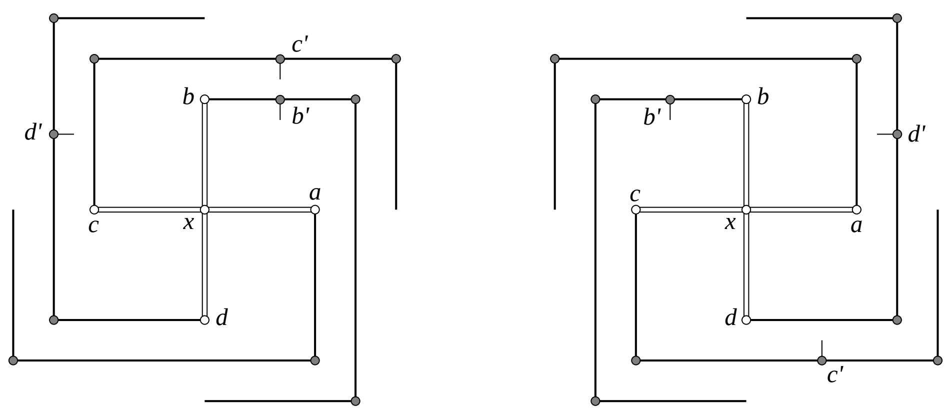

(a) *(b)*

Figure 2: Two flat states of a partially rigid orthogonal tree. The four open edges are the only ones not rigid, permitting dihedral motions.

dihedral motions. In Fig. 2, this means that $\{x, a, b, c, d\}$ remain coplanar; and we view this as the plane Π of the flat states under consideration. Note that, for example, a rotation of a about bd would maintain the 90° angles between all edges adjacent consecutively around x, but would alter the 180° angle between xa and xc, and thus is not a fixed-angle motion.

Second, the short links, or "pins," incident to vertices b', c', and d' must remain coplanar with their branch, because they are rigid. For example, the b' pin must remain coplanar with xb, for otherwise the rigid edge bb' would twist.

Third, $X_{(a)}$ and $X_{(b)}$ do indeed represent incongruent flat states of the same linkage. The purpose of the b' pin is to ensure that its relation to (say) the c' pin in the two states is not the same. Without the b' pin, a flat state congruent to $X_{(b)}$ could be obtained by a rigid motion of the entire linkage, flipping it upside-down. (The c' and d' pins play a second role, detailed below.) It is clear that state $X_{(b)}$ can be obtained from state $X_{(a)}$ by rotating the a-branch 180° about xa, and similarly for the other branches. Thus the two flat states are in the same component of configuration space, in that there are dihedral motions connecting them. We now show that no dihedral motions connect them without self-intersection, which establishes that they are in different components of the free space.

Theorem 2.1 *The two flat states in Fig. 2 of an orthogonal partially rigid fixed-angle tree cannot be reached by dihedral motions that avoid crossing links.*

Proof: Each of the four branches of the tree must be rotated 180° to achieve state $X_{(b)}$. We first argue that two opposite branches cannot rotate to the same side of the Π-plane, either both above or both below. Without loss of generality, assume both the a- and the c-branches rotate above Π. Then, as illustrated in Fig. 3, vertex a_1 must hit a point on the c_1c_2 edge, for the length aa_1 is the same as the distance from a to c_1c_2. Although this represents a self-intersection of the linkage, it does not represent a proper crossing of one link through another; rather it constitutes "grazing" contact. The purpose of the c' and d' pins is to cause a "true" crossing. For the a_1a_2 edge must pass directly through the c' pin (at c'), and, if they were physical links, disconnect it there. A variety of other rigid appendages could be added to the tree to create other types of crossings, e.g., two edges meeting at a point in the relative interior of each. Another approach is to vary the edge lengths of the tree to remove its symmetry.

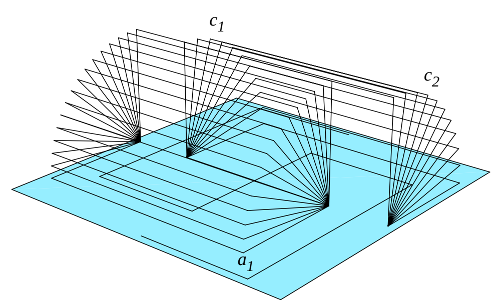

Figure 3: The a- and c-branches collide when rotated above.

Now we argue that two adjacent branches cannot rotate to the same side of Π. Consider the a- and b-branches, again without loss of generality. As it is more difficult to identify an exact pair of points on the two branches that must collide, we instead employ a topological argument. Connect a shallow rope R from a to a_3 underneath Π, and a rope S from b to b_3 that passes below R. See Fig. 4. In $X_{(a)}$, the two closed loops $A = (R, a, a_1, a_2, a_3)$ and $B = (S, b, b_1, b_2, b_3)$ are unlinked.

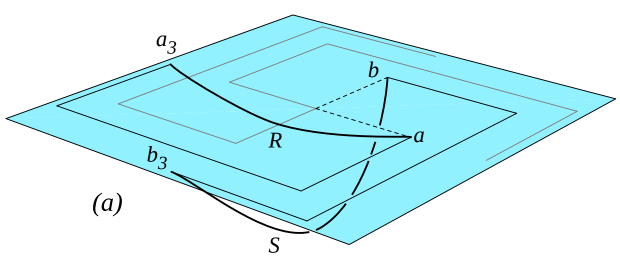

Figure 4: With the additions of the ropes R and S underneath, the a-chain is not linked with the b-chain in (a), but is linked in (b).

But in $X_{(b)}$, A and B are topologically linked. Therefore, it is not possible for the a- and b-branches to rotate above Π without passing through one another.

Finally, we apply the pigeon-hole principle. There are four branches, and two sides of Π. Thus one of the sides (say, above) must have (at least) two branches rotating through it. Whether these branches are opposite or adjacent, a collision is forced. $\square$

2.2 Orthogonal Graph

We can convert the partially rigid tree to a completely flexible graph by using extra "braces" to effectively force the partial rigidity. Thus we obtain

Corollary 2.2 *Two flat states of an orthogonal fixed-angle linkage cannot be reached by dihedral motions that avoid crossing links.*

2.3 Partially Rigid Pinned Chains

We can also convert the partially rigid tree in Fig. 2 into four partially rigid chains, each pinned at one endpoint just next to the central degree-4 vertex. Thus we obtain

Corollary 2.3 *The two flat states corresponding to Fig. 2 of four orthogonal partially rigid fixed-angle pinned chains cannot be reached by dihedral motions that avoid crossing links.*

We delay until Section 6 our examples of flat-state disconnection for 180° edge spins.

3 Nonacute Open Chains

We now turn to positive results, starting with the simplest and perhaps most elegant case of a single open chain with nonacute angles. After introducing notation in Section 3.1, we consider two algorithms establishing flat-state connectivity, in Sections 3.2 and 3.3.

3.1 Notation

We begin with some general notation for chains. An abstract polygonal chain $\mathcal{C}$ of n links is defined by its fixed sequence of link lengths, $(\ell_1, \ldots, \ell_n)$, and whether it is open or closed. For a fixed-angle chain, the $n-1$ or n angles α_i between adjacent links are also fixed. A realization C of a chain is specified by the position of its $n+1$ vertices: $v_0, v_1, \ldots, v_n$. If the chain is closed, $v_n = v_0$. The links or edges of the chain are $e_i = (v_{i-1}, v_i)$, $i = 1, \ldots, n$, so that the vector along the ith link is $v_i - v_{i-1}$. Another realization of $\mathcal{C}$ might be called C', with vertices v_i', etc. Alternatively, C_1, C_2, and so on, could be different realizations of $\mathcal{C}$. The plane in which a flat state C is embedded is called Π or the xy-plane.

3.2 Lifting One Link at a Time

The idea behind the first (unrestricted) algorithm is to lift the links of the chain one-by-one into a monotone chain in a vertical plane. Once we reach this *canonical state*, we can reverse and concatenate motions to reach any flat state from any other.

We begin by describing the case of orthogonal chains, as illustrated in Fig. 5, and the algorithm will generalize to arbitrary nonacute chains. The invariant at the beginning of each step i of the algorithm is that we have lifted the chain $e_1, \ldots, e_i$ into a monotone chain in a vertical plane, while the rest of the chain $e_{i+1}, \ldots, e_n$ remains where it began in the xy plane. Initially, $i = 0$ and the lifted chain contains no links, and we simply lift the first link e_1 to vertical by a 90° edge spin around the second link e_2. For general i, we first spin the lifted chain around its last (vertical) link e_i so that the vertical plane contains the next link to lift, e_{i+1}, and so that the chain $e_1, \ldots, e_{i+1}$ is monotone. Then we pick up e_{i+1} by a $90°$ edge spin around e_{i+2}. Throughout, the lifted chain remains monotone and contained in the positive-z halfspace, so we avoid self-intersection.

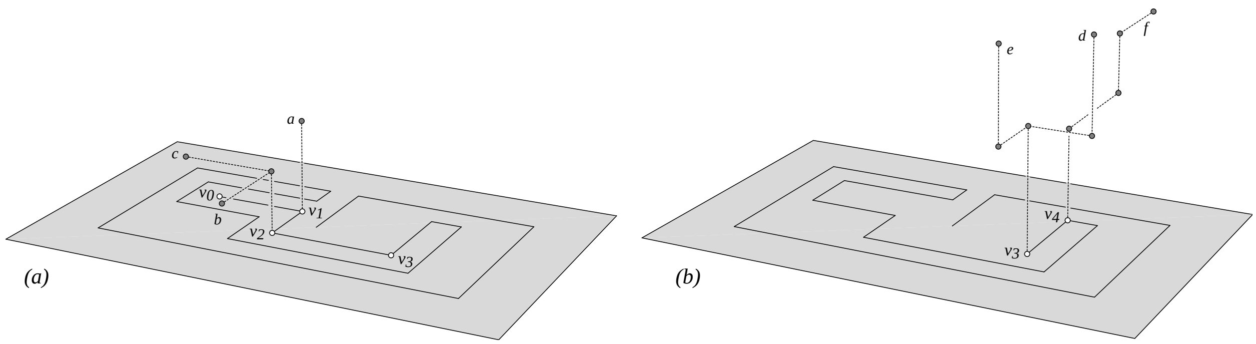

Figure 5: Picking up a planar orthogonal chain into a monotone canonical state. (a) Lifting edges $e_1 = (v_0, v_1)$ and $e_2 = (v_1, v_2)$: *a*, *b*, *c*. (b) Lifting edges e_3 and e_4: *d*, *e*, *f*.

Nonacute chains behave similarly to orthogonal chains, in particular, the canonical state is monotone, although it may no longer alternate between left and right turns. Now there may be multiple monotone states, and we must choose the state that is monotone in the z dimension. The key property is that, as the chain $e_1, \ldots, e_i$ rotates about e_{i+1}, the chain remains monotone in the z direction, so it does not penetrate the xy plane.

This algorithm proves the following result:

Theorem 3.1 *Any nonacute fixed-angle chain is flat-state connected.*

3.3 Lifting Two Links at a Time

The algorithm above makes at most 2 edge spins per link pickup, for a total of $2n$ edge spins to reach the canonical state, or $4n$ edge spins to reach an arbitrary flat state from any other. This bound is tight within an additive constant.

We can reduce the number of edge spins to $1.5n$ to reach the canonical state, or $3n$ to reach an arbitrary flat state, by lifting two edges in each step as follows. As before, in the beginning of each step, we spin the lifted chain $e_1, \ldots, e_i$ about the last link e_i to orient it to be coplanar and monotone with the next link e_{i+1}. Now we spin by $90°$ the lifted chain and the next two links e_{i+1} and e_{i+2} about the following link e_{i+3}, bringing e_{i+1} and e_{i+2} into a vertical plane, and tilting the lifted chain $e_1, \ldots, e_i$ down to a horizontal plane (parallel to the xy plane) at the top. Then we spin the old chain $e_1, \ldots, e_i$ by $90°$ around e_{i+1}, placing it back into a vertical plane, indeed the same vertical plane containing e_{i+1} and e_{i+2}, so that the new chain $e_1, \ldots, e_{i+2}$ becomes coplanar and monotone. We thus add two links to the lifted chain after at most three motions, proving the $1.5n$ upper bound; this bound is also tight up to an additive constant.

Corollary 3.2 *Any nonacute fixed-angle chain with n links can be reconfigured between two given flat states in at most $3n$ edge spins.*

4 Multiple Pinned Orthogonal Open Chains

In this section we prove that any collection of open, orthogonal chains, each with one edge pinned to the xy-plane, can be reconfigured to a canonical form, establishing that such chain collections are flat-state connected. We also require a "general position" assumption: no two vertices from different chains have a common x- or y-coordinate. Let C_i, $i = 1, \ldots, k$, be the collection of chains in the xy-plane. Each has its first edge pinned, i.e., v_0 and v_1 have fixed coordinates in the plane; but dihedral motion about this first edge is still possible (so the edge is not frozen). Call an edge parallel to the x-axis an *x-edge*, and similarly for *y-edge* and *z-edge*. The canonical form requires each chain to be a staircase in a plane parallel to the z-axis and containing its first (pinned) edge. If the first chain edge is a y-edge, the staircase is in a *yz quarter plane* in the halfspace $z > 0$ above xy; if the first chain edge is an x-edge, the staircase is in an *xz quarter plane* in the halfspace $z < 0$ below xy.

The algorithm processes independently the chains that are destined above or below the xy-plane, and keeps them on their target sides of the xy-plane, so there is no possibility of interference between the two types of chains. So henceforth we will concentrate on the chains C_i whose first edge is a y-edge, with the goal of lifting each chain C_i into a staircase S_i in a yz quarter plane. At an intermediate state, the staircase S_i is the portion of the lifted chain above the xy-plane, and C_i the portion remaining in the xy-plane. The *pivot edge* of the staircase is its first edge, which is a z-edge. Let $(\ldots, c_i, b_i, a_i)$ be the last three vertices of the chain C_i. Let a_i have coordinates (a_x, a_y); we'll use analogous notation for b_i and c_i. Vertex a_i at the foot of a staircase is its *base vertex* and the last edge of the chain, (b_i, a_i), is the staircase's *base edge*.

After each step of the algorithm, two invariants are reestablished:

1. All staircases for all chains are in (parallel) yz quarter planes;
2. The base edge for every staircase is a y-edge, i.e., is in the plane of the staircase.

We will call these two conditions the *Induction Hypothesis*.

The main idea of the algorithm is to pick up two consecutive edges of one chain, which then ensures that the next edge of that chain is a y-edge. The chain is chosen arbitrarily. To simplify the presentation, we assume without loss of generality that c_i is to the right of b_i. First, the staircases whose pivot's x-coordinates lie in the range $[b_x, c_x]$ are reoriented to avoid crossing above the (b_i, c_i) edge.

With S_i aligned with its base y-edge, the (a_i, b_i) edge can be picked up into a vertical plane without collision; see Fig. 6a. We now align S_i with (b_i, c_i), by "parting" the planes at b_x toward the left, laying all planes left of b_x down toward $-x$ (Fig. 6b), and then rotating S_i to be horizontal. Now we pick up (b_i, c_i) into a xz quarter plane, after laying down all planes right of c_x; see Fig. 6(c). Finally, reorient the xz-plane to be vertical and then restore all tilted planes to their yz orientation. We have reestablished the Induction Hypothesis. See Fig. 6(d).

Repeating this process eventually lifts every chain into parallel vertical planes, leaving only the first (pinned) y-edge of each chain in the xy-plane.

5 Unit Orthogonal Closed Chains

Our only algorithm for flat-state connectivity of closed chains is specialized to unit-length orthogonal closed chains. Despite the specialization, it is one of the most complex algorithms, and will only be sketched crudely in this abstract. The canonical form employed is an orthogonally convex polygon, justified by the first lemma:

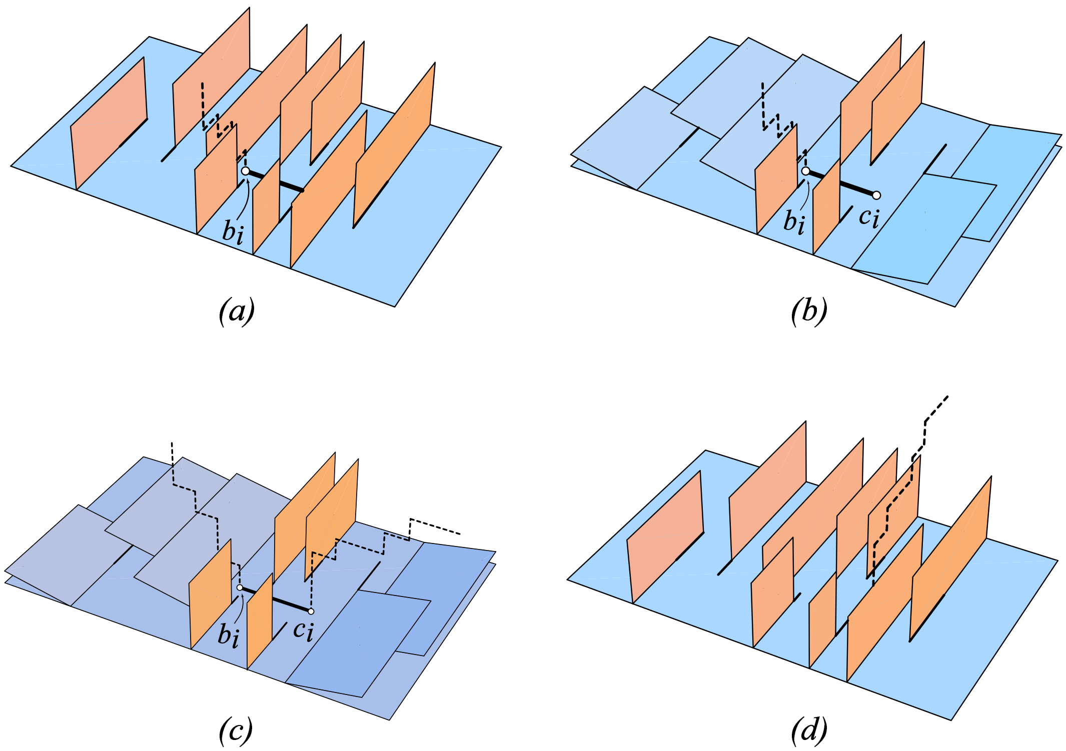

Figure 6: (a) First, y-edge (a_i, b_i) picked up; (b) Planes parted and flattened in preparation; (c) Two states of staircase shown: Aligned with the second, x-edge (b_i, c_i), and after pickup of that edge; (d) Staircase rotated into vertical plane, and flattened planes made upright.

Lemma 5.1 *Let C and D be two orthogonally convex embeddings of a unit-length orthogonal closed chain with n vertices. There is a sequence of edge spins that transforms C into D.*

The more difficult half is establishing the following:

Lemma 5.2 *Let C be a flat state of a unit-length orthogonal closed chain with n vertices. There is a sequence of edge spins that transforms C into an orthogonally convex embedding.*

This lemma is established through four motions, which we will only illustrate via an example chain, shown in Fig. 7(a). The canonical configuration will be achieved in a vertical plane, perpendicular to the horizontal plane containing the original chain. The chain is lifted starting from e_0, lifting subchains surrounding e_0 on either side. After each iteration k, let C_0 be one "half" of the lifted chain, before e_0, and C_1 the half after e_0. After iteration $k+1$, the lifted chain is $C' = C'_0 \cup e_0 \cup C'_1$, established by a combination of four motions: a *lifting move* (Fig. 8a,b), a *pocket flip* (Fig. 8c,d), *restoration moves* (Fig. 9), and *final moves*, executed if a special situation arises (Fig. 10).

These lemmas establish the following theorem:

Theorem 5.3 *Any unit-length orthogonal closed chain is flat-state connected.*

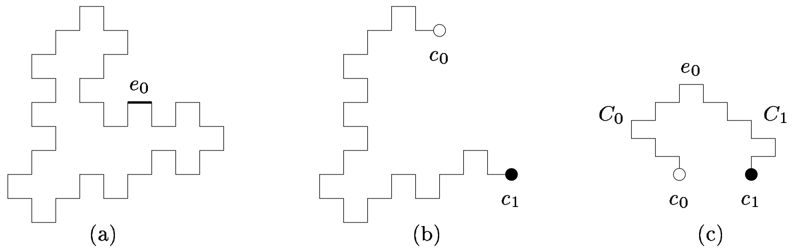

Figure 7: (a) Original C, with e_0 highlighted; (b) Unmoved portion after $k = 5$, (c) $C_0 \cup e_0 \cup C_1$ projected to $y = 0$ plane.

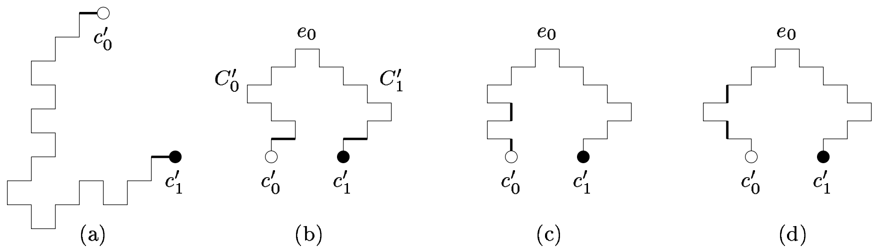

Figure 8: Lifting move, followed by two pocket flips. (a) The unmoved portion of C in the xy-plane; (b) Projection of C' onto the $y = 0$ plane; (c–d) Pocket flips applied to C'.

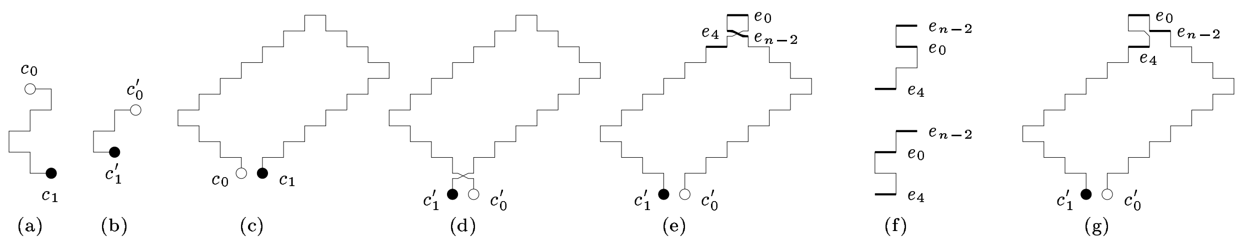

Figure 9: Restoration moves: (a–b) In xy-plane before and after lifting; (c–d) Lifted chain; (e) C_0' and C_1' are convex; (f) Flattening; (g) Planarity restored.

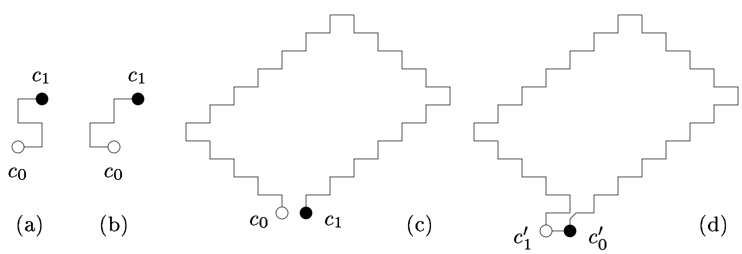

Figure 10: Final move in special separation-1 case.

6 180° Edge Spins

A natural restriction on dihedral motions is that the motion decomposes into a sequence of moves, each ending with the chain back in the xy plane—in other words, 180° edge spins. This restriction is analogous to Erdős flips in the context of locked chains [Tou99, ACD+00, BDD+01]. In this context, we can provide sharper negative results—general open chains can be flat-state disconnected—and slightly weaker positive results—orthogonal open chains are flat-state connected.

6.1 Restricted Flat-State Disconnection of Open Chains

We begin by illustrating the difficulty in reconfiguring open chains by 180° edge spins; see Fig. 11.

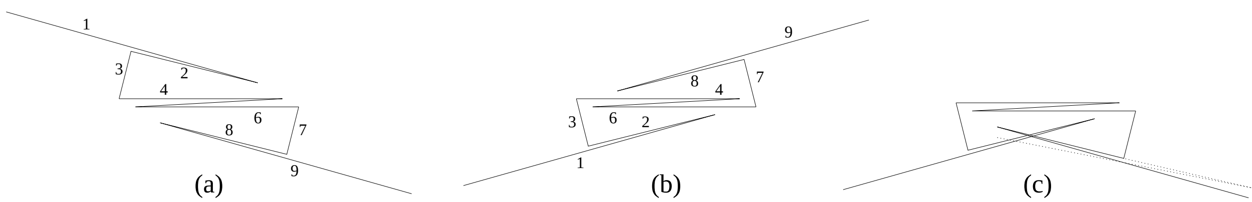

Figure 11: (a–b) Two flat states of a chain that cannot reach each other via a sequence of 180° edge spins. (c) Attempt at spinning about edge 4.

Spinning about edge 1 does nothing; spinning about edge 2 causes edges 1 and 3 to cross; spinning about edge 3 makes no important change to the flat state; spinning about edge 4 causes edges 2 and 8 to cross as shown in Fig. 11(c); spinning about edge 5 causes edges 4 and 6 to cross (in particular); and the remaining cases are symmetric. This case analysis establishes the following theorem:

Theorem 6.1 *The two incongruent flat states in Fig. 11(a–b) of a fixed-angle open chain cannot be reached by a sequence of* 180° *edge spins that avoid crossing links.*

6.2 Restricted Flat-State Connection of Orthogonal Open Chains

The main approach for proving flat-state connectivity of orthogonal chains is outlined in two figures: spin around a convex-hull edge if one exists (Fig. 12), and otherwise decompose the chain into a monotone (staircase) part and an inner part, and spin around a convex-hull edge of the inner part (Fig. 13). Such spins avoid collisions because of the empty infinite strips $R(e_1)$, $R(e_2)$, ... through the edges of the monotone part of the chain. In Fig. 13, the monotone portion of the chain is e_1, e_2, e_3, which terminates with the first edge e_3 that does not have an entire empty strip $R(e_3)$. Each spin of either type makes the chain more monotone in the sense of turning an edge whose endpoints turn in the same direction into an edge whose endpoints turn in opposite directions; hence, the number of spins is at most n. Using a balanced-tree structure to maintain information about recursive subchains, each step can be executed in $O(\log n)$ time, for a total of $O(n \log n)$ time. In addition, we show how the algorithm can be modified to keep the chain in the nonnegative-x halfspace with one vertex pinned against the $x = 0$ plane.

Theorem 6.2 *Orthogonal chains are flat-state connected even via restricted sequences of* 180° *spins that keep the chain in the nonnegative-x halfspace with one vertex pinned at $x = 0$. The sequence of $O(n)$ spins can be computed in $O(n \log n)$ time.*

7 Conclusion and Open Problems

See Table 1 for several open problems. In particular, these three classes of chains seem most interesting, with the first being the main open problem:

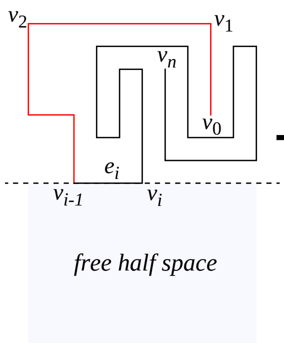

Figure 12: A dihedral rotation about a convex-hull edge resolves a violation of the canonical form.

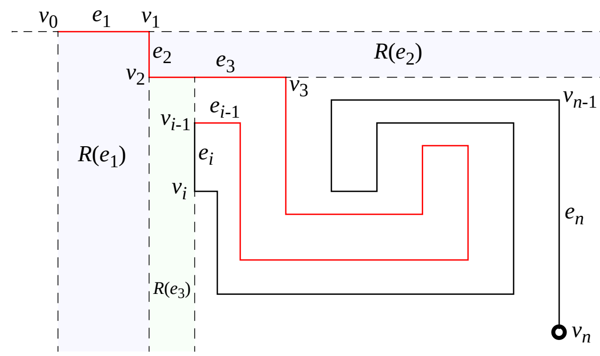

Figure 13: Determining the chain $e_1, e_2, \ldots, e_{i-1}$ that can be rotated about e_i.

1. Open chains (no restrictions).
2. Open chains with a monotone flat state.
3. Orthogonal trees (all joints flexible).

References

[ACD+00] Oswin Aichholzer, Carmen Cortés, Erik D. Demaine, Vida Dujmović, Jeff Erickson, Henk Meijer, Mark Overmars, Belén Palop, Suneeta Ramaswami, and Godfried T. Toussaint. Flipturning polygons. In *Proc. Japan Conf. Discrete Comput. Geom.*, Lecture Notes in Computer Science, Tokyo, Japan, November 2000. To appear in *Discrete Comput. Geom.*.

[ADM+02] Greg Aloupis, Erik D. Demaine, Henk Meijer, Joseph O'Rourke, Ileana Streinu, and Godfried Toussaint. Flat-state connectedness of fixed-angle chains: Special acute chains. In *Proc. 14th Canad. Conf. Comput. Geom.*, pages ??–??, August 2002.

[BDD+01] T. Biedl, E. Demaine, M. Demaine, S. Lazard, A. Lubiw, J. O'Rourke, M. Overmars, S. Robbins, I. Streinu, G. Toussaint, and S. Whitesides. Locked and unlocked polygonal chains in 3D. *Discrete Comput. Geom.*, 26(3):269–282, 2001.

[CDR00] R. Connelly, E. D. Demaine, and G. Rote. Straightening polygonal arcs and convexifying polygonal cycles. In *Proc. 41st Annu. IEEE Sympos. Found. Comput. Sci.*, pages 432–442. IEEE, November 2000.

[CJ98] J. Cantarella and H. Johnston. Nontrivial embeddings of polygonal intervals and unknots in 3-space. *J. Knot Theory Ramifications*, 7:1027–1039, 1998.

[FK97] Maxim D. Frank-Kamenetskii. *Unravelling DNA*. Addison-Wesley, 1997.

[Lal69] M. Lal. Monte Carlo computer simulations of chain molecules. *Molecular Physics*, 17:57–64, 1969.

[MJHS85] B. MacDonald, N. Jan, D. L. Hunter, and M. O. Steinitz. Polymer conformations through wiggling. *Journal of Physics A: Mathematical and General Physics*, 18:2627–2631, 1985.

[MOS90] N. Madras, A. Orlitsky, and L. A. Shepp. Monte Carlo generation of self-avoiding walks with fixed endpoints and fixed length. *Journal of Statistical Physics*, 58:159–183, 1990.

[MS93] Neal Madras and Gordon Slade. *The Self-Avoiding Walk*. Birkhâuser, 1993.

[ST00] M. Soss and G. T. Toussaint. Geometric and computational aspects of polymer reconfiguration. *J. Math. Chemistry*, 27(4):303–318, 2000.

[Str00] I. Streinu. A combinatorial approach to planar non-colliding robot arm motion planning. In *Proc. 41st Annu. IEEE Sympos. Found. Comput. Sci.* IEEE, November 2000. 443–453.

[SW88] C. E. Soteros and S. G. Whittington. Polygons and stars in a slit geometry. *Journal of Physics A: Mathematical and General Physics*, 21:L857–L861, 1988.

[Tou99] G. T. Toussaint. The Erdős-Nagy theorem and its ramifications. In *Proc. 11th Canad. Conf. Comput. Geom.*, pages 9–12, 1999.

[Whi83] S. G. Whittington. Self-avoiding walks with geometrical constraints. *Journal of Statistical Physics*, 30(2):449–456, 1983.