

Geometric and Computational Aspects of Molecular Reconfiguration

Michael Soss

School of Computer Science

McGill University

Montréal, Canada

July 2001

A thesis submitted to the
Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for
the degree of Doctor of Philosophy

Copyright ©2001 Michael Soss



National Library
of Canada

Acquisitions and
Bibliographic Services

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale
du Canada

Acquisitions et
services bibliographiques

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-75679-3

Canada

Abstract

We examine geometric problems of reconfiguring molecules modeled by polygons and polygonal chains in two and three dimensions. The molecule can be continuously reconfigured as long as the edge lengths are maintained and the object does not self-intersect.

We begin with a treatment of polygons in chapters 3 and 4. We prove that in two dimensions all convex configurations of a given polygon are reachable from any other and provide an efficient algorithm for convexifying planar monotone polygons. We also describe an algorithm to convexify three-dimensional polygons with simple projections.

We further prove that the motions to bring a convex configuration to another can be accomplished in three-dimensions by a series of restrictive moves called pivots, which are widely used in the polymer physics community. A pivot is a motion by which a contiguous subchain of the polygon is rotated about the line through its endpoints. A particular type of pivot on a planar polygon, known as a deflation, splits the polygon into two linearly separable subchains and rotates one by 180° . (A deflation can be thought of as the inverse of a flip as defined by Erdős [49].) Since a rotation of 180° about a line is equivalent to a reflection, a planar polygon remains planar after the deflation. In 1993 [164], Wegner conjectured that a planar polygon could admit only a finite number of deflations before self-intersecting. We disprove this conjecture by exhibiting a counterexample.

In chapters 5 through 7, we consider only polygonal chains in three-dimensions. As is consistent with the conventions of the chemistry and physics communities, we augment our model by allowing only those motions which maintain the angles at the vertices of the chain (and which also maintain the edge lengths and do not cause self-intersection). The chain can be reconfigured by only altering the dihedral torsion angles, which better reflects the fixed bond angles of a molecule.

We discuss several computational problems on this model. First, we describe an algorithm to determine the feasibility of the simplest possible motion—the alteration of only one dihedral angle—called a dihedral rotation. Since molecular reconfiguration is a complex process, the motion of a chain would involve a vast number of dihedral rotations. We prove lower bounds on the computational complexity of preprocessing a polygonal chain to determine the feasibility of multiple rotations.

We continue by demonstrating the intractability of several questions: (1) computing the maximum possible and minimum possible distance between the endpoints of the chain, (2) determining whether a chain can be reconfigured to certain canonical configurations, and (3) determining configurational chirality, that is, whether or not a chain can be reconfigured into its mirror image. This last problem has important implications for drug design, since the two mirror images of a molecule can have radically different properties.

Résumé

Nous étudions les problèmes géométriques liés à la reconfiguration de molécules modelées par des polygones et des chaînes polygonales en deux et trois dimensions. Une molécule peut être reconfigurée de façon continue si elle ne s'auto-intersecte pas et si les longueurs des arcs sont préservées.

Nous commençons par une étude sur les polygones dans les chapitres 3 et 4. Nous montrons qu'en deux dimensions toutes les configurations convexes d'un polygone sont accessibles depuis toutes autres et nous présentons un algorithme efficace pour convexifier tout polygone monotone planaire. Nous décrivons aussi un algorithme pour convexifier tout polygone tridimensionnel admettant une projection simple.

Nous prouvons également que les mouvements pour passer d'une configuration convexe à une autre peuvent être obtenus en trois dimensions par une série de mouvements restraints appelés les pivots, qui sont employés dans la communauté de physique des polymères. Un pivot est un mouvement consistant d'une rotation d'une sous-chaîne contigue du polygone autour de la droite définie par les extrémités de cette sous-chaîne. Un cas particulier de pivot sur un polygone planaire, appelé une déflation, sépare linéairement le polygone en deux sous-chaînes et en pivote une de 180° . (Une déflation peut être considérée comme l'inverse d'un *flip* selon Erdős [49].) Parce qu'une rotation de 180° autour d'une ligne est équivalente à une réflexion, un polygone planaire demeure planaire après une déflation. En 1993 [164], Wegner a fait la conjecture que tout polygone planaire ne peut admettre qu'un nombre fini de

déflations avant de s'auto-intersecter. Nous réfutons cette conjecture par un contre-exemple.

Du chapitre 5 au chapitre 7, nous considérons uniquement les chaînes polygonales en trois dimensions. Selon les conventions des communautés de chimie et de physique, nous augmentons le modèle en ne permettant que les mouvements qui préservent les angles aux nœuds de la chaîne (et qui également maintiennent les longueurs des arcs et ne causent aucune autointersection). De cette manière, la chaîne ne peut être reconfigurée qu'en altérant les angles dièdres de torsion. Ce modèle représente plus fidèlement les angles de liaisons fixes d'une molécule.

Nous discutons plusieurs problèmes algorithmiques liés à ce modèle. Tout d'abord, nous décrivons un algorithme pour déterminer la possibilité du mouvement le plus simple—le changement d'un seul angle dièdre—appelé une rotation dièdre. Parce que la reconfiguration moléculaire est un processus complexe, le mouvement d'une chaîne implique un nombre important de rotations dièdres. Nous prouvons des bornes inférieures sur la complexité du pré-traitement d'une chaîne polygonale pour déterminer la faisabilité de rotations multiples.

Nous continuons en montrant l'intracabilité de plusieurs problèmes: (1) calculer la distance la plus grande ou la plus petite entre les extrémités de la chaîne, (2) déterminer si une chaîne peut être reconfigurée dans certaines configurations canoniques, et (3) déterminer le chiralité d'une configuration, c'est à dire, si une chaîne peut être reconfigurée en son image miroir. Ce dernier problème a des implications importantes pour la conception de médicaments car les deux configurations chirales d'une molécule peuvent avoir des fonctions bien différentes.

Statement of Originality

Except for results whose authors are cited where first introduced, all elements of this thesis are original contributions to knowledge. Assistance has been received only where mentioned in the acknowledgements and in the following.

Portions of this thesis have appeared in sections of the following papers, and have been written with the assistance of coauthors. In all circumstances, either I personally had written the section of the original paper in question, or I have rewritten those portions which appear in this thesis.

- The work of sections 3.1 and 4.1 has appeared in: O. Aichholzer, E. Demaine, J. Erickson, F. Hurtado, M. Overmars, M. Soss, and G. Toussaint. Reconfiguring convex polygons. *Computational Geometry: Theory and Applications*, to appear in 2001.
- The work of section 3.2 has appeared in: T. Biedl, E. Demaine, S. Lazard, S. Robbins, and M. Soss. Convexifying monotone polygons. In *Lecture Notes in Computer Science no. 1741: Proceedings of the Tenth International Symposium on Algorithms and Computation*, pages 415–424.
- The work of section 4.2 has appeared in: T. Fevens, A. Hernández, A. Mesa, P. Morin, M. Soss, and G. Toussaint. Simple polygons with an infinite series of deflations. *Contributions to Algebra and Geometry*, 42(2):307–311, 2001.

- The work of section 4.3 has appeared in: J. Alberto-Calvo, D. Krizanc, P. Morin, M. Soss, and G. Toussaint. Convexifying polygons with simple projections. *Information Processing Letters*, to appear in 2001.
- The work of section 5.1 has appeared in: M. Soss and G. Toussaint. Dihedral reconfigurations of three-dimensional chains. Fourth Conference on Operations Research, Havana, 2000.
- The work of sections 5.1 and 6.5 has appeared in: M. Soss and G. Toussaint. Basic geometric problems on reconfiguring polymers. Novel Approaches in RNA Informatics, Montréal, 2000.
- The work of sections 5.1 and 6.5 has appeared in: M. Soss and G. Toussaint. Geometric and computational aspects of polymer reconfiguration. *Journal of Mathematical Chemistry*, 27(4):303–318, 2000.
- The work of section 5.2 is the subject of a manuscript in preparation: M. Soss, J. Erickson, and M. Overmars. On the difficulty of preprocessing for dihedral rotation queries.

Acknowledgements

I would first and foremost like to thank my supervisor, Godfried Toussaint, for introducing me to the subject of computational geometry and for suggesting the topic of this thesis. He has always been a source of creative ideas and possesses an infectious passion for research. Through his financial support I have traveled to many workshops and conferences, and have had the opportunity to work with researchers from around the world. My doctoral studies would have been neither as enjoyable nor as fruitful without his generosity.

I am gracious to many colleagues I have met during the course of my graduate studies whose influences have helped me in numerous ways. I would like to explicitly thank those whose collaborations appear in this thesis: Oswin Aichholzer, Therese Biedl, Prosenjit Bose, David Bremner, Jorge Alberto Calvo, Erik Demaine, Jeff Erickson, Tom Fevens, Antonio Hernández, Ferrán Hurtado, Danny Krizanc, Sylvain Lazard, Antonio Mesa, Pat Morin, Mark Overmars, Steve Robbins, Vera Sácristan, and Godfried Toussaint. The contributions of each are as indicated on page vi. (I also thank Jit, Erik, Jeff, Danny, Mark, Pat, Greg Aloupis, and Stefan Langerman, who worked with me for days at one of Godfried’s workshops trying to solve open problem 3 of chapter 8 but to no avail.) I am also grateful to David Avis, Luc Devroye, and Naoki Katoh for their guidance during the writing of section 6.4.

Thanks also to the members of the Computational Geometry Laboratory—Greg Aloupis, David Avis, Luc Devroye, Vida Dujmović, Bohdan Kaluzny, Tallman Nkgau,

Steve Robbins, Godfried Toussaint, and Sue Whitesides—for providing a stimulating research environment.

Above all, I thank my wife, Aleksandra, for her encouragement and support.

Contents

Abstract	ii
Résumé	iv
Statement of Originality	vi
Acknowledgements	viii
A Note on Stereograms	xx
1 Introduction	1
1.1 Reconfiguration in Geometry	2
1.1.1 Unfolding linkages in the plane	3
1.1.2 Convexifying polygons with flips and flipturns	6
1.1.3 Unfolding linkages in three dimensions	9
1.1.4 Kinematics	12
1.1.5 Rigidity theory	13
1.2 Reconfiguration in Polymer Physics	14
1.3 Reconfiguration in Biochemistry	20
1.3.1 Topological descriptions of molecular shape	22
1.3.2 Protein folding	23
1.3.3 Chirality	24

	xi
1.3.4 Prion diseases	26
1.4 Geometry, Polymer Physics, or Chemistry?	28
1.5 Overview	28
2 Notation and Preliminaries	30
2.1 Models of Computation	30
2.2 Trigonometric Computations	31
2.3 Elementary Geometry	33
3 Reconfiguring Polygons in Two Dimensions	36
3.1 Reconfiguring Between Convex Configurations	37
3.2 Convexifying Monotone Polygons	40
4 Reconfiguring Polygons in Three Dimensions	50
4.1 Reconfiguring Convex Polygons with Pivots	50
4.2 Deflations	56
4.2.1 A counterexample to Wegner’s conjecture	57
4.2.2 Mouth flips	62
4.3 Convexifying Polygons with Simple Projections	63
4.3.1 Springs	64
4.3.2 The Spring Algorithm	66
5 Dihedral Rotations	73
5.1 Determining the Feasibility of Dihedral Rotations	76
5.2 Preprocessing Chains for Dihedral Rotations	84
5.2.1 Single dihedral rotation queries	86
5.2.2 Multiple dihedral rotation queries	90
6 Endpoint-to-Endpoint Distance and Canonical Configurations	97
6.1 Preliminaries	99

	xii
6.2 Span of Two-Dimensional Configurations	100
6.3 Span of Three-Dimensional Configurations	111
6.4 Approximating the Maximum Span	115
6.5 Flattening	125
7 Chirality	131
7.1 Quasi-rigidity	136
7.2 Fundamental Quasi-rigid Structures	142
7.3 Complex Quasi-rigid Structures	155
7.4 Achirality is NP-hard	160
8 Conclusion	168

List of Figures

1	The principle behind stereograms.	xx
2	<i>Stereogram</i> : Three overlapping disks.	xxi
1.1	Reconfiguring a linkage by a rotation at the red joint.	2
1.2	A chain linkage, a cyclic linkage, and a tree linkage.	2
1.3	Opening a convex chain linkage.	4
1.4	Ruler Folding.	4
1.5	The locked tree linkage of Biedl et al.	5
1.6	Flipping all pockets of a polygon.	7
1.7	Béla Nagy's polygon where flipping all pockets leads to nonsimplicity.	7
1.8	Joss and Shannon's quadrilateral which can require arbitrarily many flips.	8
1.9	A deflation.	8
1.10	A deflated polygon.	9
1.11	A flipturn.	9
1.12	The locked knitting needles chain.	10
1.13	The locked polygon of Biedl et al.	10
1.14	The locked hexagon of Cantarella and Johnston.	11
1.15	The locked hexagon of Toussaint.	11
1.16	Performing a flip in three-space.	11
1.17	A four-bar mechanism with the edge $\overline{ab}$ fixed.	13

1.18	Peaucellier's Inversor.	13
1.19	Range of orientations allowed with a fixed joint-angle.	15
1.20	The motions allowed in Verdier's model.	17
1.21	A locked chain in Verdier's planar model.	17
1.22	Two locked chains in Verdier's three-dimensional model.	17
1.23	A pivot (of Lal's model) on the red edge.	18
1.24	Examples of pivots on the planar lattice.	19
1.25	Asymmetry of the carbon atom.	20
1.26	Ball-and-stick model of a protein. Proteins are composed of a long chain called a backbone, drawn with white bonds, and various side chains, shaded.	21
1.27	The two possible configurations of cyclohexane. Only the carbon atoms are drawn; the hydrogen atoms are omitted.	22
1.28	<i>Left:</i> (S)-thalidomide. <i>Right:</i> (R)-thalidomide.	24
1.29	A saccharin molecule.	25
1.30	<i>Stereogram:</i> Mislow and Bolstad's rubber glove.	26
1.31	<i>Left:</i> Proposed structure of common protein PrP ^C [81, 82]. <i>Right:</i> Proposed structured of disease-causing prion PrP ^{Sc} [82].	27
2.1	A polygon and its reflex vertices.	34
3.1	A quadrilateral of vertices labeled $+, -, +, -$	38
3.2	The motion of lemma 5.	39
3.3	<i>Left:</i> An x -monotone polygon. <i>Right:</i> A polygon that is not x -monotone.	41
3.4	A monotone polygon decomposed into two chains.	41
3.5	The definition of k as per step 2b.	43
3.6	The motion of step 2(c)iii.	43
3.7	Quadrilateral of lemma 9.	45

3.8	The region W .	47
3.9	The definitions of a, b, c, d , and m .	48
3.10	The configuration when v_k straightens.	49
4.1	A pivot on $\overline{v_i v_j}$.	51
4.2	The transformation illustrated in figure 3.2 but accomplished with three pivots, shown chronologically from left to right. <i>Top row</i> : Bird's eye view. <i>Bottom row</i> : Oblique view.	51
4.3	Illustration of the pivots used in theorem 5. <i>Top row</i> : Bird's eye view. <i>Bottom row</i> : Oblique view.	52
4.4	Two pivots performed to reconfigure a parallelogram that is not a rhombus.	54
4.5	A deflation.	57
4.6	A deflated polygon.	57
4.7	A self-intersecting embedding of $abcd$.	58
4.8	The proof of lemma 15.	59
4.9	<i>Left</i> : A mouth. <i>Right</i> : A mouth flip.	62
4.10	A star-shaped polygon that admits no mouth-flips.	62
4.11	A polygon that admits infinitely many mouth-flips.	63
4.12	A polygon with a simple orthogonal projection to the xy -plane.	64
4.13	Coplanar edges whose projection is a single edge.	65
4.14	A spring-unfolding of a planar chain.	65
4.15	Pulling a spring until it straightens.	66
4.16	Step 1a. Reconfiguring a 3-spring into a 2-spring.	68
4.17	Step 1b. Reconfiguring a 2-spring into a 1-spring.	69
4.18	Step 1c. Reconfiguring a 1-spring. <i>Top</i> : Step 1(c)i. <i>Bottom</i> : Step 1(c)ii.	69

5.1	Ball-and-stick model of a protein. Proteins are composed of a long chain called a backbone, drawn with white bonds, and various side chains, shaded.	73
5.2	Range of motion allowed with a fixed joint-angle.	74
5.3	A dihedral rotation of angle ϕ at $\overline{uv}$. (<i>Shown in stereo in figure 5.4.</i>)	75
5.4	<i>Stereogram</i> : A dihedral rotation at $\overline{uv}$. (The stationary part of the chain has been chosen to be planar in order to emphasize the three-dimensional motion; this is not a requirement of the dihedral rotation.)	75
5.5	<i>Left</i> : The tree for a set with no repeated elements, $\{2, 1, 5, 8, 7, 6\}$. <i>Right</i> : The tree after a rotation of π	79
5.6	<i>Left</i> : The tree for a set with a repeated element, $\{2, 1, 5, 8, 7, 5\}$. The endpoints of the overlapped edges are encircled. <i>Right</i> : the tree after a rotation of π . The bold segment indicates the collision.	80
5.7	Oblique view of the base and left gadget for the set $(2, 1, 5, 8, 7, 5)$. There are two pairs of edges at height $z = 5$ corresponding to the repeated element. (<i>Shown in stereo in figure 5.8.</i>)	81
5.8	<i>Stereogram</i> : Oblique view of the base and left gadget for the set $(2, 1, 5, 8, 7, 5)$	82
5.9	View of the chain from above ($z = +\infty$). Not to scale; the y -dimension is magnified for clarity.	82
5.10	Oblique view of the chain for the set $(2, 1, 5, 8, 7, 5)$. If a dihedral rotation is performed at $\overline{uv}$, the edges indicated at the arrows, which correspond to the repeated element 5, will collide. (<i>Shown in stereo in figure 5.11.</i>)	83
5.11	<i>Stereogram</i> : Oblique view of the chain for the set $(2, 1, 5, 8, 7, 5)$	84
5.12	<i>Stereogram</i> : Converting the chain into a polygon.	85
5.13	Reducing 3SUM' to Single Dihedral Rotation Query.	88

5.14	An edge spin at the indicated vertical edge.	89
5.15	<i>Left:</i> There exist no $a' \in A'$ and $c' \in C'$ such that $a' - 2b' + c' = 0$. <i>Right:</i> The 3SUM' is solved, as $a' - 2b' + c' = 0$	89
5.16	<i>Left:</i> The gadget for setting the distance between two consecutive teeth on the left comb. <i>Center:</i> The gadget for consecutive elements of the staircase. <i>Right:</i> The gadget for consecutive elements of the right comb.	93
5.17	The chain created in the proof of theorem 17.	93
5.18	The chain, moved into position for the newly revealed A' , B' , and C' . <i>(Shown in stereo in figure 5.19).</i>	95
5.19	<i>Stereogram:</i> The chain, moved into position for the newly revealed A' , B' , and C'	96
6.1	Definitions of ℓ_i and θ_i	99
6.2	The chain constructed in the proof of theorem 18.	103
6.3	Two links of length l differing in orientation by θ	103
6.4	Triangle $\triangle abc$ embedded between two planes.	105
6.5	The distance $ x'y' $ in relation to $ xy $	106
6.6	The chain created in the proof of theorem 20.	109
6.7	A chain for which the minimum planar span is less than one.	109
6.8	A chain which has no planar span less than one.	110
6.9	<i>Stereogram:</i> A chain whose maximum span occurs at a nonplanar con- figuration.	111
6.10	The maximum planar span of the chain in figure 6.9.	112
6.11	Relationship between ϕ_i , ϕ_{i-1} , and θ_i	113
6.12	Program to compute the maximum span.	113
6.13	The chain D of the proof of theorem 21.	115
6.14	The two possible planar configurations of the starting subchain con- structed in the proof of theorem 24.	127

6.15	The subchain starting at vertex a	127
6.16	The entire chain constructed in the proof of theorem 24.	127
6.17	An instance that can be flattened.	128
6.18	<i>Left:</i> An x -monotone chain. <i>Right:</i> A chain that is not x -monotone. .	129
6.19	A chain with a monotone planar configuration.	129
7.1	A rubber glove.	133
7.2	The gadget to decide Achirality.	134
7.3	<i>Stereogram:</i> Side view of the gadget to decide Achirality. (Not to scale.)	135
7.4	A staple and hook.	136
7.5	A linked staple and hook.	136
7.6	A staple-and-hook triangle.	138
7.7	The discrepancy in position of the two configurations is the largest discrepancy in position of any point, the distance between d and d' . .	139
7.8	The distance between χ and χ' is the largest discrepancy in position of any point on the chain, minimized over all global rotations and trans- lations. From its position in figure 7.7, χ' has been rotated through the third dimension (about the edge $\overline{bc}$ in figure 7.7), rotated slightly clockwise, then translated onto χ	140
7.9	A tightly locked staple-and-hook triangle.	140
7.10	Step-by-step construction of a graduated ruler.	143
7.11	A graduated ruler with smaller staple-and-hook edges.	143
7.12	Labeling of the chain in the proof of lemma 30.	143
7.13	Illustration of x and h	144
7.14	Legend for graduated rulers.	145
7.15	Building the frame of the planar net.	145
7.16	Labeling used in the proof of lemma 31.	146
7.17	A spacer.	150

7.18	Two rulers joined by three spacers.	151
7.19	Stapling the spacers to complete the planar net.	151
7.20	<i>Stereogram</i> : A cap.	152
7.21	A staple-edge.	153
7.22	Labeling of the staple-edge in the proof of lemma 34.	153
7.23	A staple-edge (magenta) attached to two hooks of a graduated ruler.	155
7.24	A spatula formed by a planar net and a staple-edge.	156
7.25	The maximum slope of the staple-edge.	156
7.26	A double spatula.	157
7.27	Simplified drawing of a double spatula.	157
7.28	A double spatula with one planar net rotated.	158
7.29	Computing the maximum distance between the two nets when one is rotated.	159
7.30	<i>Stereogram</i> : A capped edge.	159
7.31	The auxiliary chain and cap.	160
7.32	The connection of the top half of the double spatula and one auxiliary chain. (Not to scale.)	161
7.33	The double spatula connected to one auxiliary chain. (Not to scale.)	162
7.34	View from above of the long staple-edge of the double spatula and one auxiliary chain. The planar nets of the double spatula are not shown.	163
7.35	Placement of the spacers of the double spatula.	164
7.36	Each link of the central chain is trapped from above by at least two spacers of the double spatula.	165
7.37	The chain constructed to prove that Achirality is NP-hard. (Not to scale.)	167
8.1	A chain and two planar configurations.	170
8.2	Moving between the two configurations.	170

A Note on Stereograms

Because the majority of this thesis deals with polygons and chains in three-dimensions, several of the figures herein are stereograms. A stereogram works by the principle illustrated in figure 1. By crossing one's eyes, the left eye and the right eye see different images. When the images overlap in the brain, one creates the illusion of depth by manipulating the distances between various components. In figure 1 the right protrusion will seem farther away than the left since one would have to relax the eyes slightly to bring it into focus, as would be consistent with looking at an object far away.

Figure 2 provides a practice image. Begin to cross your eyes, and you'll see double. Four copies of the image will be visible, since you are seeing two with each eye. Continue to cross your eyes until two of the images overlap, resulting in three images.




Figure 1: The principle behind stereograms.




Figure 2: *Stereogram*: Three overlapping disks.

The middle one is a composite image; your left eye is focusing on the right image of the stereogram and your right eye on the left image. If you can focus your eyes on the middle image of the three disks, they will appear to overlap one another.

Seeing stereograms quickly and easily requires practice. Because many readers may have difficulty seeing stereograms, whenever a three-dimensional image can be implied by a standard drawing, it is drawn by a standard single image. In section 4.1, the drawings are so vital to understanding the motions involved that they have been illustrated by projections to different viewpoints as opposed to stereograms.

Only in cases where the third dimension contains valuable information has a stereogram been drawn, and in most cases a standard alternate image has been presented. Even if one cannot see the three-dimensional image in a stereogram, either of the images intended for each eye provides a two-dimensional projection. All stereograms are clearly marked with the label, “*Stereogram*,” as in figure 2.

Chapter 1

Introduction

Most objects of interest in the world can be manipulated and altered. Toys have moving parts; compasses can be flexed at their hinges to draw circles of various radii; this very book can be opened and closed by bending it at its spine.

Suppose we constructed a physical model of a graph using stiff rods for edges and connecting them at the vertices with freely moving (ball-and-socket) joints. The graph could then be *reconfigured*, or moved, in any continuous manner from one *configuration*, or embedding, to a second configuration, so long as no physical constraint of the structure were violated. In this case, we could insist that no joint was broken, that the lengths of the bars remained fixed, and that no two bars intersected throughout the motion. We refer to such a flexible structure as a *linkage*. Alluding to this physical example, we will often refer to the edges of the graph as *links* and the vertices as *joints*. One possible reconfiguration of a linkage is illustrated in figure 1.1.

This thesis deals with reconfiguring objects from one position to another. We will begin by examining problems long studied by the geometry community, and move to problems of reconfiguration in physics and biochemistry.




Figure 1.1: Reconfiguring a linkage by a rotation at the red joint.




Figure 1.2: A chain linkage, a cyclic linkage, and a tree linkage.

1.1 Reconfiguration in Geometry

The geometry community has long considered the problem of determining if it is possible to reconfigure a linkage from one configuration (embedding) to another while maintaining the edge lengths and avoiding self-intersections during the motion. We can simplify this problem by instead asking if a linkage can be brought to some canonical form. Suppose a linkage can be brought from a configuration A to some canonical form C , and also can be brought from a configuration B to C . Because all the motions are reversible, it follows that the linkage can be continuously reconfigured from A to B . Thus the problem of unfolding linkages is born. In particular, this question is most often asked for the three classes¹ of linkages in figure 1.2:

- Given a chain linkage, can it be straightened?

¹There is conflicting terminology over these classes in the geometry and engineering disciplines. We will use solely the terminology of the former. A *chain*, often referred to as an *open chain* in engineering, denotes a simple path with two endpoints. Likewise, we will use the terms *cycle* or *polygon* to denote what is often referred to as a *closed chain* in engineering.

- Given a cyclic (polygon) linkage, can it be convexified (brought to a convex position)?
- Given a tree linkage, can it be flattened?

The answers, of course, depend on the motions allowed (i.e., reconfiguration with respect to a specific set of operations) and the space in which the linkage lies. We briefly examine the history surrounding these problems.

1.1.1 Unfolding linkages in the plane

Augustin Cauchy appears to have been the first mathematician to consider the problem of straightening a chain linkage. As a lemma for his celebrated theorem regarding the rigidity of polyhedra, he attempted to prove the following in 1813.

Lemma 1. (due to Cauchy [29]) *If some joints of a convex chain are opened, the distance between the endpoints increases and the chain remains convex.*

We consider a joint to be opening if its angle approaches π , in other words, if the two edges incident to it approach a straight angle. Such a motion is illustrated in figure 1.3, where the red vertices have been opened. Cauchy’s original proof was incorrect, although its flaw was unnoticed until 1934 when Steinitz and Rademacher [146] published a correction. Lyusternik [98] independently published a similar proof in 1966, and a short, elegant proof was found by Schoenberg and Zaremba [142] in the following year. As a lemma in a paper on intersections of planes with polytopes, O’Rourke [116] has since further generalized Cauchy’s lemma to include specific non-convex configurations. The survey paper by Toussaint [156] provides Cauchy’s proof and a short history of the problem.

In the computer science community, reconfiguration of chain linkages first appeared in the context of modeling the possible motions of robot arms. John Hopcroft, Deborah Joseph, and Sue Whitesides [78] considered the case where the linkage is permitted




Figure 1.3: Opening a convex chain linkage.




Figure 1.4: Ruler Folding.

to self-intersect, since the limbs of a robot arm can pass under one another. They designed algorithms for deciding whether two configurations are mutually accessible when the chain is in some bounded region. They further proved that determining whether a chain could be flattened onto a line segment of specified length, a problem known as Ruler Folding (illustrated in figure 1.4), is NP-complete.

The question of whether a planar chain linkage can be straightened or a planar cyclic linkage can be convexified, without self-intersections, has arisen several times since. Robert Connelly, Erik Demaine, and Günter Rote [37] researched the origins of the problem and found that it was independently posed by Stephen Schanuel in the early 1970's, by Ulf Grenander [67] in 1987, by William Lenhart and Sue Whitesides [94, 95, 169] in 1991, and by Joseph Mitchell in 1992. Grenander, Chow, and Keenan [66] also considered a cyclic linkage whose angles are fixed but whose lengths are variable, and proved that any embedding could be reconfigured into any other.

Biedl et al. [21] demonstrated that not all tree linkages can be unfolded by exhibiting the *locked* (unable to be unfolded) tree of figure 1.5, but the question for chain linkages and cyclic linkages received a flurry of interest in the computational geom-




Figure 1.5: The locked tree linkage of Biedl et al.

etry community and remained open for several years. Its difficulty led researchers to consider special cases. In 1998, Everett, Lazard, Robbins, Schröder, and Whitesides [51] showed that all star-shaped polygons can be convexified by an expansive motion, and in the following year Biedl, Demaine, Lazard, Robbins, and Soss [22] provided an algorithm to convexify monotone polygons. (The latter result is described in section 3.2.)

Attempts were made to construct chains that were believed to be locked [110], but unfolding motions were found for each example. The difficulty of these folding problems was further illustrated by Arkin, Fekete, Mitchell, and Skiena [9], who demonstrated the intractability of folding problems which stem from wire bending and sheet metal folding, and again recently by Arkin, Bender, Demaine, Demaine, Mitchell, Sethia, and Skiena [8].

Using tools in rigidity theory Connelly, Demaine, and Rote [37] were finally able to prove in January, 2000, that chains and cycles can be unfolded. However, describing their motions involves integrating vector fields, and thus their algorithm is not programmable in conventional models of computation. Several months later,

Streinu [149] discovered a simpler method which computes $O(n^2)$ motions to convexify a polygon, although the complexity of computing each motion is still unclear.

A general algorithm for motion planning, which can determine if an object can be brought from a starting configuration to a target configuration, was developed in 1983 by Schwartz and Sharir [144], but the complexity is doubly exponential in the degrees of freedom. In the case of linkages, this is at least as large as the number of joints. This result complemented the proof by Reif [128] in 1979 that deciding if an arbitrary hinged object (which could include polyhedral segments) could be moved from a starting configuration to a target is PSPACE-complete. Clearly more specialized algorithms are necessary in the case of simple linkages.

1.1.2 Convexifying polygons with flips and flipturns

In 1935, Paul Erdős [49] posed the following problem.

Given any simple polygon P which is not convex, draw the smallest convex polygon P' which contains P . This convex polygon P' will contain the area P and certain additional areas. Reflect each of these areas with respect to the corresponding added side, thus obtaining a new polygon P_1 . If P_1 is not convex, repeat the process, obtaining a new polygon P_2 . Prove that after a finite number of such steps a polygon P_n will be obtained which will be convex.

This operation, which we will call a *flip*, is illustrated in figure 1.6. With the aid of modern geometric vocabulary, we can reword Erdős' problem statement to be more accessible.

Given a nonconvex simple polygon, consider its convex hull (drawn in red). Subtracting the polygon from its convex hull yields several polygons called *pockets* (shaded). Reflect each of these pockets across its *lid*, that




Figure 1.6: Flipping all pockets of a polygon.




Figure 1.7: Béla Nagy's polygon where flipping all pockets leads to nonsimplicity.

is, the edge it shares with the convex hull. We call this operation a *flip*.

Prove that after a finite number of flips the polygon will be convex.

The first of many proofs of Erdős' conjecture was published four years later by Béla Nagy [111]. Nagy showed that flipping all pockets simultaneously might lead to a nonsimple polygon, as in figure 1.7, and modified the problem so that only one pocket is flipped at a time.

This problem had since been independently discovered and solved by several mathematicians: in 1957 by Reshetnyak [131] and by Yusupov [172]; in 1959 by Kazarinoff and Bing [23, 87, 88]; in 1973 by Joss and Shannon [70]; in 1981 by Kaluza [86]; in 1993 by Wegner [164]; and in 1999 by Biedl et al. [20]. The paper by Toussaint [157] provides a detailed account.

Joss and Shannon [70] demonstrated that although finitely many flips suffice for convexification, the number required cannot be bounded by a function of the number of edges in the polygon. They presented the quadrilateral in figure 1.8. Note that




Figure 1.8: Joss and Shannon's quadrilateral which can require arbitrarily many flips.




Figure 1.9: A deflation.

there is only one flip possible at any one time since there is at most one reflex vertex. For any integer k , one can make the smallest edge tiny enough as to require at least k flips to convexify the polygon. Wegner [164] and Biedl et al. [20] also discovered the same quadrilateral.

Wegner [164] also posed the inverse problem. Select a line that passes through the polygon in exactly two vertices, and reflect one of the two subchains across this line. If the resulting polygon does not self-intersect, the operation is called a *deflation*. This is the inverse of the *flip* operation, since the reflected subchain is now a pocket of the convex hull of the resulting polygon and the original line is its lid. Alluding to the concept that convex polygons do not admit flips, Wegner defined a *deflated polygon* as a polygon that does not admit any deflations. He conjectured that all polygons would be deflated after finitely many deflations. Figure 1.9 illustrates the deflation operation, and figure 1.10 illustrates a polygon which is deflated. We exhibit a counterexample to Wegner's conjecture in section 4.2.

Joss and Shannon [70] also considered a variant of the flip operation. Instead of reflecting a pocket, one can also rotate the pocket by 180 degrees about the center of its lid, as illustrated in figure 1.11. We call this operation a *flipturn*. Note that after a flipturn, each edge is rotated by 180 degrees. Unlike in a flip, each edge retains




Figure 1.10: A deflated polygon.




Figure 1.11: A flipturn.

its original slope. If we consider each edge as a vector, a flipturn has the effect of permuting the order of the edges. Since each flipturn increases the area of the polygon, we will never visit the same permutation twice. Therefore the polygon will be convex after at most $(n - 1)!$ flipturns, where n is the number of edges of the polygon. Joss and Shannon conjectured that $n^2/4$ flipturns would suffice. In 2000, Ahn et al. [2] proved that a polygon is convex after any sequence of $\lceil n(s - 1)/2 \rceil - s$ flipturns, where s is the number of distinct slopes of the edges. For arbitrary polygons, this value is $(n^2 - 3n)/2$. In the special case of orthogonal polygons, Aichholzer et al. [4] proved that there always exists a convexifying sequence of $5(n - 4)/6$ flipturns.

1.1.3 Unfolding linkages in three dimensions

The geometry of three-space is quite different from that of two-space, and indeed not all three-dimensional chains can be straightened. Cantarella and Johnston [27] and Biedl et al. [20] independently proved that the chain of figure 1.12 is locked. This




Figure 1.12: The locked knitting needles chain.




Figure 1.13: The locked polygon of Biedl et al.

chain is often referred to as *knitting needles* due to its resemblance to the tools of the same name. If the first and last links are long enough, the endpoints cannot be brought near the other four joints. One can then prove that the linkage behaves much like a trefoil knot.

In the case of three-dimensional polygons, the goal is to place the polygon into a planar convex configuration. Clearly a knotted polygon cannot be convexified, but Biedl et al. [20] demonstrated a class of locked unknots by joining two knitting needles to form the polygon of figure 1.13. Cantarella and Johnston [27] also proved that there exist locked unknots and presented the class of locked unknotted hexagons illustrated in figure 1.14. They conjectured that the configuration space of unknotted hexagons had three classes: the convex hexagon and the left-handed and right-handed versions of their locked polygon. Toussaint [158] discovered the hexagon illustrated in figure 1.15, bringing the conjectured number of classes to five.

Biedl et al. [20] have shown that planar polygons can be convexified by motions in three-space. They rediscovered the flip operation with the idea of rotating a pocket




Figure 1.14: The locked hexagon of Cantarella and Johnston.




Figure 1.15: The locked hexagon of Toussaint.

in three-space about its lid, illustrated in figure 1.16. Realizing that a polygon may require an unbounded number of flips with respect to its number of edges, they described a linear-time algorithm to convexify a planar polygon using more complicated motions. Aronov, Goodman, and Pollack [10] generalized the result to include crossing polygons several months later. (They considered a reconfiguration of a crossing polygon to be valid if it introduces no new crossings during the motion.)

We can generalize the flip operation by selecting two points of a three-dimensional polygon and rotating one of the subchains about the line through these two points. This operation was first used in 1945 by Choquet [33] in an application known as curve stretching. A curve c' is a *stretched* version of c if for every two points on c , the arc length between them is maintained and the Euclidean distance between them is either maintained or increased. One can imagine c as a rope, and c' as a position




Figure 1.16: Performing a flip in three-space.

of the same rope, but “spread apart.” A polygon which is convexified by flips is therefore a stretched version of the original polygon, since each flip either maintains or increases the pairwise distances between vertices.

Sallee [140] proved that for every three-dimensional curve, there exists a stretched version which is planar and convex. Robertson and Wegner have also studied this operation (which they refer to as *inflation*) in the plane [135, 136, 164], and Wegner has explored stretching curves on the sphere [165, 166]. Millett [108] has also used similar motions in three dimensions to convexify knots, but for his purposes allowed the polygon to self-intersect during the motion.

In 1999, Cocan and O’Rourke [35] proved that all chains, polygons, and trees can be unfolded in dimensions four and higher.

1.1.4 Kinematics

The study of kinematic linkages appears to have been first mathematically codified in 1874 by the engineer Franz Reuleaux in his work *Theoretische Kinematik* [132]. Kinematics is the study of motion, and in the context of linkages, is the study of how certain joints move in concert with the motion of other joints.

Reuleaux defined a *mechanism* as “[a] closed kinematic chain, of which one link is thus made stationary [132, p. 47].” A major component of his work is in describing the mathematical relationships between the curves drawn by each joint as one moves the linkage in a specified fashion. Even the simple four-bar linkage [134] shown in figure 1.17 was the subject of intense study in the nineteenth century. (This linkage was often called a three-bar linkage since there was no consensus at that time on whether to count the stationary edge as part of the structure.) Here one can easily compute the path traveled by the points p and q as the angle θ is altered.

Perhaps the most famous early linkage is Peaucellier’s Inversor of figure 1.18, developed in 1864 by the French engineer of the same name. If the points a and b are



Figure 1.17: A four-bar mechanism with the edge $\overline{ab}$ fixed.

Figure 1.18: Peaucellier's Inversor.

fixed in space such that distance between a and b is the same as that between b and p , then as p moves along the red circle, q moves along the blue line. Thus the linkage *inverts* radial motion to linear motion and vice versa.

The paper by Farouki [53] and the book by McCarthy [103] provide an in-depth historical background and further detail the connections between geometry and kinematics.

1.1.5 Rigidity theory

The field of rigidity theory, the mathematical study of flexibility of structural frameworks, has seen a renaissance in recent years. This field is most concerned with de-

termining the degrees of motion available in a graph whose edges behave as distance constraints.

Whiteley [168] has performed much research in the field of applying rigidity theoretic tools to discover the degrees of freedom available in molecular models. Jacobs et al. [85], and Thorpe et al. [155] have further used such tools to determine rigid substructures in proteins.

The scope of this thesis will be more algorithmic than topological, and thus we will not delve far into rigidity theory here. Books by Graver, Servatius, and Servatius [65] and by Thorpe and Duxbury [154] provide a thorough introduction to the field.

1.2 Reconfiguration in Polymer Physics

Flexible polymer molecules are long chains of atoms, often called monomers. These chains are held together by chemical bonds, about which a certain amount of rotation is generally possible. In essence, the model is a long chain linkage. As with the linkages discussed above, the movements allowed give rise to numerous possible configurations, some highly elongated and some tightly coiled [17]. The enormous number of possibilities suggested a statistical approach to describe the possible configurations of the chain. It was discovered that many physical properties of a polymer such as how it diffuses and scatters light can be related to statistical averages such as the mean squared distance between the endpoints of the chain. The first reference to this concept appears to be in a short letter by Henry Eyring [52] in 1932.

In 1934, Werner von Kuhn [90] proposed a model in which a long chain of unit length edges is allowed to assume any configuration, possibly self-intersecting. To obtain a random configuration, he chose each link to assume any orientation with uniform probability. Since this is equivalent to a random walk in three dimensions [31], Kuhn was able to show without great difficulty that the mean squared end-to-end distance is $O(n\ell^2)$, where ℓ is the length of each link.




Figure 1.19: Range of orientations allowed with a fixed joint-angle.

Unfortunately, Kuhn's model did not perform very well in predicting polymer properties. Benoit [17] and Taylor [151] independently corrected his model so that successive elements of the polymer chain were not randomly oriented. They recognized that not only should the lengths of the links relate to the lengths of the monomers, but the angles at the joints of the linkage should relate to the bond angles between monomers. Rather than allowing a link to assume any orientation, it was bound to those that preserve the bond angle with the previous link. This resulted in a range illustrated in figure 1.19. It seems that neither physicist was aware of Eyring's work, as his model is the same as theirs.

In 1943, Kuhn and Kuhn [91] recognized that the normal random-walk model was severely deficient, since it implied that the atoms could overlap. They realized that self-intersections, or even positions in which atoms are very close together, should not be allowed. Computing the mean squared distance between the endpoints under the restriction against two atoms occupying the same region of space is often referred to as the excluded volume problem.

Subsequent attempts by Flory [57], Hermans [75], Grimley [69], and Bueche [25] to treat the excluded volume problem as a small perturbation of the normal random walk were not entirely successful. Other attempts, in which the distribution function was

sought [173] contained complicated functions, the evaluations of which were highly infeasible and did not provide an accurate model [163].

In 1961, Sykes [150] proposed modeling polymers as self-avoiding walks on a cubic lattice and began his work by counting the number of configurations for small numbers of links. His attempts to compute this number for large walks did not meet much success [45], and even today this well-known problem remains unsolved and a focus of mathematical research [101]. It became evident that Monte Carlo methods would prove useful in finding an approximation to the mean squared end-to-end distance.

Early methods of generating self-avoiding random walks took quite long to finish [137, 163]. The difficulty was due to the approach; a walk was generated by choosing a random direction at each step. If a self-intersection arose, the entire configuration was discarded and the process begun anew. Note that it is not correct to simply remove the self-intersecting step and choose a new direction, as this would erroneously increase the likelihood of the portion of the configuration attained up to that point.

In 1969, Peter Verdier [160] proposed a faster method of computing random walks. Rather than randomly walking on the lattice and hoping for a sequence of non-intersecting steps, he began with a valid walk and performed a series of random reconfigurations. Any single joint could be moved to a vacant lattice point so long as the edges remained unit length and the chain did not self-intersect. These conditions imply that the only motions possible are those illustrated in figure 1.20.

An interesting problem arose concerning this new method of generating walks. It is not obvious that there exists a sequence of reconfigurations by which any configuration can be brought to any other, and thus it is not obvious whether all valid configurations have a positive probability of being generated. In the physics community, if all configurations are attainable, the model is said to be *ergodic*. Not surprisingly, Verdier's model lacks ergodicity. In the plane, the chain in figure 1.21




Figure 1.20: The motions allowed in Verdier's model.




Figure 1.21: A locked chain in Verdier's planar model.

is locked; in three-dimensions, the chains in figure 1.22 are locked. Therefore neither could be generated by Verdier's simulations unless it was selected as the starting configuration. Other simple methods such as removing the first link and replacing it on the other side [89, 162] (known as *reptation* due to its resemblance to the motion of a slithering snake) also suffered from locked configurations.




Figure 1.22: Two locked chains in Verdier's three-dimensional model.




Figure 1.23: A pivot (of Lal’s model) on the red edge.

The same year as Verdier, Moti Lal [92] proposed a different operation for randomly reconfiguring a walk which yields an ergodic model. His method was to select an edge at random, defining two subchains, then to reflect one of these subchains across the line through the edge. This operation has been since called a *pivot* throughout the literature. If the resulting chain intersected itself, the pivot was rejected and a new edge was randomly chosen. Otherwise, the procedure was repeated from the new configuration. Lal defined his model on the hexagonal lattice, although the pivot operation works just as well in a variety of spaces. Figure 1.23 illustrates a pivot on the red edge, moving half of the chain to the position indicated in blue.

This algorithm was independently reinvented in 1976 by Olaj and Pelinka [113] and again nine years later by MacDonald et al. [99]. The pivot algorithm remains, in the words of Madras and Sokal [102], “the most efficient algorithm yet invented for estimating the mean squared end-to-end distance of a random walk on the lattice.” Its efficiency has led it to be a major focus of study during the past fifteen years [28, 99, 100, 101, 102, 129, 130], although expanded by an operation. In addition to Lal’s original pivot, one may select two points on the chain and cut out the portion between them. The portion is then replaced, either reflected about a line through its cut points or rotated by 180 degrees, as in figure 1.24. As long as the resulting chain does not self-intersect and lies on the lattice, the pivot is accepted. This new operation easily




Figure 1.24: Examples of pivots on the planar lattice.

generalizes to polygons and is remarkably similar to flips and flipturns.

In order to demonstrate ergodicity, Madras and Sokal [102] proved that in any dimension all lattice chains can be straightened by pivots. Madras, Orlitsky, and Shepp [100] extended this result to include convexifying polygons. The physics community was proving that chains and polygons on the lattice could be unfolded with flips and flipturns!

Many in the physics community were concerned that the restriction of the model to the lattice was too far removed from reality. So called *off-lattice* or *continuum* models were constructed. John Curro [41, 42] extended Lal's idea of a pivot to continuous space by selecting an edge of the chain and rotating one portion of the chain about a line through the edge. As in Lal's model, the joint-angle at the edge stays fixed, and if the result of the motion is not self-intersecting, it is accepted. This motion is often referred to as a *dihedral rotation*. This model has been used with success by several researchers since its introduction [41, 42, 60, 104, 147, 148]. Note that the range of possible configurations is nearly identical to that used by Benoit and Taylor! The model with fixed edge lengths and fixed joint-angles is once again brought to the forefront. It also arises in several other settings which we will soon describe.




Figure 1.25: Asymmetry of the carbon atom.

1.3 Reconfiguration in Biochemistry

The field of stereochemistry, the study of three-dimensional molecular geometry, was first explored by Louis Pasteur [118] in 1860. Chemists had observed since the early 1800's that when polarized light passed through solutions of organic compounds, in some samples the light was diffracted to the right and in others to the left. Pasteur theorized that this was due to asymmetry in the molecule. Thirteen years later, Johannes Wislicenus [170] theorized that the difference between the two samples was due to a difference in the spatial arrangement of their atoms, and in the following year Joseph-Achilles Le Bel [15] and Jacobus van't Hoff [77] independently attributed this phenomenon to the asymmetry of the carbon atom. The carbon atom admits four bonds in a tetrahedral pattern. If the groups bonding with the carbon atom are distinct, there are two orientations possible, illustrated in figure 1.25. The field of stereochemistry was not well received until Wislicenus [171] demonstrated that a variety of unknown relationships could be explained by this theory.

Ever since, chemists have realized that the three-dimensional shapes of molecules determine how they interact with the ambient world. Drug design and discovery is becoming increasingly based on identifying molecular structure [12], as shape determines how molecules can dock with another molecule's receptive sites [54, 93]. In particular, the structure of a protein determines which molecules it can manipulate and therefore its biological function [80].




Figure 1.26: Ball-and-stick model of a protein. Proteins are composed of a long chain called a backbone, drawn with white bonds, and various side chains, shaded.

One popular molecular model is the ball-and-stick model of figure 1.26. With stiff bonds and rather rigid bond lengths and bond angles, the configuration² of a molecule is essentially described by dihedral angles [64, 141, 143]. (Given a four-vertex subchain (t, u, v, w) , the *dihedral angle* at $\overline{uv}$ is the angle between the planes determined by $\triangle tuv$ and $\triangle uvw$.) In the case of a protein, the molecule is flexible at the dihedral angles ϕ and ψ about the bonds adjacent to certain carbon atoms known as α -carbons [143]. As in the off-lattice model described in section 1.2, the bond lengths and bond angles are fixed, but the molecule is free to rotate about several of its bonds. Thus the only motions allowed are dihedral rotations. For example, the five rightmost atoms of the molecule in figure 1.26 are free to spin around the bond to their left, thereby changing the dihedral angle ψ . This model has been successful in solving a variety of problems in pharmacology [34] and biochemistry [141].

This model has been studied numerous times from a mathematical viewpoint [85, 155]. One of the earlier studies was in 1987 by Richard Randell [127]. His research provided topological explanations of chemical phenomena, such as demonstrating that the configuration space of cyclohexane, the six-atom carbon ring C_6H_{12} , has two

²In the chemistry community, the term “conformation” is often used in conjunction with “configuration,” where the two terms have slightly different meanings. In keeping with the terminology of the geometry community, we will exclusively use the latter.




Figure 1.27: The two possible configurations of cyclohexane. Only the carbon atoms are drawn; the hydrogen atoms are omitted.

distinct path-connected components [126]. (Randell later discovered that the same calculations had been made from an algebraic viewpoint a century earlier [138].) Cyclohexane takes the shape of an equilateral hexagon with all bond angles fixed at about 109.5° . The *boat* configuration, on the left of figure 1.27, is flexible and can be reconfigured to other rotationally symmetric embeddings. The *chair* configuration on the right is rigid and cannot be moved. Randell also proves other interesting results, including that a carbon ring of fewer than 11 atoms must be unknotted. Emiris and Mourrain [48] have recently generalized some of Randell's work to larger molecules.

1.3.1 Topological descriptions of molecular shape

Molecular shape has also often been described from a topological viewpoint [106]. The most common approach is to determine if the molecular graph can be continuously deformed from one embedding to another. If two graphs are not topologically equivalent, then certainly one embedding cannot be reconfigured to another when geometric and chemical constraints are added. This approach has led to the application of knot theory to molecular structure [55, 56, 107, 161]. The results therein have direct implications in problems where the molecule in question is a ring, such as mitochondrial DNA [59]. Since we will deal mostly with chain molecules, topological approaches will not be as helpful.

1.3.2 Protein folding

Proteins, like polymers, can be described as a linear chain, although of amino acids instead of monomers. As is evident in figure 1.26, the linear backbone of a protein (denoted by white bonds) has a greater contribution to its overall shape than the much smaller side chains (shaded). One of the remarkable properties of a protein is its ability to fold up into a low-energy resting shape, called its native configuration. Interestingly enough, a protein folds into the same shape every time. It has long been believed by many that the native configuration of a protein is its configuration of minimum energy [7]. Several models have been invented and since studied to describe the intramolecular forces that determine stable configurations [3, 30, 44, 72, 73], but finding the minimum energy configuration of proteins in each model has been shown to be NP-hard [19, 38, 58, 159].

Determining how proteins fold would allow us to compute their shapes and biological functions from their amino acid sequence. The implications for biochemistry and pharmacology are enormous, and a great deal of research has focused on this very problem. The ability of proteins to quickly fold to their native configurations despite the intractability of the protein folding problem has come to be known as Levinthal's paradox [96]. Several theories have arisen to explain this, including that the native configuration is not in fact the minimum energy configuration but rather a local extremum. This approach places the focus of study onto the folding itself by seeking any stable structure rather than the globally minimum energy configuration. Attempts to simulate the folding process have had moderate success [145] although the running times are usually quite large.

For brevity, we will not delve into the numerous details of protein folding. Neumaier [112] has written a thorough survey which is accessible to those of a mathematical background.



Figure 1.28: *Left*: (S)-thalidomide. *Right*: (R)-thalidomide.

1.3.3 Chirality

Molecules fold into a stable structure of low energy, but there may be several symmetric low energy states. A molecule may be synthesized and fold into one configuration, but its mirror image is symmetric and therefore equally as stable. If the configuration of a molecule differs from its mirror image, the molecule is said to be *chiral*.

The chirality of a molecule has profound effects on its interactions. The molecule on the right of figure 1.28, (R)-thalidomide, was prescribed to pregnant women in the 1960's as a cure for morning sickness. Due to its chirality, (R)-thalidomide and its mirror image, (S)-thalidomide, are equally stable configurations. During synthesis both forms were created with equal probability, resulting in a racemic (50:50) mixture. While (R)-thalidomide is effective against morning sickness, (S)-thalidomide is responsible for causing horrific birth defects [153].

In contrast, the artificial sweetener saccharin is illustrated in figure 1.29. Note the horizontal plane of symmetry through the molecule. The saccharin molecule is identical to its mirror image, and is said to be *achiral*. The native configuration of saccharin is therefore unique, and there is no danger of an alternate but symmetric form.

To underscore the importance of chirality studies, it suffices to realize that nearly half of today's synthetic pharmaceuticals are chiral. Ninety percent of these are sold as racemic mixtures due to the enormous cost of producing a pure substance [47].

Yet, determining the chirality of a rigid molecule is not sufficient to explain its





Figure 1.29: A saccharin molecule.

properties. When a compound exists as two symmetric configurations that interconvert rapidly, these forms become equally populated [152]. Recent experiments have shown that even if (R)-thalidomide were taken in its pure form, the energy barrier between the two forms is low enough that chemicals in the body would quickly reconfigure (R)-thalidomide into (S)-thalidomide [153].

Thus one is not only interested in detecting whether a rigid object is chiral, but also in determining whether a molecule can be continuously reconfigured into its mirror image. A wonderful example of this phenomenon is due to Mislow and Bolstad [109], who in 1955 exhibited the molecule of figure 1.30. The rigid representation of their molecule is chiral but it can easily be reconfigured into its mirror image. The central bond is rigid, but if middle portion is rotated (pivoted, in fact) about the red bonds by 90 degrees, the mirror image is attained. It is rather interesting that despite this reconfiguration, the molecule cannot be brought to a configuration admitting a plane of symmetry! Mislow and Bolstad called this type of molecule a *rubber glove*. One can turn a left-handed glove inside out to make a right-handed glove, but it is impossible to bring the glove to a symmetric configuration.

We will examine computational problems concerning chirality and rubber gloves in chapter 7.

Figure 1.30: *Stereogram*: Mislow and Bolstad's rubber glove.

1.3.4 Prion diseases

A survey of reconfiguration in biochemistry would be incomplete without a short discussion of prions.

Several diseases, including a debilitating sheep ailment known as scrapie, its variant Bovine Spongiform Encephalopathy (also known as Mad Cow Disease), human Creutzfeldt-Jakob Disease, Cystic Fibrosis, and countless others had until recently been far beyond the understanding of biologists.

After early experiments in the mid-1900's found scrapie to be transmissible among sheep [39] and goats [40, 119], the search began to find the infectious agent. Since all known pathogens at the time were micro-organisms, the community searched for a virus. Attempts to destroy the infectious properties of samples by radiation [5, 6], heat [83], and detergents [83, 84] all proved ineffective. It was discovered that a protein always accompanied the agent, and it was conjectured that the virus was somehow able to hide inside a thick protein membrane [63].

A mathematician named J. S. Griffith [68] theorized that the agent was a self-




Figure 1.31: *Left*: Proposed structure of common protein PrP^{C} [81, 82]. *Right*: Proposed structured of disease-causing prion PrP^{Sc} [82].

replicating protein, and he outlined three possible such mechanisms. This idea was completely novel, and since Griffith's work was purely theoretical and had no experimental basis, it attracted little attention. Not for fifteen years did someone else arrive at the same conclusion, Stanley Prusiner [121, 122, 123] in 1982. Prusiner termed the agent a *prion*, for *infectious protein*.

Prusiner's conjecture was that the infectious protein already existed inside healthy organisms, but with a different configuration. Illustrated on the left of figure 1.31 is a common protein normally produced by the body, PrP^{C} . A proposed structure of the disease-causing agent, the prion PrP^{Sc} , is illustrated on the right [81, 82]. These two molecules have the same chemical makeup, differing solely in their foldings [13, 81]. The main difference between the two is the conversion of two α -helices to β -sheets [13, 117], a structural change which greatly increases the stability of the configuration [36], which further explains its resistance to heat and radiation.

PrP^{Sc} is not a living organism, so it cannot propagate by reproduction as can viruses or bacteria. When it comes into contact with the naturally occurring PrP^{C} , it reconfigures the latter into its own configuration! The two disease-causing molecules contact two common molecules and create four, and so on, resulting in exponential

growth. The radical idea of a self-replicating protein gained acceptance when experiments demonstrated that mice which were genetically engineered to lack the PrP^{C} protein were impervious to the effects of PrP^{Sc} [26, 124, 139, 167].

A recent focus of research in biochemistry is to discover the mechanism by which PrP^{C} folds into PrP^{Sc} . Such an understanding would bring great insight into how proteins fold and would have implications toward the treatment of prion diseases [125].

1.4 Geometry, Polymer Physics, or Chemistry?

This scope of this thesis is to provide a geometric foundation for several computational problems mentioned above rather than serve as a treatise in chemistry or polymer physics. In the course of designing algorithms and proving lower bounds, we will often build chains and polygons which may not look like everyday run-of-the-mill molecules. As such, our models will be defined in geometric terms, and our discussions will tend toward geometric concerns rather than chemical ones.

1.5 Overview

Chapter 2 begins our discussion with an explanation of the notation and terminology used throughout the thesis.

The topics are covered in the same order as introduced in this chapter. We deal with problems raised by the geometry community in chapters 3 and 4, and move to those of physics and chemistry in chapters 5 through 7.

In chapter 3 we exhibit two algorithms for reconfiguring polygons in the plane. We demonstrate that all convex configurations of a given polygon are reachable from any other and provide an efficient algorithm for convexifying monotone polygons.

Chapter 4 begins our study of linkages in three dimensions. We prove that the motions used to bring a convex configuration to another can be accomplished through

the use of pivots, and disprove Wegner's conjecture that no polygon admits an infinite sequence of deflations. We further demonstrate that all polygons which admit simple orthogonal projections can be convexified.

Starting with chapter 5, we focus on three-dimensional chains with fixed joint-angles. We present algorithms to decide the feasibility of performing dihedral rotations without self-intersection and demonstrate the impossibility of preprocessing a chain to compute such questions efficiently.

In chapter 6 we study the problem of computing the distribution of possible end-to-end distances of a chain where the joint-angles are fixed. In particular, we prove that finding the minimum and maximum of the distribution is NP-complete when the embedding of the chain is restricted to the plane. We also examine the hardness of these questions when the chain is allowed to assume any three-dimensional configuration.

We continue in the same chapter with the study of bringing such chains into a canonical form. Using earlier results, it is shown that determining whether a chain can be brought to a planar non-crossing position is NP-hard, as is the variant of bringing it into a planar monotone position.

Results from throughout the the thesis are brought together in chapter 7, where we examine the problem of determining the chirality of a polygonal chain with fixed joint-angles and of determining whether an achiral chain is a rubber glove.

Chapter 8 concludes the thesis and discusses future directions.

Chapter 2

Notation and Preliminaries

We briefly summarize the models of computation used herein and provide a basic background in geometry.

2.1 Models of Computation

The model of computation assumed in the analysis and design of algorithms is the *extended real RAM model*. In this model, each memory location can hold a single real number, and the following operations can be performed in unit time.

- Addition, subtraction, multiplication, or division of two real numbers
- Computing the square root of a real number
- Comparison of two real numbers
- Indirect addressing of memory locations

We also present a number of NP-hardness proofs in later chapters, and for these we need to restrict ourselves to a model consistent with the power of a Turing machine. For these proofs we are limited to rational numbers and simple arithmetic, and must consider that storing a large number N or a small number $1/N$ requires $\log N$ space.

Readers who are not familiar with elementary complexity theory will be adequately prepared by either the introductory text of Hopcroft and Ullman [79] or that of Lewis and Papadimitriou [97].

2.2 Trigonometric Computations

In the extended real RAM model, trigonometric functions such as $\sin x$ are not computable. We must resort instead to approximations.

Lemma 2. *The functions $\sin x$ and $\cos x$ can be computed within an error of ε in time $O(x + \log(1/\varepsilon))$ for any $\varepsilon > 0$. Furthermore, the approximation can be guaranteed to be either greater or less than the true value.*

Proof. The functions $\sin x$ and $\cos x$ are expressed by the following MacLaurin expansions.

$$\begin{aligned}\sin x &= x - \frac{x^3}{3!} + \frac{x^5}{5!} - \frac{x^7}{7!} + \frac{x^9}{9!} + \cdots \\ \cos x &= 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \frac{x^6}{6!} + \frac{x^8}{8!} + \cdots\end{aligned}$$

The term $x^{k+2}/(k+2)!$ is at most half of $x^k/k!$ for all $k \geq \sqrt{2}x$. Beyond this term, the remainder of the expansion is less than a geometric series where each term is less than half its predecessor. After computing the first $\sqrt{2}x$ terms, we need compute only $O(\log(1/\varepsilon))$ additional terms to reach one smaller than ε . The remainder of the expansion is less than ε , and the sign of the first term determines whether the remainder is positive or negative. Each term is computable in constant time from its predecessor via a multiplication by $x^2/(i(i-1))$, so the approximation can be computed in time $O(x + \log(1/\varepsilon))$. ■

Since this thesis relies heavily on trigonometric computation, it is necessary to realize that one cannot exactly compute either cosines or arccosines. For clarity

of notation, we use $\widehat{\cos} x$ to denote the approximate cosine of x computed by the technique of the above lemma, and $\cos x$ to represent the true value. (We will not use sines very often in this thesis, but computation of cosines is ubiquitous. Thus we restrict the remainder of this discussion to cosines.)

We may wish to consider the error not just in terms of the discrepancy between $\widehat{\cos} x$ and $\cos x$, but rather in the arccosines of these values. In other words, $\cos x$ is not truly $\widehat{\cos} x$, but there is some value γ such that $\cos(x + \gamma) = \widehat{\cos} x$, that is, $x + \gamma = \arccos(\widehat{\cos} x)$. We wish to bound the error γ .

Lemma 3. *For $0 \leq x \leq \pi$, let $\widehat{\cos} x$ be an approximation which is less than $\cos x$ by at most ε . Then $\arccos(\widehat{\cos} x) - x \leq \sqrt{2\varepsilon}$.*

Proof. By the MacLaurin series for $\cos x$, we obtain the following.

$$\begin{aligned}
 \varepsilon &= \cos x - \widehat{\cos} x \\
 \varepsilon &= \cos x - \cos(x + \gamma) \\
 \varepsilon &= \left(1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \dots\right) - \left(1 - \frac{(x + \gamma)^2}{2!} + \frac{(x + \gamma)^4}{4!} - \dots\right) \\
 \varepsilon &= \frac{(x + \gamma)^2 - x^2}{2!} - \frac{(x + \gamma)^4 - x^4}{4!} - \dots \\
 \varepsilon &\geq \frac{(x + \gamma)^2 - x^2}{2} \\
 \varepsilon &\geq \frac{\gamma^2 + x\gamma}{2} \\
 \varepsilon &\geq \frac{\gamma^2}{2} \\
 \sqrt{2\varepsilon} &\geq \gamma
 \end{aligned}$$

Therefore $\arccos(\widehat{\cos} x)$ is at most $\sqrt{2\varepsilon}$ greater than x . ■

Recall that we can exactly obtain the sines and cosines of an angle of a triangle by computing the cross-product or dot-product of the adjacent sides. Errors enter the

computation only when we wish to convert measures of angles into their sines and cosines or vice versa.

2.3 Elementary Geometry

We denote points by lowercase letters, such as p and q . When appropriate, the *edge* or line segment between two points is written as $\overline{pq}$. This is in contrast to the infinite line $\overleftrightarrow{pq}$ that passes through p and q . The *ray*, or *halfline*, from p through q is the portion of a line with an endpoint at p and extending through q , and is written as $\overrightarrow{pq}$. We denote the distance between these points as $|pq|$.

A *chain* is a path of line segments, or *edges*. We call a chain *simple* if it does not self-intersect (other than at the endpoints shared by adjacent edges). For brevity, we imply simplicity unless otherwise stated. We denote a chain of $n - 1$ edges by the vertices of the chain, $(v_1, v_2, \dots, v_n)$, such that adjacent pairs of vertices share the edges $\overline{v_i v_{i+1}}$, for $1 \leq i \leq n - 1$.

A *polygon* is a cycle of edges. We denote a polygon of n edges by the vertices of the chain, $(v_1, v_2, \dots, v_n)$, such that adjacent pairs of vertices share the edges $\overline{v_i v_{i+1}}$ for $1 \leq i \leq n$. Because a polygon is a cycle, we consider all arithmetic on indices to be modulo n . Thus we also include the edge $\overline{v_n v_1}$ by the expression $\overline{v_i v_{i+1}}$. We also imply simplicity when referring to polygons unless otherwise specified.

In two dimensions a polygon divides the plane into two regions, its interior and its exterior. We call a polygon *convex* if for any two points in its interior, the line segment between them also lies inside. We also refer to vertices of the polygon as *convex*, *straight*, or *reflex*. A vertex of a polygon is *convex* if its internal angle is less than π , *straight* if its angle is equal to π , and *reflex* otherwise. In figure 2.1, the reflex vertices are indicated in red. We obtain an equivalent characterization of a convex polygon by defining it to be a simple polygon which has no reflex vertices. For the most part, we will ignore straight vertices of a polygon. If v_i is a straight vertex, we




Figure 2.1: A polygon and its reflex vertices.

will often consider $\overline{v_{i-1}v_{i+1}}$ as a single edge.

In three dimensions the notions of interior, exterior, and convexity are meaningless, since a polygon does not divide the space. We will use this terminology only if the polygon is planar, and consider in this case that the notions apply with respect to the plane in which the polygon lies.

An *embedding* is a realization of an abstract graph in space which is consistent with given constraints. For the purposes of this thesis, we will consider an embedding of a polygon to be any non-crossing polygon which maintains the edge lengths, their ordering, and when in two dimensions, their orientation (clockwise or counterclockwise). Embeddings of chains are similarly defined, except that chains have no orientations. In chapters 5 and later, we discuss chains and polygons with fixed vertex-angles. In this case we also require an embedding to preserve these angles.

We also refer to embeddings as *configurations*, and the two may be used interchangeably. A polygon can be *reconfigured*, or moved, from one configuration to another in a continuous manner. The continuous manner depends on the motion being discussed; we will study several throughout this discourse. In the analyses of algorithms, we consider a *move* to be a finitely described motion which alters each vertex monotonically, either always increasing or decreasing their angles. Some authors use the term *simple move* or *complex move* to refer to moves which affect $O(1)$ joints or $\omega(1)$ joints respectively. We will not use this terminology here but instead refer directly to the number of joints affected.

A more thorough introduction to terminology can be found in computational geometry texts. Books by O'Rourke [114], by de Berg, van Kreveld, Overmars, and Schwarzkopf [18], and by Preparata and Shamos [120] should provide a more than adequate coverage. The survey paper by O'Rourke [115] provides an introduction to problems of folding and unfolding.

Chapter 3

Reconfiguring Polygons in Two Dimensions

We begin our discussion in section 3.1 with an algorithm to reconfigure convex polygons. In the unfolding of polygons, convex polygons are often used as a canonical target configuration. Since all planar polygons can be convexified [37, 149], it would only remain to show that any convex polygon can be reconfigured into any other convex embedding to complete the proof that any planar polygon can be brought to any other of its embeddings. Lenhart and Whitesides [95] proved this last fact for the case where the edges are allowed to cross. We prove this is also true when self-intersections are not permitted and also detail a much simpler motion to do so.

The algorithms to convexify polygons by Connelly, Demaine, and Rote [37] and by Streinu [149] are highly complex. Much simpler algorithms exist for special cases, such as when the polygon is star-shaped [51]. In section 3.2, we detail a simple algorithm to convexify monotone polygons.

3.1 Reconfiguring Between Convex Configurations


The general idea in the proof of Lenhart and Whitesides [95] is to show that every convex polygon can be reconfigured into another canonical state, a triangle. We describe an algorithm to move a convex polygon into any other convex embedding without the aid of an intermediate canonical form. In particular, each vertex-angle varies monotonically with time, either always increasing or decreasing. In this sense, the motion brings the polygon directly from the source configuration to the target. Our motion is of the simplest type possible; it can be decomposed into a linear number of moves which each change only four angles. In two dimensions, any move must change at least four angles, since three vertices define a triangle, which is rigid.

Consider two convex configurations of a polygon, a source S and a target T . We label each angle of S with a $+$ if it needs to expand to match the corresponding angle in T , with a $-$ if it needs to shrink, and with a 0 if they already match. This labeling scheme is used by Cauchy [29] in his theorem about the rigidity of convex polyhedra. We will also require his key lemma about alternations between labelings of $+$ and $-$.

Lemma 4. (due to Cauchy [29]) *In a $+, -, 0$ labeling of two distinct convex configurations, there are at least four alternations between $+$ and $-$, ignoring 0 's.*

Proof. Because the configurations are distinct, not all labels are 0 . Because a polygon is a cycle, the number of alternations between $+$ and $-$ is even. It cannot be zero, because the sum of the internal angles of a polygon remains constant. It cannot be two as there would be a chain of increasing angles and a chain of decreasing angles which share their endpoints. By lemma 1 the endpoints of the former chain would move further apart while the endpoints of the latter chain move closer together. ■

Our motion will flex quadrilaterals of vertices v_i, v_j, v_k, v_l labeled $+, -, +, -$ as in figure 3.1. We move the pair of vertices marked $-$ apart until one angle matches the

Figure 3.1: A quadrilateral of vertices labeled $+$, $-$, $+$, $-$.

desired value in T . The motion is illustrated in figure 3.2.

Lemma 5. *Let v_1, v_2, v_3, v_4 be a convex quadrilateral. Moving v_1 away from v_3 results in $\angle v_1$ and $\angle v_3$ shrinking and in $\angle v_2$ and $\angle v_4$ expanding, until one of the angles reaches 0 or π .*

Proof. By Euclid's proposition I.25¹, if v_1 and v_3 move apart, the angles $\angle v_2$ and $\angle v_4$ must be increasing. Because no angle moves past 0 or π , we maintain a convex configuration. By lemma 4 there must be four alternations of $+$ and $-$ with respect to all future configurations visited by the motion. Thus v_1 and v_3 must be shrinking. ■

We prove two theorems in this section. The first states that there exists a motion between any two convex configurations, and the second computes it in linear time.

Theorem 1. *Given two convex configurations of a polygon, there is a motion between them that alters each vertex-angle monotonically and which involves at most $n - 3$*

¹"If two triangles have the two sides equal to two sides respectively, but have the base greater than the base, they will also have the one of the angles contained by the equal straight lines greater than the other." [74]




Figure 3.2: The motion of lemma 5.

moves that each change only four angles at once.

Proof. Consider two configurations S and T , and label the vertices of S as described. By lemma 4 there exists a convex quadrilateral of four vertices v_i, v_j, v_k, v_l in cyclic order around the polygon labeled $+, -, +, -$. Keeping all other vertex-angles fixed, we move the vertices marked $-$ apart as per lemma 5 until one of the four angles matches its corresponding angle in T (which must occur before it reaches 0 or π). We repeat this process until all angles match, resulting in a sequence of motions from S to T . Each motion changes the label of an angle from $+$ or $-$ to 0. There cannot be fewer than four signed labels in two distinct configurations by lemma 4, thus $n - 3$ moves suffice. ■

Theorem 2. *Computing the motion in theorem 1 can be done in linear time.*

Proof. We decompose the polygon into continuous blocks of vertices where the labels are either $+$ or 0 and blocks where the labels are $-$ or 0. Let the blocks be as large as possible so that the polygon alternates between *plus-blocks* of all $+$ or 0 vertices and *minus-blocks* of all $-$ or 0 vertices.

We maintain four vertices v_i, v_j, v_k, v_l as the first signed vertices of four adjacent

blocks. Thus they determine a quadrilateral whose vertices are labeled $+, -, +, -$. Each motion flexes the quadrilateral $v_i v_j v_k v_l$ until one of the labels becomes 0. (We can precisely qualify the stopping condition with a small number of trigonometric expressions; we delay this qualification until the proof of lemma 12 where we deal with a more general case.) We then alter the indices i, j, k, l so that v_i, v_j, v_k, v_l are again the first signed vertices of four adjacent blocks.

If v_i becomes 0, we increment its index until v_i is a non-0 vertex. If the new v_i is also labeled $+$, then the above rule again holds. If the new v_i is labeled $-$, then v_i and v_j are in the same block. We advance j, k , and l so that they are the first vertices of adjacent blocks. This is easily accomplished by setting $j \leftarrow k$ and $k \leftarrow l$, and incrementing l until v_l is a $+$ vertex of the next block.

If some other vertex, say v_j , becomes 0, we perform a similar operation. The difference is that if v_j becomes 0 in a block that only had one signed element, the blocks containing v_i, v_j , and v_k are of the same sign and are therefore merged. We set $j \leftarrow l$ and increment k and l to find the beginning of the next two blocks.

There are at most $n - 3$ motions on quadrilaterals, each of which can be performed in constant time. Thus the computation time is dominated by the incrementing of the indices. Assume that initially $i = 1$. Because i advances only when v_i is labeled 0, at any point all vertices $v_1, \dots, v_{i-1}$ must be marked 0. No index can be incremented more than n times, since once it is brought to 1, it would be incremented until it reached i . This would imply there are fewer than four blocks, and thus by lemma 4 the polygon would have reached its target. An index is incremented at most n times, so the algorithm runs in linear time. ■

3.2 Convexifying Monotone Polygons

A polygon is *monotone* with respect to a line ℓ if the intersection of the interior of the polygon with every line perpendicular to ℓ is connected. In other words, each such




Figure 3.3: *Left*: An x -monotone polygon. *Right*: A polygon that is not x -monotone.




Figure 3.4: A monotone polygon decomposed into two chains.

intersection is either empty, a single point, or a line segment. A chain is *monotone* with respect to a line ℓ if the intersection of the chain with every line perpendicular to ℓ is connected.

When ℓ is the x -axis, then the intersection of the chain or polygon with any vertical line is connected, and we call the polygon or chain *x -monotone*. For our purposes, we will consider all monotone objects to be x -monotone, since one can always rotate the object in question until it is monotone with respect to the x -axis. An x -monotone polygon is illustrated in figure 3.3.

We may also consider a monotone polygon as being composed of two monotone chains, an upper chain and a lower chain. In figure 3.4, because the edge $\overline{v_9v_{10}}$ is vertical, one may either label $(v_1, v_2, \dots, v_9)$ as the upper chain and $(v_9, v_{10}, \dots, v_1)$ as the lower, or $(v_1, v_2, \dots, v_{10})$ as the upper chain and $(v_{10}, \dots, v_1)$ as the lower.

The following algorithm convexifies a monotone polygon.

Algorithm Convexify-Monotone

- Repeat until the polygon is convex:
 1. Label the rightmost reflex joint as v_1 .
 2. If v_1 is on the lower chain, then
 - (a) Label the vertices of the polygon counterclockwise from v_1 .
 - (b) Compute the largest index k such that $v_2, \dots, v_{k-1}$ are below the line $\overleftrightarrow{v_n v_1}$ (or to the right of the line, should it be vertical). (See figure 3.5.)
 - (c) Repeat until v_n, v_1, v_{k-1} , or v_k is straightened:
 - i. Fix the positions of joints $v_k, v_{k+1}, \dots, v_n$.
 - ii. Fix all joint angles except those at v_n, v_1, v_{k-1} , and v_k .
 - iii. Rotate v_1 clockwise about v_n , which uniquely defines the motion of the angles at v_n, v_{k-1} , and v_k and the motion of the vertices $v_2, \dots, v_{k-1}$, until either one of the above joints straightens or v_n, v_1, v_{k-1} become collinear, whichever occurs first. (See figure 3.6.)
 - iv. Update the coordinates of v_1 and v_{k-1} .
 - v. If v_n, v_1 , and v_{k-1} have become collinear, then decrement k and update the coordinates of the new v_{k-1} .
 - vi. If a joint has straightened, delete the straightened vertex and fuse the two adjacent edges into one long edge. (From this point on, the two edges will move in concert as a single edge.)
 - (d) Update the coordinates of any vertices that have moved.
 3. If v_1 is on the upper chain, the procedure is similar, except replace “clockwise” with “counterclockwise” and “above” with “below.”



Figure 3.5: The definition of k as per step 2b.

Figure 3.6: The motion of step 2(c)iii.

Let us first prove that each step in the algorithm is well-defined. The concern here is to ensure that the value k always exists.

Lemma 6. *The value k of step 2b is well-defined and is at least 3.*

Proof. Because v_1 is reflex, the ray $\overrightarrow{v_nv_1}$ must enter the interior of the polygon and cross it again at some other edge. This edge cannot be $\overline{v_1v_2}$, so v_2 must be below the line $\overleftrightarrow{v_nv_1}$. Some chain $(v_2, \dots, v_{k-1})$ is below $\overleftrightarrow{v_nv_1}$, so $k \geq 3$. ■

We now prove the correctness of the algorithm. Throughout the remainder of this section, we assume that v_1 is on the lower chain of the polygon. The arguments when v_1 is on the upper chain are symmetric.

Lemma 7. *At no point during the motion does any convex vertex become reflex or any reflex vertex become convex.*

Proof. All vertex-angles are fixed except v_n, v_1, v_{k-1} , and v_k . If any of these straighten, the motion is halted in step 2(c)iii. ■

Lemma 8. *Vertices $v_2, \dots, v_{k-1}$ are convex.*

Proof. Because v_n is either to the left or below v_1 , the ray $\overrightarrow{v_nv_1}$ must cross the polygon directly above or to the right of v_1 . Because the polygon is monotone, the polygon cannot cross the vertical halfline below v_1 . Thus $(v_2, \dots, v_{k-1})$ is either to the right of v_1 or directly below v_1 . If a vertex is directly below v_1 , due to monotonicity and the fact that we ignore straight joints, it must be v_2 . Then v_3 must be to the right of v_2 , and thus v_2 is convex. Since all other vertices lie to the right of v_1 , the rightmost reflex vertex, $v_2, \dots, v_{k-1}$ are convex. ■

The following lemma is illustrated in figure 3.7.




Figure 3.7: Quadrilateral of lemma 9.

Lemma 9. (due to Biedl et al. [20]) *Let $v_1v_2v_3v_4$ (in counterclockwise order) be a simple quadrilateral with v_1 reflex. If the linkage is moved so that v_1 rotates clockwise about v_4 with $\overline{v_3v_4}$ fixed, then the angle at v_1 decreases, and the angles at v_2 , v_3 , and v_4 increase. In other words, all angles approach π .*

Note that $v_nv_1v_{k-1}v_k$ forms a simple quadrilateral with v_1 reflex. We obtain the following result.

Lemma 10. *During the motion of step 2(c)iii, v_{k-1} rotates counterclockwise about v_k , and the angles at v_1 and v_{k-1} approach π .*

We now demonstrate that the line segment $\overline{v_1v_{k-1}}$ rotates counterclockwise. This can also be thought of as examining the motion of v_1 relative to the reference frame of v_{k-1} , determined by fixing the position of v_{k-1} and fixing the coordinate axes. We can visualize this by translating the linkage throughout the motion so that v_{k-1} always maintains its position.

Lemma 11. *Vertex v_1 rotates counterclockwise about v_{k-1} , and v_{k-1} rotates counterclockwise about v_1 .*

Proof. Relative to the reference frame of v_{k-1} , v_k rotates counterclockwise about v_{k-1} . The angle $\angle v_1v_{k-1}v_k$ is increasing by lemma 10, so v_1 is also rotating counterclockwise about v_{k-1} . By symmetry v_{k-1} rotates counterclockwise about v_1 . ■

We are now ready to prove the main theorem of this section.


Theorem 3. *Algorithm Convexify-Monotone correctly convexifies a monotone polygon, and at all times the polygon remains simple and monotone.*

Proof. We first prove that the polygon remains simple throughout the motion. Since the subchain $(v_k, v_{k+1}, \dots, v_n)$ remains fixed at all times, it cannot self-intersect. We first demonstrate that $(v_n, v_1, \dots, v_k)$ does not self-intersect and then prove that these two subchains do not intersect.

To show that $(v_n, v_1, \dots, v_k)$ does not self-intersect, consider the motions performed in step 2(c)iii. We treat $v_n v_1 v_{k-1} v_k$ as a quadrilateral, and move v_1 counterclockwise about v_n . Note that $(v_1, v_2, \dots, v_k)$ forms a convex chain, and that v_{k-1} is opening. Recall Cauchy's lemma (lemma 1) which states that if some joints of a convex chain are opened, the distance between the endpoints increases and the chain remains convex. Therefore the chain $v_1, \dots, v_k$ remains convex and therefore simple. To prove that the edge $\overline{v_n v_1}$ is not involved in an intersection, note that the chain $(v_2, \dots, v_{k-1})$ is rotating counterclockwise about v_1 by lemma 11. If any part of the chain $(v_2, \dots, v_k)$ were to strike $\overline{v_n v_1}$, it would imply that at least one edge of this chain were above $\overrightarrow{v_n v_1}$, a contradiction on the definition of k .

We now prove that $(v_n, v_1, \dots, v_k)$ does not intersect the rest of the polygon. Let W be the region in space bounded by the halfline extending directly downward from v_n , the line segment $\overline{v_n v_k}$, and the ray $\overrightarrow{v_{k+1} v_k}$. The region may either be bounded should $\overrightarrow{v_{k+1} v_k}$ proceed to the left and intersect the halfline below v_n , or it may be unbounded should $\overrightarrow{v_{k+1} v_k}$ proceed to the right. The latter case is illustrated in figure 3.8.

By lemma 11, $v_2, \dots, v_{k-1}$ rotate counterclockwise about v_1 and therefore cannot cross the vertical halfline directly beneath this vertex. If v_1 were to touch the vertical halfline below v_n , v_1 would have become a convex vertex since v_2 lies to its right. This is impossible by lemma 7, so v_1 cannot escape W through this boundary. Neither can $v_2, \dots, v_{k-1}$ since these points always lie to the right of v_1 . No point can cross

Figure 3.8: The region W .

the ray $\overrightarrow{v_{k+1}v_k}$ because by convexity, the first such point would be v_{k-1} , causing v_k to straighten. Lastly, no point can touch the line $\overleftrightarrow{v_n v_k}$, since this would imply the existence of a reflex vertex between v_1 and v_k .


Now that we have proven that the polygon remains simple, its monotonicity easily follows. The fixed subpolygon, $(v_k, v_{k+1}, \dots, v_n)$, has not moved and remains monotone. Vertex v_1 is always to the right of v_n and to the left of v_2 , vertex v_{k-1} is to the right of v_k , and the rest of the polygon is convex.

The only statement left to prove is that the algorithm ends with a convex polygon. After each iteration of the motion of step 2c either k is decremented or a vertex is straightened. Since k can be decremented at most n times, it follows that a vertex is straightened for every n iterations of this step. Once a vertex is straightened, it remains straight, so this can occur at most $n - 3$ times. Thus the polygon consists of only convex and straight vertices after at most $n(n - 3)$ iterations of step 2c. ■

We end this section with an analysis of the time complexity of the algorithm.

Lemma 12. *Each iteration of step 2c can be performed in $O(1)$ time.*

Proof. Steps 2(c)i and 2(c)ii involve no computation, and steps 2(c)iv and 2(c)v

Figure 3.9: The definitions of a, b, c, d , and m .

involve updating $O(1)$ coordinate positions based on known edge lengths and vertex-angles. The most work is performed in step 2(c)iii, where the computation involves calculating the halting event for the motion. The motion is stopped when the first of the following four events occur: v_n straightens, v_{k-1} straightens, v_k straightens, or v_n, v_1, v_{k-1} become collinear. Note that we do not have to consider the case where v_1 straightens, since this is identical to v_n, v_1, v_{k-1} becoming collinear when $k - 1 = 2$.

For brevity of notation, let $a = |v_n v_1|$, $b = |v_1 v_{k-1}|$, $c = |v_{k-1} v_k|$, and $d = |v_k v_n|$, as illustrated in figure 3.9. These quantities will remain fixed throughout the motion. Let $m = |v_k v_1|$, a quantity which increases during the motion. To discover which of the four halting events occur first, we compute the value of m which would occur at each event. Whichever occurs at the smallest value of m halts the motion.


Working through the trigonometry yields the following descriptions of the events.

- Vertex v_n straightens if v_n is convex and $m^2 = a^2 + d^2 + 2ad \cos \angle v_{n-1} v_n v_1$.
- Vertex v_{k-1} straightens if $m^2 = b^2 + c^2 + 2bc \cos \angle v_{k-2} v_{k-1} v_1$.
- Vertex v_k straightens if

$$m^2 = (a \cos \gamma \cos \delta - a \sin \gamma \sin \delta)^2 + (a \cos \gamma \sin \delta + a \sin \gamma \cos \delta)^2,$$

where $\gamma = \angle v_1 v_n v_{k-1}$ and $\delta = \angle v_k v_{k-1} v_n$ at the moment when v_k is straightened.

This configuration is illustrated in figure 3.10, where e denotes $|v_n v_{k-1}|$. These

Figure 3.10: The configuration when v_k straightens.

variables can be computed by the following equations.

$$\begin{aligned}
 e &= \sqrt{(d \cos \angle v_{k+1} v_k v_n + c)^2 + (d \sin \angle v_{k+1} v_k v_n)^2} \\
 \cos \gamma &= (a^2 + e^2 - b^2)/2ae \\
 \cos \delta &= (c^2 + e^2 - d^2)/2ce
 \end{aligned}$$

- Vertices v_n, v_1, v_{k-1} become collinear if $m^2 = a^2 + d^2 - a(a+b) + (a(c^2 - d^2)/(a+b))$.

All of the above expressions can be computed in constant time. ■

Theorem 4. *Algorithm Convexify-Monotone convexifies a monotone polygon using $O(n^2)$ motions, each of which open four joints at once. Furthermore, all $O(n^2)$ motions can be computed in $O(n^2)$ time.*

Proof. By lemma 12 each move can be computed in $O(1)$ time. Each such move either straightens a vertex or decrements k . Since k can be decremented at most n times for each vertex straightened, there are most $O(n^2)$ moves which in total are computed in $O(n^2)$ time. ■

Chapter 4

Reconfiguring Polygons in Three Dimensions

In the preceding chapter, we demonstrated how to reconfigure a planar convex polygon to any other of its planar convex embeddings. We begin chapter 4 by showing that the same result can be achieved by a series of simpler three-dimensional moves known as pivots.

In section 4.2 we discuss a specialized type of pivot called a deflation. In particular, we disprove a conjecture of Wegner by demonstrating a class of polygons that admit infinitely many such motions.

Section 4.3 concludes the chapter with an algorithm to convexify a polygon in three dimensions that admits a simple orthogonal projection.

4.1 Reconfiguring Convex Polygons with Pivots

We show that a convex polygon can be reconfigured to any other convex embedding by the use of pivots. Let v_i and v_j be two vertices of a polygon. In the polymer physics literature, a *pivot* on the diagonal $\overline{v_i v_j}$ is a motion where the portion of the



Figure 4.1: A pivot on $\overline{v_i v_j}$.

Figure 4.2: The transformation illustrated in figure 3.2 but accomplished with three pivots, shown chronologically from left to right. *Top row*: Bird's eye view. *Bottom row*: Oblique view.

polygon between v_i and v_j (denoted in this chapter as $[v_i, v_j]$) is rotated about $\overline{v_i v_j}$. An example of a pivot is illustrated in figure 4.1.

We study the problem of reconfiguring convex polygons discussed in section 3.1 via this restrictive move. A pivot is in some sense the simplest move possible in three dimensions as it affects only two joints at once. Figure 4.2 demonstrates reconfiguring a convex polygon with pivots by simulating the transformation performed in figure 3.2 on page 39.

In 1994, Millett [108] proved that any *equilateral* convex planar polygon can be reconfigured to any other convex embedding via pivots. We demonstrate that this




Figure 4.3: Illustration of the pivots used in theorem 5. *Top row:* Bird's eye view. *Bottom row:* Oblique view.

procedure also works for arbitrary convex polygons, although an unbounded number of pivots (as a function of the number of edges) may be required.

Theorem 5. *There is a sequence of pivots between any two planar convex configurations of a polygon.*

Proof. We use similar logic as in the proof of theorem 1 in that we first locate a quadrangle $v_1v_2v_3v_4$ whose vertices can be labeled $-, +, -, +$, respectively. We simply need to show that the quadrangle motion of lemma 5 can be simulated by pivots. Suitable motions are described as follows and illustrated in figure 4.3.

We first pivot on $\overline{v_1v_3}$, rotating the subchain containing v_2 by $\pi/2$. Our polygon is now in the position of the second illustration of figure 4.3. We now bring the polygon into a “folded convex” position, where it lies in the union of two planes, folded along the crease determined by v_2 and v_4 . We pivot the subchains $[v_4, v_1]$ and $[v_1, v_2]$ into the plane determined by $v_4v_1v_2$, and pivot the subchains $[v_2, v_3]$ and $[v_3, v_4]$ into the plane determined by $v_2v_3v_4$. This brings us to the third quadrangle of the figure. To prove that none of these four pivots cause a collision, first note that $[v_3, v_1]$ and $[v_1, v_3]$ cannot collide because these chains remain on opposite sides of a

vertical plane through v_1 and v_3 . Furthermore, $[v_3, v_4]$ and $[v_4, v_1]$ cannot collide by convexity because the sum of the angles of the two pivots is less than π . The case of $[v_1, v_2]$ and $[v_2, v_3]$ is symmetric.

A final pivot along $\overline{v_2v_4}$ brings the subchain $[v_2, v_4]$ into the same plane as the rest of the polygon. We note that if it is desired to place the polygon in its original plane, then rather than pivoting $[v_4, v_1]$ downward into the plane of $v_4v_1v_2$, we can instead pivot the rest of the polygon into the plane of $[v_4, v_1]$.

We have shown that no collisions occur during these pivots, but it remains to show that any desired quadrangle can be achieved through the repetition of these motions. Consider again the first quadrangle of the top row of figure 4.3. Let x be the closest point from v_2 on the line through v_1 and v_3 . By the law of cosines, the distance between v_2 and v_4 is given by

$$(v_2v_4)^2 = (v_2x)^2 + (v_4x)^2 - 2(v_2x)(v_4x) \cos \angle v_2xv_4.$$

After the first pivot (second quadrangle of the figure), $\angle v_2xv_4$ is $\pi/2$, so the last term is equal to zero. Therefore, after each series of pivots, v_2 and v_4 come closer, and their squared distance decreases by the original value of $|2(v_2x)(v_4x) \cos \angle v_2xv_4|$. Thus we always make considerable progress toward our target configuration and will eventually reach it, unless in our target either v_2x , v_4x , or $\cos \angle v_2xv_4$ is zero. In each of these cases, we will show that either v_2 or v_4 is collinear with v_1 and v_3 . The target configuration cannot have both v_2x and v_4x equal to zero, because then the polygon would self-intersect. If only v_2x (or v_4x) is zero, then v_1, v_2 (or v_4), and v_3 are collinear. If $\cos \angle v_2xv_4$ is zero, then v_1, v_4 , and v_3 are collinear because $\overline{v_2x}$ is perpendicular to the line through v_1 and v_3 .

Assume without loss of generality that if one of v_2x , v_4x , or $\cos \angle v_2xv_4$ is zero, then v_1, v_2 , and v_3 are collinear. In this case, v_2 is the only vertex between v_1 and v_3 by convexity of the target configuration. Let the distance from v_4 to $\overrightarrow{v_1v_3}$ be greater than ε in the target. When v_1, v_2 , and v_3 are close enough to collinear (when $|v_2x| < \varepsilon$)




Figure 4.4: Two pivots performed to reconfigure a parallelogram that is not a rhombus.

a pivot about v_1v_3 of any angle, even π , will not cause any self-intersections. If we then pivot until v_2v_4 is the same distance as v_4x , and perform the remaining pivots to restore planarity of the polygon, $v_1v_2v_3$ can be made collinear or as close to collinear as desired. ■

The geometric progression of the above proof hints at the notion that there may be some polygons for which the number of pivots required to move between any two arbitrary configurations may not be bounded by a function of the number of edges. Indeed we show this to be the case. Before a direct proof of this statement, we require the following lemma. Figure 4.4 serves as a useful visual aid during the course of the proof of lemma 13.

Lemma 13. *Let $v_1v_2v_3v_4$ be a planar convex quadrangle. After two pivots, suppose the quadrangle is once again planar and convex, resulting in a configuration $v''_1v''_2v''_3v''_4$. Then $\angle v''_2v''_1v''_4$ will be at least the value of the expression $|\angle v_2v_1v_3 - \angle v_4v_1v_3|$.*

Proof. If both pivots are on the diagonal $\overline{v_2v_4}$, then the angle at v_1 has not changed. We break the remaining possibilities into two cases: the case in which the pivot on $\overline{v_1v_3}$ is the first pivot (or both), and the case in which it is preceded by a pivot on

$\overline{v_2v_4}$.

If the first is on $\overline{v_1v_3}$, then the pivot occurs on a planar polygon. (If both pivots are on $\overline{v_1v_3}$, we can merge them into a single pivot and thus the argument is identical.) Ignoring intersections for the time being, let v_2 rotate freely around the diagonal $\overline{v_1v_3}$. The point v_2 traces out a circle in space centered on v_1v_3 ; thus $\angle v_2v_1v_3$ is constant. Because $\angle v_4v_1v_3$ does not vary during the pivot, the resulting $\angle v'_2v'_1v'_4$ and thus $\angle v''_2v''_1v''_4$ are at least the difference of these two angles.

If the pivot on $\overline{v_2v_4}$ occurs first, then the next pivot must occur on $\overline{v_1v_3}$ to bring the quadrangle into a planar position. We can also visualize this as the triangle $\triangle v_1v_3v_4$ rotating about $\overline{v_1v_3}$ until it is coplanar with the triangle $\triangle v_1v_3v_2$. In this case, the distance v_2v_4 , which was constant during the previous pivot, is now increasing. By Euclid's proposition I.25, $\angle v_2v_1v_4$ must have increased. ■

The next theorem follows easily from lemma 13.

Theorem 6. *There exist polygons which require arbitrarily many pivots to achieve a goal configuration.*

Proof. Examine the leftmost parallelogram in figure 4.4, where v_1 is at an acute angle. Because this polygon is a parallelogram, its planar configurations are all convex. These can be made as flat as desired, i.e., in which $\angle v_2v_1v_4$ is arbitrarily close to zero. Furthermore, because the polygon is not a rhombus, $\angle v_2v_1v_3 \neq \angle v_4v_1v_3$. By the law of sines,

$$\sin \angle v_2v_1v_3 = \frac{|v_2v_3|}{|v_3v_4|} \sin \angle v_4v_1v_3.$$

Since $\sin \theta$ is a concave monotonic function passing through the origin for $0 \leq \theta \leq \pi/2$, it follows that $k \sin \theta < \sin k\theta$ for all $k < 1$. Assume that $|v_2v_3| < |v_3v_4|$.

$$\begin{aligned}
\sin \angle v_2 v_1 v_3 &= \frac{|v_2 v_3|}{|v_3 v_4|} \sin \angle v_4 v_1 v_3 \\
\sin \angle v_2 v_1 v_3 &< \sin \left[\frac{|v_2 v_3|}{|v_3 v_4|} \angle v_4 v_1 v_3 \right] \\
\angle v_2 v_1 v_3 &< \frac{|v_2 v_3|}{|v_3 v_4|} \angle v_4 v_1 v_3
\end{aligned}$$

Because $\angle v_2 v_1 v_4 = \angle v_2 v_1 v_3 + \angle v_4 v_1 v_3$, as $\angle v_2 v_1 v_4$ approaches zero every two pivots can only reduce $\angle v_2 v_1 v_4$ to

$$\angle v_2'' v_1'' v_4'' \geq \frac{|v_3 v_4| - |v_2 v_3|}{|v_3 v_4| + |v_2 v_3|} \angle v_2 v_1 v_4.$$

Thus $\angle v_2 v_1 v_4$ approaches zero by a constant fraction each time. We can choose a goal configuration with a small enough $\angle v_2 v_1 v_4$ to require any number of pivots desired. (We note that although one cannot achieve a configuration in which $\angle v_2 v_1 v_4 = 0$, this is not a valid configuration as the polygon would be flat and therefore self-intersecting.) While this proves the theorem for the case in which every two pivots restores the polygon to a planar configuration, we have not directly proved the theorem for arbitrary pivots. This is easily remedied by considering each pivot as a pair of pivots on the same diagonal, the first to bring the quadrangle into a planar non-intersecting position and the second to produce the original pivot as desired. ■

4.2 Deflations

In 1993, Bernd Wegner [164] posed a problem concerning deflations, the inverse of the flip operation of Erdős. Recall that a deflation selects a line that passes through the polygon in exactly two vertices and reflects one of the two subchains across this line. Similar to the concept that convex polygons do not admit flips, Wegner defined a *deflated polygon* as a polygon that does not admit any deflations. He conjectured that




Figure 4.5: A deflation.




Figure 4.6: A deflated polygon.


any polygon would be deflated after finitely many deflations. Figure 4.5 illustrates the deflation operation, and figure 4.6 illustrates a polygon which is deflated. We exhibit a counterexample to Wegner's conjecture and discuss its ramifications to a convexifying operation known as a mouth-flip.

4.2.1 A counterexample to Wegner's conjecture

Consider a quadrilateral $abcd$ such that $|ab| + |cd| = |ad| + |bc|$ and that no two adjacent edges have the same length. We prove that such a quadrilateral can be deflated infinitely many times.

Lemma 14. *There is no self-intersecting embedding of quadrilateral $abcd$ except when a, b, c , and d are collinear.*

Proof. Consider any self-intersecting embedding of $abcd$. We assume without loss of generality that the edges $\overline{ad}$ and $\overline{bc}$ cross. Let x be the point of intersection as illus-

Figure 4.7: A self-intersecting embedding of $abcd$.

trated in figure 4.7. No two adjacent edges have the same length, so x is distinct from a, b, c, d . Because the vertices are not all collinear, at least one of $\angle axb$ or $\angle cxd$ is not a straight angle. Therefore $(|ax| + |xb|) + (|cx| + |xd|) > |ab| + |cd|$. By regrouping the terms of this expression, we obtain $(|ax| + |xd|) + (|cx| + |xb|) > |ab| + |cd|$. Since a, x , and d are collinear, as are c, x , and b , this implies that $|ad| + |bc| > |ab| + |cd|$, a contradiction. ■

This lemma is sufficient to prove that the quadrilateral $abcd$ admits infinitely many deflations.

Theorem 7. *A quadrilateral where the sums of the lengths of the opposite edges are equal and where no two adjacent edges have the same length admits infinitely many deflations.*

Proof. By lemma 14 a deflation that causes a self-intersection would have to bring all vertices collinear. Since only one vertex moves for each deflation, the embedding prior to this deflation would need exactly three vertices to be collinear. This would be a self-intersecting embedding, which by lemma 14 does not exist. ■

In fact, all polygons that admit an infinite sequence of deflations have a similar subpolygon. We prove that in any polygon that admits infinitely many deflations, at some point in the reconfiguration, two vertices v_1 and v_k (k even) are brought into a




Figure 4.8: The proof of lemma 15.

position such that $|v_1v_2| + |v_3v_4| + \cdots + |v_{k-1}v_k| = |v_2v_3| + |v_4v_5| + \cdots + |v_kv_1|$.

The following terminology will facilitate our discussion. In a deflation in which a subchain $(v_1, v_2, \dots, v_k)$ is reflected across the line $\overleftrightarrow{v_1v_k}$, we refer to $\overleftrightarrow{v_1v_k}$ as the *line of deflation* and to v_1 and v_k as the *anchors* of the deflation. We also say that the vertices $v_2, \dots, v_{k-1}$ are *moved* by the deflation. Note that the anchors of a deflation are not moved by it.

We first demonstrate that when a vertex is used as the anchor of a deflation, its angle tends toward 0 or 2π .

Lemma 15. *Let $\angle v_i$ denote the internal angle at some vertex v_i of the polygon. Then any deflation of the polygon of which v_i is an anchor causes the value $|\pi - \angle v_i|$ to either increase or remain constant.*

Proof. Let the deflation occur on the line $\overleftrightarrow{v_i v_j}$, flipping the subchain $(v_i, v_{i+1}, \dots, v_j)$ as illustrated in figure 4.8. The only angles $\angle v$ for which the value $|\pi - \angle v|$ can be changed by this operation are v_i and v_j ; all others are either maintained or reflected. If v_{i-1} lies on $\overleftrightarrow{v_i v_j}$, then $|\pi - \angle v_i|$ does not change. If v_{i-1} and v_{i+1} are initially on opposite sides of the line, then after the reflection they are on the same side. Thus the distance $|v_{i-1}v_{i+1}|$ must decrease, and by Euclid's Proposition I.25 the value $|\pi - \angle v_i|$ also decreases. ■

Lemma 16. *Each deflation causes the area of the convex hull of the polygon to de-*

crease.

Proof. The line of deflation $\overleftrightarrow{v_1 v_k}$ separates the polygon into two subpolygons, $(v_1, v_2, v_3, \dots, v_k, v_1)$ on one side of the line and $(v_k, v_{k+1}, \dots, v_1, v_k)$ on the other. After a deflation, the former is placed inside the latter. The area of the convex hull of the entire polygon is decreased by at least the area of the convex hull of $(v_1, v_2, v_3, \dots, v_n, v_1)$. ■

Lemma 17. *Let v_1 and v_k be the anchors for infinitely many deflations. Then the area of the convex hull of the subpolygon $(v_1, v_2, v_3, \dots, v_k)$ moved by the deflations must tend to zero as the number of deflations grows to infinity.*

Proof. The lemma follows immediately from lemma 16. ■

Lemma 18. *Let v_i be a vertex that is moved by deflation infinitely many times. Then v_i is the anchor for infinitely many deflations, and the internal angle at v_i approaches 0 and 2π ; that is, the value $|\pi - \angle v_i|$ approaches π .*

Proof. Each time the vertex v_i is moved by deflation, the triangle $\triangle v_{i-1} v_i v_{i+1}$ must be contained in the convex hull of the flipped subpolygon. This triangle has area $\frac{1}{2} |v_{i-1} v_i| |v_i v_{i+1}| \sin \angle v_i$. Since the distances are fixed, the sine of the angle must be decreasing. A vertex cannot approach π by lemma 15, so $|\pi - \angle v_i|$ must be approaching π . A vertex-angle can change only if it is an anchor of a deflation. ■

Lemma 19. *For any infinite sequence of deflations on a polygon, there are at least two vertices that are moved by deflation only finitely many times.*

Proof. To achieve a contradiction, suppose there are fewer than two vertices that are moved by deflation infinitely many times. Then all internal angles of the polygon

other than at these vertices approach 0 or 2π . The sum of the internal angles of a polygon of n vertices is $(n - 2)\pi$, which is 0 modulo 2π . Therefore, if there is only one vertex that moves finitely many times, its angle must also approach 0 or 2π . All edges of the flipped subpolygon become more and more parallel as the deflations continue. At some point in the infinite sequence of deflations all edges become nearly horizontal, and the convex hull becomes nearly a line segment. Assume without loss of generality that this line segment is horizontal.

A line of deflation must separate the two edges incident to each of its anchors; otherwise it would not split the polygon into two subpolygons. Since all internal angles are close to 0 or 2π , only horizontal lines can separate two consecutive edges. Therefore all lines of deflation must be horizontal.

Consider the leftmost and rightmost vertices of the polygon. These vertices cannot be moved by a horizontal line of deflation because they would move vertically, and there is no subpolygon extending left or right enough to contain them. Therefore the leftmost and rightmost vertices cannot be moved. ■

Lemma 20. *Let $v_2, v_3, \dots, v_{k-1}$ be vertices that are moved infinitely many times such that v_1 and v_k are moved only finitely many times. Then k is an even number. Furthermore, after the last deflation which moves v_1 or v_k , $|v_1v_2| + |v_3v_4| + \dots + |v_{k-1}v_k| = |v_2v_3| + |v_4v_5| + \dots + |v_kv_1|$.*

Proof. Since the internal angles at $v_2, v_3, \dots, v_{k-1}$ approach 0 or 2π as the number of deflations grows, the edges will alternate between those that are directed toward v_k and those that are directed toward v_1 . The sum of the lengths of all edges directed toward v_k minus those directed toward v_1 must equal the net distance traversed, $|v_1v_k|$. ■



Figure 4.9: *Left*: A mouth. *Right*: A mouth flip.

Figure 4.10: A star-shaped polygon that admits no mouth-flips.

4.2.2 Mouth flips

The existence of polygons that admit infinitely many deflations has implications for a convexifying motion known as a *mouth-flip*. A mouth of a polygon is three consecutive vertices such that the triangle formed by them is empty and exterior to the polygon, such as in figure 4.9. In a mouth flip, the two edges of the mouth are reflected across the line through its first and third vertices.

Millett [108] demonstrated that any equilateral star-shaped polygon can be convexified with finitely many mouth-flips, and Toussaint [157] later pointed out that $n!$ suffice. It is also known [43] that there exist non-equilateral star-shaped polygons that do not admit a single mouth-flip, such as in figure 4.10.

We can also construct star-shaped polygons that cannot be convexified by mouth-flips, even though they admit infinitely many. Consider the polygon $abcde$ of figure 4.11 where $|ab| + |cd| = |ad| + |bc|$, with b and c in the pocket of the convex hull determined by $\overline{ad}$. Then $abcd$ forms a quadrilateral that cannot be deflated. The polygon can be convexified only if b and c can cross $\overline{ad}$, which is impossible by theorem 7.




Figure 4.11: A polygon that admits infinitely many mouth-flips.

4.3 Convexifying Polygons with Simple Projections

We leave our discussion of pivots and planar polygons and turn to more general motions and polygons. In chapter 1, we demonstrated locked chains and polygons, illustrated in figures 1.12, 1.13, 1.14, and 1.15. We now prove that if a polygon admits a simple orthogonal projection to some plane, it can be convexified. An *orthogonal projection* of a polygon onto some plane $\mathcal{P}$ is the set of all points $q \in \mathcal{P}$ for which some point of the polygon lies on the line through q orthogonal to $\mathcal{P}$. Intuitively, one may think of the projection as the shadow the polygon would cast on the ground when a light source is infinitely high, as in figure 4.12. We call a projection *simple* if no two points of the polygon project to the same point on the plane.

Theorem 8. *A polygon which has a simple orthogonal projection can be reconfigured to a polygon with a convex orthogonal projection.*

Proof. Assume the polygon has a simple projection onto the xy -plane. By keeping the z -coordinate of each vertex fixed, we can reconfigure the projection on the xy -plane by making the polygon in three-space track the motions as necessary. By Connelly, Demaine, and Rote [37] and Streinu [149], any planar polygon can be convexified. If the projection does not self-intersect during the motion, surely the polygon does not self-intersect. Thus we can reconfigure any polygon with a simple orthogonal projection into a polygon with a convex orthogonal projection. ■




Figure 4.12: A polygon with a simple orthogonal projection to the xy -plane.

Our algorithm to convexify a polygon with a simple projection first reconfigures it into a polygon with a convex projection. (For brevity, we will often omit the word, “orthogonal,” and write only, “simple projection.”) The algorithm then repeatedly reconfigures the polygon while maintaining the convexity of its projection. The overall motion is a simple combination of basic primitives which are easy to compute and visualize. We now explain our first primitive motion, which straightens a monotonic polygonal chain lying on a single vertical plane.

4.3.1 Springs

Consider a polygonal chain which is entirely contained in a vertical plane, such as the one illustrated in figure 4.13. If we wish to straighten this series of edges, we need only to pull the rightmost vertex to the right, moving only the last two vertices and changing the angle at the third, as in figure 4.14. When the second vertex straightens, we freeze that joint permanently as straight, thereby effectively deleting




Figure 4.13: Coplanar edges whose projection is a single edge.




Figure 4.14: A spring-unfolding of a planar chain.

a vertex, resulting in a polygonal chain of one fewer vertex. At each step the two right edges move until straightened, at which point they are frozen to act henceforth as a single edge. This causes the three marked vertices to flex; the black vertices move and are opened while the white vertex opens but stays in position. The procedure is then repeated with the two rightmost edges (including the newly fused edge) until either the entire chain is straightened or the motion is stopped. It is vital to realize that the projection of the polygonal chain is, at all times, in the same vertical plane. Thus the projection is always a line segment, growing throughout the motion.

As this motion resembles the straightening of a spring (as in figure 4.15), we refer to this straightening as a *spring-unfolding* of a polygonal chain. In keeping with this terminology, we define a *spring* as a consecutive sequence of at least two non-collinear




Figure 4.15: Pulling a spring until it straightens.

edges whose simple orthogonal projection is a line segment. Thus the polygonal chain of figure 4.13 is a spring.

It is a simple matter to compute when each vertex straightens given a RAM computer which can compute square roots (and therefore distances). Because only two edges move at once, each motion, which straightens a vertex, can be computed in constant time. Thus a spring of n vertices can be straightened in $O(n)$ time.

4.3.2 The Spring Algorithm

Our algorithm requires that the convex projection be a triangle. Changing a polygon from one convex position to another is not difficult and can be accomplished in linear time as per our algorithm of section 3.1. By keeping the z -coordinate of each vertex fixed, we can reconfigure the projection on the xy -plane while making the polygon in three-space track this motion.

Once the projection is a triangle, all edges of the polygon lie along three vertical planes. Therefore the polygon is composed of three or fewer springs. We adopt the term k -spring as a triangle which contains exactly k springs. For example, a polygon with a triangular projection $\triangle abc$ that has only one edge projecting onto ab , but several edges projecting onto each ac and bc , is a 2-spring.

The Spring Algorithm works just as one would intuitively straighten a triangle made of springs. From a high-level standpoint, we repeatedly pull on vertices, straightening a spring each time. For low-level details, it will become apparent that we can easily compute which vertices of a spring move and at which time since their

motions are uniquely linked to the high-level description of the springs.

The following is our algorithm in a high-level form.

The Spring Algorithm

1. While the polygon has more than four edges,
 - (a) If the polygon is a 3-spring, reconfigure it into a 2-spring (as per figure 4.16).
 - Fix the position of the spring ab . Keeping the height of c fixed, pull c orthogonally away from ab until one of the springs (ac or bc) straightens. At any moment, each of the three springs remain in a single vertical plane, and thus the projection is always a triangle during this step.
 - (b) Else if the polygon is a 2-spring, reconfigure it into a 1-spring (as per figure 4.17).
 - Let $\overline{ab}$ be the non-spring edge; this edge remains fixed. Keeping the height of c fixed, pull c orthogonally away from $\overline{ab}$ until one of the springs (ac or bc) straightens. As in the previous step, each of the three springs remains in a single vertical plane, so that the projection is always a triangle throughout this step.
 - (c) Else if the polygon is a 1-spring, reconfigure it into a 2- or 3-spring or a quadrilateral (as per figure 4.18).
 - Let $\overline{ab}, \overline{bc}$ be the non-spring edges, labeled such that the angle at c is acute. Fix in space the median¹ vertex d of the spring ac . Keeping c at

¹Suppose the spring ac contains k vertices. By the phrase, “median vertex d of the edge ac ,” we mean a vertex d such that each ad and dc contain roughly $k/2$ vertices. To be exact, the two chains ad and dc would contain $\lceil k/2 \rceil$ and $\lceil (k+1)/2 \rceil$ vertices, respectively.




Figure 4.16: Step 1a. Reconfiguring a 3-spring into a 2-spring.

a fixed height, rotate $\overline{bc}$ to open the internal angle at b while unfolding the spring cd until one of the two following possibilities occurs.

- i. If the edges $\overline{ab}$ and $\overline{bc}$ achieve collinear projections, we attain a 3-spring, unless $\overline{ab}$ and $\overline{bc}$ happen to be collinear in three-space, in which case we attain a 2-spring.
 - ii. If the spring cd straightens, rotate $\overline{cd}$ to open the internal angle at c while keeping d at a fixed height and unfolding the spring ad until one of the three following conditions is satisfied. The edges $\overline{bc}$ and $\overline{cd}$ could achieve collinear projections, resulting in a 2-spring (or 1-spring should b, c , and d be collinear). The edge $\overline{ab}$ and the spring ad could achieve collinear projections, resulting in a 1-spring. Lastly, the spring ad could straighten, resulting in a three-dimensional quadrilateral.
2. If the polygon has four edges, label it $abcd$. Rotate c about the diagonal $\overline{bd}$ until the quadrilateral is planar. Convexify the planar quadrilateral.
 3. Return the convexified polygon.

We first prove the correctness of the algorithm. Since each step involves the straightening of vertices until a triangle or quadrilateral is obtained, certainly the algorithm will finish. We now prove the hypotheses on which the algorithm is based, that the polygon never self-intersects and that the projection is a triangle or convex




Figure 4.17: Step 1b. Reconfiguring a 2-spring into a 1-spring.




Figure 4.18: Step 1c. Reconfiguring a 1-spring. *Top*: Step 1(c)i. *Bottom*: Step 1(c)ii.

quadrilateral at all times. No spring can intersect itself, so intersections could only be between two springs, resulting in a non-convex projection. Note that during the execution of steps 1a and 1b the projection is a triangle, so there is no possibility for a non-convex projection. During step 1c, the vertex c is rotated about b , causing the projection to be a quadrilateral. Consider the motion of c during this step. Because $\angle acb$ is acute, the quadrilateral will stay convex. (If c were obtuse, then the angle at d would initially become reflex.) From this point on, if any vertex of the projection straightens, it becomes a triangle. Therefore the projection is convex at all times.

Each high-level motion (steps 1a, 1b, and 1c) is easily computable. We sum the lengths of all edges in the springs to discover their length when straightened to see which spring straightens first. For the low-level description of the motions of individual vertices, we can determine the position of each during the spring-unfoldings since we know how the length of each spring increases over time. As noted in section 4.3.1, a spring-unfolding of a chain can be computed in time linear in the number of vertices in the spring, and the spring-unfolding is determined solely by the motion of the endpoints of the spring. Therefore each high-level motion can be computed in time linear in the number of vertices of the polygon.

During every two or three iterations of the while loop (step 1), the polygon will be in a 1-spring configuration. If the polygon has n (unstraightened) vertices at this point in the algorithm, $n - 1$ of these will be contained in the spring. During the execution of step 1c, the spring is separated into two nearly equal parts, each of which contains no less than $\frac{n-4}{2}$ vertices. After two or three more iterations of step 1, the polygon is again a 1-spring, so at least one of these two springs of $\frac{n-4}{2}$ vertices must have been straightened. After at most every three iterations, the polygon's complexity, and thus the time complexity of each iteration, is approximately halved. By this geometric progression, the entire algorithm finishes in linear time.

We further note that at no point in time are more than seven vertices opening or

closing. During step 1 only two adjacent edges of the projection are in motion at any one time. The three endpoints adjacent to these two edges are in motion, and for each of two edges in the projection corresponding to springs, two additional vertices on each are involved in the spring-unfoldings. Therefore the total number of vertices changing at any one time is seven.

Theorem 9. *A polygon which admits a convex orthogonal projection can be convexified in $O(n)$ time with $O(n)$ motions that each alter $O(1)$ vertices.*

We can apply theorem 9 to a variety of previous results on convexifying polygons in the plane, by first convexifying the projection and then applying the Spring Algorithm. Building on Streinu's algorithm [149] to convexify a planar polygon, we obtain theorem 10.

Theorem 10. *A polygon which admits a simple orthogonal projection can be convexified with $O(n^2)$ motions.*

There are currently two classes of planar polygons which have special convexifying algorithms. We demonstrated how to convexify monotone polygons in section 3.2, and Everett, Lazard, Robbins, Schröder, and Whitesides [51] have designed an algorithm to convexify star-shaped polygons. Applying theorem 9 to these results yields the following two theorems.

Theorem 11. *A polygon which admits a monotonic orthogonal projection can be convexified in $O(n^2)$ time with $O(n^2)$ motions that each alter $O(1)$ vertices.*

Theorem 12. *A polygon which admits a star-shaped orthogonal projection can be convexified in $O(n^2)$ time with $O(n)$ motions that each alter $O(n)$ vertices.*

In light of the above, it is important to be able to determine if a polygon has an orthogonal projection or one of the mentioned special cases. Bose, Gómez, Ramos, and Toussaint [24] have an algorithm to determine whether a polygon has a simple

orthogonal projection in $O(n^4)$ time and another which determines whether it has a monotonic projection in $O(n^2)$ time. The former algorithm was also independently discovered by Barequet, Dickerson, and Eppstein [14]. No algorithm to determine whether a polygon admits a star-shaped projection has yet been published.

Chapter 5

Dihedral Rotations

For the remainder of the thesis, we deal with chain linkages in three dimensions. Recall the molecular model of figure 5.1 used by the chemistry and physics communities. Proteins and other polymers consist of long backbones with short side chains, and their structures are often approximated by chain linkages. As in the model we have used in previous chapters, the chain can be reconfigured so long as it does not self-intersect and the edge lengths are preserved, but we impose the additional restriction that the vertex-angles are maintained throughout the motion. Thus the range of motion permitted at a vertex of the chain is as illustrated in figure 5.2.




Figure 5.1: Ball-and-stick model of a protein. Proteins are composed of a long chain called a backbone, drawn with white bonds, and various side chains, shaded.




Figure 5.2: Range of motion allowed with a fixed joint-angle.

The simplest reconfiguration of such a chain is a *dihedral rotation*, illustrated in figure 5.3 and in stereo in figure 5.4. Select some interior edge $\overline{uv}$ of the chain. This defines two subchains A and B . We keep them individually rigid and, leaving A fixed in space, rotate B around uv (while maintaining the fixed vertex-angle at v) by some angle ϕ . Due to the fixed vertex-angle at v , B spins about an axis of rotation collinear with the edge $\overline{uv}$. The effect of such a motion is to change the *dihedral angle* at $\overline{uv}$. Given a four-vertex subchain (t, u, v, w) , the dihedral angle at $\overline{uv}$ is the angle between the planes determined by $\triangle tuv$ and $\triangle uvw$. Thus we call the motion a dihedral rotation of angle ϕ at $\overline{uv}$. Because the chain must remain simple during the motion, we call the rotation *feasible* if it can be performed without self-intersection.

We consider a configuration of such a chain to be any embedding which does not self-intersect and preserves the lengths of the edges and the angles at the vertices. We begin our discussion with determining the feasibility of a dihedral rotation and continue with the matter of preprocessing a chain to perform such computations more quickly.




Figure 5.3: A dihedral rotation of angle ϕ at $\overline{uv}$. (Shown in stereo in figure 5.4.)




Figure 5.4: *Stereogram*: A dihedral rotation at $\overline{uv}$. (The stationary part of the chain has been chosen to be planar in order to emphasize the three-dimensional motion; this is not a requirement of the dihedral rotation.)

5.1 Determining the Feasibility of Dihedral Rotations

We formalize the problem as follows.

Dihedral Rotation: *Given a three-dimensional chain, a selected edge $\overline{uv}$, and an angle ϕ , can a dihedral rotation of angle ϕ be performed at $\overline{uv}$ without causing the chain to self-intersect?*

We sketch a quadratic-time algorithm to answer the above question and provide a lower bound on its complexity.

Theorem 13. *The Dihedral Rotation problem is solvable in $O(n^2)$ time and $O(n)$ space, where n is the number of edges in the chain.*

Proof. The main idea is to spin one of the subchains around $\overline{uv}$ completely (by angle 2π), and examine all self-intersections which occur along the way. We compute the angle of rotation at which intersection occurs during the motion, and if none occur before angle ϕ is reached, the dihedral rotation can be performed.

Consider any plane P incident to $\overline{uv}$. Without considering intersections between the two subchains, rotate subchain A around $\overline{uv}$ and trace where the subchain would sweep through P during the motion. Perform the same for subchain B . Each line segment rotating about uv sweeps a pair of hyperbolic arcs¹ in P , so we have two planar arrangements of $O(n)$ hyperbolic arcs each.

Rotating subchain A is equivalent to rotating subchain B , modulo a rotation of the entire chain, because both are about the same axis. If executing a rotation of 2π

¹ Assume $\overline{uv}$ is on the x -axis and that P is the xy -plane. Rotating a point (x, y, z) about $\overline{uv}$ sweeps through P at $(x, \sqrt{y^2 + z^2}, 0)$ and $(x, -\sqrt{y^2 + z^2}, 0)$. A line parametrized by $(x, y, z) + t(a, b, c)$ sweeps out two curves $(x + at, \pm\sqrt{(bt + y)^2 + (ct + z)^2}, 0)$, which asymptotically approach the lines $(x + at, \pm t\sqrt{b^2 + c^2}, 0)$.

causes an intersection, there will be an intersection of the two arrangements. We can compute if any arcs corresponding to A intersect an arc corresponding with B with brute force in $O(n^2)$ time. Faster yet, we can use the algorithm of Balaban [11]. This yields a time complexity of $O(n \log n + k)$, where k is the sum of the self-intersections in each arrangement and the intersections between arrangements. (Although k could be as large as $\Theta(n^2)$ in the worst case, it would prove faster in some instances than brute force.) Regardless of the method used to detect intersections, as each intersecting pair of arcs is detected, we determine the angle of rotation required before such an intersection would occur during the dihedral rotation. Thus the entire algorithm finishes in $O(n^2)$ time. Furthermore, once we detect a pair of intersecting arcs and determine the angle of rotation required, we will never again need to examine that intersection. Therefore we do not need to store intersections already detected, so we require only linear space by using brute force or Balaban's algorithm. ■

Rather than asking the Dihedral Rotation problem for an arbitrary ϕ , one may wish to know if a complete rotation (of angle 2π) can be performed. This special case can be handled much more efficiently.

Theorem 14. *For $\phi = 2\pi$, the Dihedral Rotation problem is solvable in deterministic time $O(n2^{\alpha(n)} \log^2 n)$ and space $O(n2^{\alpha(n)})$ and in expected time $O(n2^{\alpha(n)} \log n)$ and space $O(n2^{\alpha(n)})$, where $\alpha(n)$ is the slow-growing, nearly constant, inverse Ackermann function.*

Proof. We repeat the technique used in the proof of theorem 13. However, for this choice of ϕ , we do not have to check to see which intersection occurs first in our arrangements. If any intersection at all exists between the arrangements, the dihedral rotation cannot be performed. Given two arrangements of hyperbolic arcs which pairwise intersect at most twice, we can detect an intersection between them in the stated deterministic time and space using algorithms of Agarwal and Sharir [1]

and of Guibas, Sharir, and Sifrony [71]. By replacing the latter with the algorithm of Chazelle, Edelsbrunner, Guibas, Sharir, and Snoeyink [32], we can achieve the stated expected time and space. ■

The Dihedral Rotation problem can thus be solved in quadratic time, and in expected time $O(n2^{\alpha(n)} \log n)$ for the special case where $\phi = 2\pi$. We now prove that any algorithm which solves this problem requires $\Omega(n \log n)$ time in the worst case.

Theorem 15. *The time complexity of the Dihedral Rotation problem on a chain of n edges is $\Omega(n \log n)$ under the algebraic decision tree model of computation.*

Proof. We perform a reduction to Dihedral Rotation from Element Uniqueness.

Element Uniqueness: *Given a set $S = \{s_1, \dots, s_n\}$, is every element unique? In other words, does $i \neq j$ imply that $s_i \neq s_j$?*

In 1983, Ben-Or [16] proved that the time complexity of the Element Uniqueness problem is known to have an $\Omega(n \log n)$ lower bound under the algebraic decision tree model. Given a set S , we create a chain and select an edge in linear time such that answering the Dihedral Rotation problem also answers the Element Uniqueness problem on S .

Let N be the number of elements in the set S . We consider all elements to be strictly positive. If this is not the case, we can add some suitable integer to each element in the set in linear time.

As the chain we will build is difficult to visualize, it will be far easier first to explain how a tree can be constructed in the xz -plane to answer our query. We will not address the issue of constructing the tree but will later show how to construct a chain in linear time which behaves in an identical manner.




Figure 5.5: *Left*: The tree for a set with no repeated elements, $\{2, 1, 5, 8, 7, 6\}$. *Right*: The tree after a rotation of π .

Construction of the tree

We discuss the tree in three parts, the *base* and the left and right *gadgets*. The base and the left gadget will remain stationary while the right gadget spins about an edge on the base. If and only if element uniqueness holds, the two gadgets will not collide.

We begin our construction by drawing the *base*, a four-vertex chain from $(0, 0, 0)$ to $(0, 0, -1)$ to $u = (N + \frac{3}{2}, 0, -1)$ to $v = (N + \frac{3}{2}, 0, 0)$.

For the left gadget, we first draw a vertical edge from $(0, 0, 0)$ to $(0, 0, \max\{S\})$. We call this edge the *stem* of the gadget. For each s_i , we connect a new edge $\overline{(0, 0, s_i)(i, 0, s_i)}$ to the stem. In other words, from the stem we build an edge of length i extending to the right at height s_i . If s_i is not unique, we have two overlapping edges, but for the sake of argument we allow the intersection. (We will avoid this problem later when we create the chain.)

We create the right gadget similarly, except each edge has length $N - i + 1$. We first draw the stem from $(N + \frac{3}{2}, 0, 0)$ to $(N + \frac{3}{2}, 0, \max\{S\})$. For each s_i , we connect a new edge $\overline{(N + \frac{3}{2}, 0, s_i)(2N + \frac{3}{2} - i, 0, s_i)}$ to the stem. In other words, from the stem we build an edge of length $N - i + 1$ extending to the right at height s_i . (Again, ignore the problems of self-intersection if s_i is not unique.)

Examples are shown in figures 5.5 and 5.6.

We now perform a complete rotation (of angle 2π) at $\overline{uv}$. Since the tree is orthogonal, every edge will stay at the same height during the motion. Therefore the only




Figure 5.6: *Left*: The tree for a set with a repeated element, $\{2, 1, 5, 8, 7, 5\}$. The endpoints of the overlapped edges are encircled. *Right*: the tree after a rotation of π . The bold segment indicates the collision.

possibility of a collision is that two edges of the same height, and thus corresponding to elements of the same value, collide. Recall that the two stems are distance $N + \frac{3}{2}$ apart. Therefore if and only if the lengths of two edges of the same height on opposite gadgets sum to at least $N + \frac{3}{2}$ will there be a collision. Since the lengths of the two edges corresponding to the same element sum to $N + 1$, we have a collision only if two edges corresponding to different elements collide. This will necessarily happen if two elements s_i and s_j ($i < j$) have the same value. The left gadget edge corresponding to s_j has length j , and the right gadget edge corresponding to s_i has length $N - i$. Since $i < j$, the two edges have total length at least $N + 2$. If $s_i = s_j$, the edges would be at the same height and therefore collide.

Construction of the chain

We create a three-dimensional chain that does not self-intersect and behaves exactly like the tree. In fact, from the viewpoint of $y = \infty$, the tree and the chain look identical.

The base is the same, a four-vertex subchain from $(0, 0, 0)$ to $(0, 0, -1)$ to $u = (N + \frac{3}{2}, 0, -1)$ to $v = (N + \frac{3}{2}, 0, 0)$. For the left gadget, we create the first edge similarly to the tree by drawing an edge up from $(0, 0, 0)$ to $(0, 0, s_1)$ and another horizontally to $(1, 0, s_1)$. To avoid intersections in the chain we exploit the third dimension, y . We draw an edge back to a point just in front of the stem, at $(0, -\frac{1}{4N}, s_1)$. Now we are




Figure 5.7: Oblique view of the base and left gadget for the set $(2, 1, 5, 8, 7, 5)$. There are two pairs of edges at height $z = 5$ corresponding to the repeated element. (*Shown in stereo in figure 5.8.*)

free to draw a vertical edge from $(0, -\frac{1}{4N}, s_1)$ to $(0, -\frac{1}{4N}, s_2)$, which places us at height s_2 in preparation for the next two edges, to $(2, 0, s_2)$ and back to $(0, -\frac{2}{4N}, s_2)$. We continue this pattern drawing edges to $(0, -\frac{i-1}{4N}, s_i)$, $(i, 0, s_i)$, $(0, -\frac{i}{4N}, s_i)$, and so on until the gadget is complete. We construct the right chain in the exact same fashion, except that we reverse the labeling of the elements from $s_1, \dots, s_n$ to $s_n, \dots, s_1$.

An illustration of the left gadget and the base is in figure 5.7 and shown in stereo in figure 5.8; a bird's eye view of the whole chain is in figure 5.9.

We perform a complete rotation (of angle 2π) at $\overline{uv}$. Since all edges are within radius N of their respective gadget stem, and the stems are $N + \frac{3}{2}$ apart, no part of either gadget will enter a cylinder of radius $\frac{1}{2}$ around the stem of the opposite gadget. Because all the vertical edges are contained in a cylinder of radius $\frac{1}{4}$ around the stems, only the horizontal edges corresponding to elements in the set S can collide. Thus the behavior of the chain exactly mimics that of the tree. The entire chain is illustrated in figure 5.10 and shown in stereo in figure 5.11. ■

As an aside, this also implies a lower bound for the computing the feasibility of a




Figure 5.8: *Stereogram*: Oblique view of the base and left gadget for the set $(2, 1, 5, 8, 7, 5)$.



Figure 5.9: View of the chain from above ($z = +\infty$). Not to scale; the y -dimension is magnified for clarity.




Figure 5.10: Oblique view of the chain for the set $(2, 1, 5, 8, 7, 5)$. If a dihedral rotation is performed at $\overline{uv}$, the edges indicated at the arrows, which correspond to the repeated element 5, will collide. (*Shown in stereo in figure 5.11.*)




Figure 5.11: *Stereogram*: Oblique view of the chain for the set $(2, 1, 5, 8, 7, 5)$.

pivot operation of section 4.1. We can easily convert the above chain into a polygon with the addition of a few edges to obtain the polygon of figure 5.12. Performing a pivot on the diagonal determined by the red vertices has the same effect as a dihedral rotation at the edge $\overline{uv}$ of the chain. Thus determining the feasibility of a pivot also requires $\Omega(n \log n)$ time in the algebraic decision tree model.

5.2 Preprocessing Chains for Dihedral Rotations

The previous results apply to single motions, but the folding of a molecule is a complex process. If it were modeled with dihedral rotations, one would expect at least hundreds or thousands of motions. In this section we study the complexity of determining the feasibility of a sequence of dihedral rotations. To compute each motion as if it were a separate problem seems inefficient as the chain always maintains its edge lengths and vertex-angles. An intuitive goal would be to preprocess the chain so




Figure 5.12: *Stereogram*: Converting the chain into a polygon.

that the feasibility of each ensuing dihedral rotation could be determined in $o(n \log n)$ time.

We show two problems concerning multiple dihedral rotations to be 3SUM-hard. A problem is 3SUM-hard if there is a subquadratic reduction from the following.

3SUM: *Given a set of integers, do there exist three elements that sum to zero?*

3SUM-hardness was introduced by Gajentaan and Overmars [61] to provide evidence in support of conjectured $\Omega(n^2)$ lower bounds for several problems. The best known algorithm for the 3SUM problem on n integers runs in time $\Theta(n^2)$. Quadratic lower bounds have been proven in restricted models by Erickson [50], but the strongest lower bound yet proven in any general model of computation is $\Omega(n \log n)$.

In section 5.2.1, we consider the case where we wish to preprocess a chain such that we can quickly answer a single dihedral rotation query whose location (the edge on

which the rotation occurs) is not known in advance. We show that preprocessing the chain and answering n queries is 3SUM-hard. Thus $\Omega(n^2)$ preprocessing is almost certainly required to achieve sublinear query time.

In section 5.2.2, we prove a much stronger result. If we wish to determine the feasibility of a sequence of dihedral rotations, then answering each query in sublinear time is impossible after any amount of preprocessing unless there exists a family of decision trees of subquadratic depth that solve the 3SUM problem. Since this seems unlikely in light of current research [50], answering a sequence of n queries is almost certainly 3SUM-hard regardless of the preprocessing.

5.2.1 Single dihedral rotation queries

Suppose we are given a chain and are asked to preprocess so that we can quickly determine if a dihedral rotation is feasible when given an edge and an angle ϕ . We are not interested in altering the chain but only in whether this one rotation could be performed.

Single Dihedral Rotation Query: *Preprocess a chain so that a single arbitrary dihedral rotation query can be answered efficiently.*

We are interested in both the preprocessing time and the time needed to answer a query. There is an inherent tradeoff between the two durations. We could spend $\Omega(n^2 \log n)$ time computing the degrees of freedom for all possible dihedral rotations. If we store the results in a table, then any query could be answered simply by looking up the result. On the other hand, with no preprocessing, each query requires $\Omega(n \log n)$ time.

Our discussion provides strong evidence for the following conjecture, which demonstrates the link between the time spent preprocessing and the time spent answering each query.

Conjecture. *In any scheme to preprocess a chain of n edges to answer Single Dihedral Rotation Queries, either*

- *the preprocessing requires $\Omega(n^2)$ time, or*
- *each ensuing query requires $\Omega(n)$ time.*

In other words, we conjecture that unless quadratic time is spent preprocessing in the worst case, then the query time will be at least linear. We support this by proving that the preprocessing and n queries is 3SUM-hard. Rather than reducing 3SUM to Single Dihedral Rotation Query (SDRQ), we instead reduce from the following problem, often referred to as 3SUM'.

3SUM': *Given three sets of integers, do there exist three elements, one from each set, that sum to zero?*

Although 3SUM is the more famous problem of the pair, 3SUM' will be easier to use in achieving our reduction to SDRQ. This poses no additional complication since the two problems are reducible to one another in linear time, only changing the complexity of the input by a constant factor [61]. Therefore a reduction from a 3SUM problem on n integers is equivalent to a reduction from a 3SUM' problem on three sets of which the largest contains n integers.

Because the time complexity of the 3SUM (or 3SUM') problem is unknown, we adopt the convention of using $3SUM(n)$ to denote the time complexity of the 3SUM problem on n integers.

Theorem 16. *In any scheme to preprocess a chain of n edges to answer Single Dihedral Rotation Queries, the preprocessing and n queries is 3SUM-hard.*

Proof. We reduce 3SUM' to SDRQ. In particular, given any 3SUM' problem, we create, in $O(n \log n)$ time, a chain such that $O(n)$ dihedral rotations queries answer




Figure 5.13: Reducing 3SUM' to Single Dihedral Rotation Query.

the 3SUM' problem. If time P is spent preprocessing and time Q is necessary per query, then $P + nQ = \Omega(3SUM(n))$.

We will need to transform the three sets A , B , and C for our reduction. Let m be the maximum absolute value of any element in $A \cup B \cup C$. We build three sets A' , B' , and C' as follows. We create A' by subtracting $2m$ from every element in A , B' by multiplying every element of B by $-\frac{1}{2}$, and C' by adding $2m$ to every element of C . Thus $A' = \{a - 2m : a \in A\}$, $B' = \{-b/2 : b \in B\}$, $C' = \{c + 2m : c \in C\}$. If and only if some triplet $(a \in A) + (b \in B) + (c \in C) = 0$, there is some triplet $(a - 2m \in A') - 2(b/2 \in B') + (c + 2m \in C') = a' - 2b' + c' = 0$. Thus we have an equivalent 3SUM' problem on A' , $-2B'$, C' .

We create a planar chain as illustrated in figure 5.13. (The explanation and illustration of the chain are much easier under the assumption that the sets are sorted, although one can create a chain that does not require sorting by exploiting the third dimension.) The chain has a *comb* with teeth facing up on the left end, an axis-parallel *staircase* in the middle, and a *comb* with teeth facing down on the right. For each element $a' \in A'$, we create a very slim upward tooth with x -coordinate a' on the left comb. We then create the orthogonal staircase such that a vertical edge exists with x -coordinate b' for every $b' \in B'$, and finally, for each element $c' \in C'$ we create a very slim downward tooth with x -coordinate c' on the right comb.

We now ask $O(n)$ queries; namely, can a dihedral rotation of angle 2π be performed at each vertical edge in the orthogonal staircase? Since the chain is planar, the only possibility for an intersection is when the rotation has reached π . At this point, one comb and part of the staircase have been reflected across the vertical edge, as in



Figure 5.14: An edge spin at the indicated vertical edge.




Figure 5.15: *Left*: There exist no $a' \in A'$ and $c' \in C'$ such that $a' - 2b' + c' = 0$. *Right*: The $3SUM'$ is solved, as $a' - 2b' + c' = 0$.

figure 5.14.

Because the rotation is performed at a vertical edge, no edge changes height. Thus the staircase cannot self-intersect. Each comb stays individually rigid, so neither comb can self-intersect. Furthermore, because each vertical edge in the staircase is at distance at most m from another staircase edge, but at distance at least $3m/2$ from any edge of a comb, the dihedral rotation cannot cause a comb and the staircase to intersect. Therefore the only possible intersection during the rotation occurs between the two combs. Since the height of an edge is maintained throughout the motion, intersections are only possible at the teeth.

Suppose two teeth with x -coordinates a' and c' intersect during a dihedral rotation at a staircase edge with x -coordinate b' . Since we have in essence performed a reflection across the line $x = b'$, an intersection occurs if and only if the distance from the tooth at a' to the line $x = b'$ is equal to the distance from the tooth at c' to the same line. This implies that $b' - a' = c' - b'$, or differently stated, that $a' - 2b' + c' = 0$. Therefore, if and only if there exist some $a' \in A'$ and $c' \in C'$ such that $a' - 2b' + c' = 0$, the chain self-intersects when we perform a dihedral rotation on the vertical edge at x -coordinate b' . Figure 5.15 illustrates two examples.

We perform a query on each vertical edge of the staircase. If a dihedral rotation

of angle 2π is possible at every such edge, then there exist no a', b', c' for which $a' - 2b' + c' = 0$. If one of the rotations is not possible, it must be due to a potential collision, which implies the existence of such a triplet a', b', c' .

Since the preprocessing (of time complexity P) and $O(n)$ queries (of time complexity Q each) answer the 3SUM' question, $P + nQ = \Omega(3\text{SUM}(n))$. ■

5.2.2 Multiple dihedral rotation queries

As mentioned earlier, we are often not interested in one motion but rather in a long sequence of motions. Instead of determining the feasibility of a single dihedral rotation, we would like to perform rotations and determine quickly if each is feasible. In other words, we might ask, “Can we perform a dihedral rotation at edge e_1 by angle ϕ_1 , then at edge e_2 by angle ϕ_2 , then at edge e_3 by angle ϕ_3 , and so on?” This leads us to the following problem.

Multiple Dihedral Rotation Query: *Preprocess a chain such that one can repeatedly determine the feasibility of dihedral rotations and perform them if feasible, thereby altering the chain after each query.*

Again, we are interested in both the preprocessing time and the time to answer each query. We demonstrate that there does not exist a scheme in which the time to answer a query is sublinear unless there exists a family of subquadratic decision trees to solve 3SUM. This seems unlikely in light of results by Erickson [50] which demonstrate a quadratic bound on the size of such a tree, although in a restricted model of computation. Even if such a family of decision trees does exist, the preprocessing time would be at least the time required to construct the n^{th} tree.

The concept of using a family of decision trees to analyze complexity in this fashion was introduced by Meyer auf der Heide [105], who proved that for the NP-complete

problem KNAPSACK of input length n there exists a linear search tree of depth $O(n^4 \log n)$ which solves it. However, this result is not easily exploitable by any general algorithm for KNAPSACK since there is no known polynomial algorithm to build each tree. To use such a scheme directly, one would have to construct the n^{th} linear search tree, which requires exponential space. Because the algorithms are different for each input length, the family of decision trees is called a *nonuniform algorithm*.

It may seem contradictory that a problem whose complexity is believed to be super-polynomial (unless $P = NP$) can be solved by a family of decision trees of polynomial depth. The difference lies in the fact that an algorithm that solves KNAPSACK must accept an input string of *any* length and output whether the input string satisfies the conditions of the problem. A nonuniform algorithm is a collection of trees such that for any *fixed* length there exists a tree which accepts an input string of that length and outputs whether the string satisfies the conditions of the problem. To even simulate a general algorithm with a nonuniform algorithm one would need infinitely many trees—or rather to build the appropriate tree after counting the length of the input. Thus the existence of an infinite family of subquadratic-depth decision trees which solve 3SUM would not contradict the conjecture that 3SUM has a time complexity of $\Omega(n^2)$.

We now demonstrate the link between a subquadratic nonuniform algorithm for 3SUM' and the Multiple Dihedral Rotation Query problem. If such a nonuniform algorithm exists, let $3SUMTREE(n)$ denote the time complexity of constructing the decision tree which solves 3SUM' where the largest set has at most n elements.

Theorem 17. *If 3SUM has quadratic time complexity, then in any scheme to preprocess a chain of n edges to determine the feasibility of a sequence of dihedral rotations, either*

- *the preprocessing requires $\Omega(3SUMTREE(n))$ time, or*

- *answering a dihedral rotation query requires linear time in the worst case.*

Furthermore, if there does not exist a subquadratic nonuniform algorithm which solves 3SUM, then there does not exist a scheme to preprocess a chain of n edges such that each dihedral rotation query can be answered in sublinear time.

Proof. We reduce the construction of a subquadratic nonuniform algorithm for 3SUM' on sets of size n to Multiple Dihedral Rotation Query with the following approach. Suppose we are given the number n and asked to construct an algorithm for the 3SUM' problem where each set has n elements. We create a chain whose structure is determined by the number n , and spend time P preprocessing it to answer multiple dihedral rotation queries. When the preprocessing has finished, the sets A , B , and C are revealed. We perform $O(n)$ dihedral rotations, each in time Q , to move the chain into a configuration similar to the one in figure 5.13 for the three sets. After $O(n)$ additional rotations, as in the proof of theorem 16, the 3SUM' is solved. If $nQ = o(n^2)$, then the 3SUM' on A , B , and C has been solved in subquadratic time. Therefore constructing a subquadratic nonuniform algorithm for 3SUM' has been reduced to preprocessing the chain. If $Q = o(n)$, P must be $\Omega(\text{3SUMTREE}(n))$. If there exists no subquadratic nonuniform algorithm for 3SUM', then Q must be of linear complexity.

We build a planar chain for the given n . In figure 5.16 are the three basic building blocks, which we refer to as *gadgets*. On the left is the gadget used to construct the left comb. With a dihedral rotation at $\overline{uv}$, we can bring a'_1 and a'_2 any distance apart as we choose, to a maximum of the length of the gadget when fully extended. Dihedral rotations at α_1 and α_2 bring a_1 and a_2 (and the rest of the chain) back into the original plane. Thus the only end result after the three rotations is to shrink the distance between the two teeth, leaving the rest of the chain unaffected except for a translation. The center gadget performs a similar function, allowing the distance between two steps in the staircase to be selected with a dihedral rotation at $\overline{uv}$,




Figure 5.16: *Left*: The gadget for setting the distance between two consecutive teeth on the left comb. *Center*: The gadget for consecutive elements of the staircase. *Right*: The gadget for consecutive elements of the right comb.




Figure 5.17: The chain created in the proof of theorem 17.

followed by rotations at β_1 and β_2 to bring the remainder of the chain back to the same plane. The right gadget carries out the same function for the right comb.

The full chain, constructed of these gadgets (and two long steps of length 3.5), is illustrated in figure 5.17. We now perform any preprocessing, of time P , to determine quickly the feasibility of a future sequence of dihedral rotations. Then we are given the three sets, A , B , and C , for which the 3SUM' is asked.

Our reduction will be easier if we assume that all the sets A , B , and C each have the same number of elements. If some set has k fewer elements than another, then we can add elements $\{i/4n\}$, for $i \in \{1 \dots k\}$, to the smaller set without affecting the outcome of the 3SUM'. (Because the other elements of the sets are integers, no one, two, or three of these fractions can contribute to a triplet of elements which sum to zero).

As in the proof of theorem 16, we transform the sets A , B , and C . Let m be the maximum absolute value of any element in $A \cup B \cup C$. We create A' by dividing each element in A by $2m$ and then subtracting 2.5; we create B' by multiplying every

element of B by $\frac{-1}{4m}$, and C' dividing each element in C by $2m$ and then adding 2.5. Thus $A' = \{a/2m - 2.5 : a \in A\}$, $B' = \{-b/4m : b \in B\}$, and $C' = \{c/2m + 2.5 : c \in C\}$. Note that if and only if some triplet $(a \in A) + (b \in B) + (c \in C) = 0$, there is some triplet $(a/2m - 2.5 \in A') - 2(b/4m \in B') + (c/2m + 2.5 \in C') = a' - 2b' + c' = 0$. We have now have an equivalent 3SUM' problem on $A', -2B', C'$.

The sets A' , B' , and C' have the property that every element in A' is between -2 and -3 ; every element in B' is between $-\frac{1}{4}$ and $\frac{1}{4}$; every element of C' is between 2 and 3. The teeth of the comb can be separated by up to distance 1, and the steps of the staircase by up to $\frac{1}{2}$. Furthermore, the distance between the combs and the staircase can be made to be up to 3.5. Therefore the teeth of the left comb can be placed such that the x -coordinates of the teeth correspond to the elements of A' . The same applies for the vertical steps of the staircase (corresponding to B'), and the teeth of the right comb (corresponding to C').

We sort the three sets, and place the leftmost tooth of the left comb at the minimum value in A' . Then with three dihedral rotations, we bring the second tooth to the x -coordinate of the second value in A' . We repeat this procedure until the left comb has teeth at all x -coordinates of A' . We use the long triplet of edges (marked with length 3.5) to bring the first vertical edge of the staircase into position, and likewise bring all the stairs into position. We bring the right comb into position in the same manner. An example of the final configuration is in figure 5.18 and shown in stereo in figure 5.19.

We can prove that the chain does not self-intersect during these dihedral rotations by examining the gadgets in figure 5.16. Note that our three motions perform a rotation at α_1 (or β_1, γ_1) by an angle ϕ in one direction, a rotation at $\overline{uv}$ by 2ϕ in the other direction, and a rotation at α_2 (or β_2, γ_2) by ϕ to bring the remainder of the chain back into its original plane. Since our sets are sorted, each successive tooth or stair proceeds to the right, thus $\theta < \pi$. Therefore the chain stays monotonic




Figure 5.18: The chain, moved into position for the newly revealed A' , B' , and C' . (Shown in stereo in figure 5.19).

in the x -direction (always proceeding either vertically or to the right), so it cannot self-intersect during any of these motions.

Once the chain is set for A' , B' , and C' , we perform dihedral rotations of angle 2π at every vertical edge in the staircase which corresponds to an element of B' . Just as in the proof of theorem 16, the chain self-intersects if and only if there exists a triplet $a' - 2b' + c' = 0$, which answers the $3\text{SUM}'$ problem.

We spent time P preprocessing the chain before the sets were revealed, and $\Theta(n \log n)$ time sorting and time $10nQ$ for $10n$ dihedral rotations ($9n$ to set the $3n$ gadgets, then n to answer the $3\text{SUM}'$, each taking time Q) after the sets were revealed. If $P \neq \Omega(3\text{SUMTREE}(n))$, then $n \log n + nQ = \Omega(3\text{SUM}(n))$, and thus $Q = \Omega(3\text{SUM}(n)/n)$. ■




Figure 5.19: *Stereogram*: The chain, moved into position for the newly revealed A' , B' , and C' .

Chapter 6

Endpoint-to-Endpoint Distance and Canonical Configurations

In section 1.2 we discussed how flexible polymers are often modeled as chains where the only motions permitted are dihedral rotations. The space of possible configurations of such polymers is vast and has been researched by several physicists and statisticians. Its most often studied property is the mean squared distance between the endpoints of the chain¹, which relates directly to many physical properties such as the extent to which the polymer diffuses and scatters light [17, 52, 151]. In 1948, the physicist William Taylor [151] discussed the need to describe the possible configurations.

It is by now a familiar concept that long-chain molecules (in non-crystalline states) may assume more or less coiled and continuously variable configurations as the result of internal rotations or torsional motions about the individual bonds of the chain. As a result of this indefiniteness it becomes necessary to discuss the properties of such molecules on a statistical ba-

¹Mean squared distance between endpoints has been studied in [17, 25, 41, 42, 52, 57, 60, 69, 75, 90, 91, 92, 99, 102, 104, 113, 129, 130, 147, 148, 151, 173].

sis. [151]


During the same year as Taylor, Benoit [17] independently discovered the same results and explicitly made reference to the minimum and maximum possible distances between the endpoints of the chain.

Une telle molécule pourra donc présenter un grand nombre de configurations différentes que l'agitation thermique lui fera prendre par hasard. Il en résulte que la distance entre les deux extrémités de la molécule, lorsque celle-ci est par exemple en solution diluée, variera continuellement entre deux valeurs, l'une minimum, obtenue quand les extrémités se touchent, l'autre, maximum, obtenue quand la chaîne est complètement étirée². [17]

The problems alluded to by Benoit, that of finding the configurations with minimum and maximum distances between the endpoints, is an underlying subproblem in any attempt to compute the distribution of the distance between the endpoints of the chain [173]. The model most often used in polymer physics restricts the chain to the lattice, where each link is thus of unit length and each vertex-angle either 0 or $\pi/2$.

We consider the general case where the chain may assume any configuration which maintains its arbitrary edge lengths and vertex-angles. We begin by proving the problems to be NP-hard when the chain is restricted to two dimensions. In the sections following, we discuss the difficulty of finding the maximum and minimum distance between the endpoints over all three-dimensional configurations and present an algorithm which efficiently approximates the maximum possible distance. These

²Such a molecule will therefore be able to exhibit a large number of different configurations which are randomly determined by thermal agitation. This causes the distance between the two endpoints of the molecule, when for example in diluted solution, to vary continuously between two values: the minimum, obtained when the endpoints touch, and the maximum, obtained when the chain is completely taut.

Figure 6.1: Definitions of ℓ_i and θ_i .

results imply the difficulty of determining whether a chain can be brought into a planar configuration, a subject which we discuss in greater detail in section 6.5.

6.1 Preliminaries

To avoid the cumbersome phrase, “distance between the endpoints,” we henceforth refer to the distance between the endpoints of a chain as its *span*. We also use of a few conventions with regard describing the lengths of links and the angles between them. Let the chain be composed of n vertices and $n - 1$ links. For brevity, we write the length $|v_i v_{i+1}|$ as ℓ_i (for $1 \leq i \leq n - 1$). We write the *turning angle* at vertex v_i as θ_i (for $2 \leq i \leq n - 1$). The turning angle at v_i is the angle between the rays $\overrightarrow{v_{i-1}v_i}$ and $\overrightarrow{v_i v_{i+1}}$, and is equal to $\pi - \angle v_{i-1}v_i v_{i+1}$. We refer to the turning angles in this chapter rather than the interior angles between consecutive links to simplify the mathematics used herein. The definitions of these quantities are illustrated in figure 6.1.

In order to discuss computing the configuration space of the chains, we must decide how the chains are to be represented. We will remain consistent with earlier chapters where the chain is represented as a list of vertices, $v_1, v_2, \dots, v_n$. Therefore it is not possible to create a chain with specific turning angles, since this would require exact computations of trigonometric functions, except for special cases such as multiples of $\pi/2$. In this chapter, if we create a chain, we will refer to the lengths of the edges

and the cosines of the turning angles, which can be exactly translated to and from the coordinates of the vertices (in a model which allows the computation of square roots).

6.2 Span of Two-Dimensional Configurations

We begin by considering the two-dimensional analogue of the problem, that of finding the maximum and minimum possible spans over all *planar* configurations of the chain. While this is not the problem usually considered by the physics community, it provides insight into the difficulty of the three-dimensional version and will be the backdrop for results in section 6.5 and chapter 7.

Maximum Planar Span: *Given a polygonal chain and a number k , is there a (possibly self-crossing) planar configuration whose span is at least k ?*

Minimum Planar Span: *Given a polygonal chain and a number k , is there a (possibly self-crossing) planar configuration whose span is at most k ?*

We show these two decision problems to be NP-hard. For the time being, we ignore self-intersections of the chain. We will later demonstrate that the problems remain NP-hard even when the chain is restricted to be simple.

Because the chain is planar, there are a finite number of configurations of the chain. As we traverse the chain from one endpoint to the other, at each vertex we have either a left turn or a right turn; due to the fixed vertex-angles there is no other freedom available. A non-deterministic algorithm needs only to select one of the 2^{n-2} configurations and calculate its span to determine if it is at least or at most k . Thus it would seem that the two problems are in NP, which begs the question of proving that they are NP-complete as opposed to just NP-hard. However, it is important to

realize that because square roots are not computable by a Turing machine, computing Euclidean distance is not possible within NP.

To prove that the two problems are NP-hard, we perform a reduction from the problem Partition.

Partition: *Given a set of positive integers S , can it be partitioned into two disjoint sets S_a and S_b such that $\Sigma(s : s \in S_a) = \Sigma(s : s \in S_b)$?*

Although the reductions for the two problems both rely on Partition, the method of proof for each is quite different. We first consider Maximum Planar Span then Minimum Planar Span.

We require lemmas 2 and 3, which are proven on pages 31 and 32.

Lemma 2. *The functions $\sin x$ and $\cos x$ can be computed within an error of ε in time $O(x + \log(1/\varepsilon))$ for any $\varepsilon > 0$. Furthermore, the approximation can be guaranteed to be either greater or less than the true value.*

Lemma 3. *For $0 \leq x \leq \pi$, let $\widehat{\cos} x$ be an approximation which is less than $\cos x$ by at most ε . Then $\arccos(\widehat{\cos} x) - x \leq \sqrt{2\varepsilon}$.*

We are now ready to prove the following theorem.

Theorem 18. *Maximum Planar Span is NP-hard.*

Proof. Given a set of integers as the instance to a Partition problem, we create a chain C and compute a number k such that the Partition problem is equivalent to the question, “Is the maximum planar span of C at least k ?”

Let $S = \{s_1, s_2, \dots, s_n\}$ be the n integers in the Partition problem, and let σ be the sum $s_1 + s_2 + \dots + s_n$.

We first create a chain of $n + 1$ edges (and thus $n + 2$ vertices), where each link is of less than unit length, such that the turning angles at the vertices $v_2, v_3, \dots, v_{n+1}$ are proportional to the integers of S . Let $\theta_2 = s_1/\sigma$, $\theta_3 = s_2/\sigma$, et cetera.

We have chosen to represent the chain by lengths of edges and cosines of the turning angles rather than the measures of the angles themselves. Therefore we must compute the cosines of these angles. By lemma 2 we can approximate the cosines within an error of $1/(32n^2\sigma^2)$ within time $O(\log n + \log \sigma)$. By lemma 3, the turning angles θ_i will differ from s_{i-1}/σ by at most $1/(4n\sigma)$.

The total curvature of the chain is no more than $1 + 1/(4\sigma)$ radians. Therefore the two extreme edges of the chain are within $1/(4\sigma)$ radians of being parallel if and only if there exists a partition of the turning angles, such that the sum of the measures of the angles turning left is equal to the sum of the measures of those turning right (minus the possible accumulated error of $1/(4\sigma)$). If no partition exists, then the orientations of the two extreme edges must differ by at least angle $1/\sigma$ minus the possible error, or $3/(4\sigma)$.

Now let us expand the two extreme links in the chain to some large length l . We now have a chain of one long link $\overline{v_1v_2}$ of length l , followed by $n - 1$ short links of length one, followed by another long link $\overline{v_{n+1}v_{n+2}}$ of length l . This situation is illustrated in figure 6.2. Note that since the total curvature of the chain is less than $\pi/2$, l_1 lies strictly to the left of l_{n+1} . This implies that if the two extreme links differ in orientation by some angle ϕ , then by elementary trigonometry the difference in x -coordinates of the endpoints, and thus the span of the chain, must be at least $l + l \cos \phi$. A visual proof is provided in figure 6.3. Also illustrated in the figure is the distance between the endpoints of the links in this position, $2l \cos(\theta/2)$. Adding the $n - 1$ short links could at most lengthen the span by the sum of their lengths. Thus the span of the chain can be no more than $2l \cos(\theta/2) + n - 1$.

If $\phi \leq 1/(4\sigma)$, then the span of the chain must be at least $l + l \cos(1/(4\sigma))$. If there is no partition of the turning angles, then the extreme links must differ in orientation by at least $3/(4\sigma)$ radians, and the span of the chain would be at most $2l \cos(3/(8\sigma)) + n - 1$. We wish to choose l carefully such that if the span is at least





Figure 6.2: The chain constructed in the proof of theorem 18.

Figure 6.3: Two links of length l differing in orientation by θ .

$l + l \cos(1/(4\sigma))$ then the two extreme links must be within $1/(4\sigma)$ of parallel.

For reasons to be made clearer in chapter 7, we select l such that even if the span is $l + l \cos(1/(4\sigma)) - 1$, the extreme links must be within $1/(4\sigma)$ of parallel. This is obtained if $l > n/[\cos(1/(4\sigma)) - \cos(3/(8\sigma))]$.


$$\begin{aligned}
 l &> \frac{n}{\cos(1/(4\sigma)) - \cos(3/(8\sigma))} \\
 l[\cos(1/(4\sigma)) - \cos(3/(8\sigma))] &> n \\
 l \cos(1/(4\sigma)) &> l \cos(3/(8\sigma)) + n \\
 l + l \cos(1/(4\sigma)) &> l + l \cos(3/(8\sigma)) + n \\
 l + l \cos(1/(4\sigma)) &> 2l \cos(3/(8\sigma)) + n \\
 l + l \cos(1/(4\sigma)) - 1 &> 2l \cos(3/(8\sigma)) + n - 1
 \end{aligned}$$

Because we cannot compute differences of cosines on a Turing machine, we instead compute an lower bound.

$$\begin{aligned}
 \cos(1/(4\sigma)) &= 1 - \frac{1}{2!(4^2)\sigma^2} + \frac{1}{4!(4^4)\sigma^4} - \frac{1}{6!(4^6)\sigma^6} + \dots \\
 \cos(1/(4\sigma)) &> 1 - \frac{1}{2!(4^2)\sigma^2} = 1 - \frac{1}{32\sigma^2} \\
 \cos(3/(8\sigma)) &= 1 - \frac{3^2}{2!(8^2)\sigma^2} + \frac{3^4}{4!(8^4)\sigma^4} - \frac{3^6}{6!(8^6)\sigma^6} + \dots \\
 \cos(3/(8\sigma)) &< 1 - \frac{3^2}{(2)2!(8^2)\sigma^2} = 1 - \frac{9}{256\sigma^2} \\
 \cos(1/(4\sigma)) - \cos(3/(8\sigma)) &> \frac{1}{256\sigma^2}
 \end{aligned}$$

Substituting this expression into the earlier inequality for l , we obtain that if $l = 256n\sigma^2$, that the two links are within $1/(4\sigma)$ of parallel if and only if the span of the chain is at least $l[2 - 1/(32\sigma^2)]$. We can bound this expression by

$$l[2 - 1/(32\sigma^2)] = (256n\sigma^2)[2 - 1/(32\sigma^2)] = (512n\sigma^2) - (8/n).$$

Figure 6.4: Triangle $\triangle abc$ embedded between two planes.


If the two links are not parallel, then the span is at most $(512n\sigma^2) - (8/n) - 1$.

The question, “Is the maximum planar span of the chain at least $(512n\sigma^2) - (8/n)$?” is equivalent to the question, “Can the chain be laid in the plane such that the two extreme edges lie almost parallel?” This last question is itself equivalent to the Partition problem on S . Clearly the chain can be constructed in polynomial time since the number of vertices is polynomial in n , and all values are polynomial in σ . Thus Maximum Planar Span is NP-hard. ■

We can generalize the above theorem to prove the hardness of finding the maximum span of a three-dimensional chain when it is constrained to lie between two close parallel planes. The intractability of this problem will be a vital lemma in chapter 7.

Maximum Span Between Two Planes: *Given a polygonal chain, a number k , and a number w , is there a (possibly self-crossing) configuration which lies between two parallel planes distance w apart, and whose span is at least k ?*

Our reduction will use the same chain as that for the proof of theorem 18. Recall that all turning angles in that chain are less than one radian. We now prove a lemma about triangles embedded between two planes, illustrated in figure 6.4, where one of the turning angles of the triangle is less than one radian (and thus the vertex-angle is greater than $\pi - 1$).

Figure 6.5: The distance $|x'y'|$ in relation to $|xy|$.

Lemma 21. Let $\triangle abc$ be a triangle where $\pi - 1 \leq \angle abc \leq \pi$. Suppose $\triangle abc$ is embedded between two horizontal planes separated by a distance w , where w is small enough such that $w \leq |ab|/2$, $w \leq |bc|/2$, and $w \leq \sqrt{|ab||bc|}$. Let $\triangle a'b'c'$ be the orthogonal projection of $\triangle abc$ onto the horizontal plane. Then $\angle abc - w\sqrt{2}/\sqrt{|ab||bc|} \leq \angle a'b'c' \leq \angle abc + w\sqrt{2}/\sqrt{|ab||bc|}$.

Proof. Let $\overline{xy}$ be any line segment between the two planes. Then the length of the edge $\overline{x'y'}$ in the projection is at least $\sqrt{|xy|^2 - w^2}$. Figure 6.5 provides a visual proof.

Consider the lengths $|a'b'|$, $|b'c'|$, and $|a'c'|$.

$$\begin{aligned} \sqrt{|ab|^2 - w^2} &\leq |a'b'| \leq |ab| \\ \sqrt{|bc|^2 - w^2} &\leq |b'c'| \leq |bc| \\ \sqrt{|ac|^2 - w^2} &\leq |a'c'| \leq |ac| \end{aligned}$$

By the law of cosines, $|a'c'|^2 = |a'b'|^2 + |b'c'|^2 - 2|a'b'||b'c'| \cos \angle a'b'c'$. We obtain the following bounds for $\cos \angle a'b'c'$. Since $\angle abc > \pi - 1$,

$$|ac|^2 \geq |ab|^2 + |bc|^2 - 2 \cos(\pi - 1) |ab||bc| \geq |ab|^2 + |bc|^2 + |ab||bc|.$$

Since $\sqrt{|ac|^2 - w^2} \leq |a'c'|$ and $w^2 < |ab||bc|$, we know that $\cos \angle a'b'c'$ is negative.

$$\begin{aligned}
\cos \angle a'b'c' &= -\frac{|a'b'|^2 + |b'c'|^2 - |a'c'|^2}{2|a'b'||b'c'|} \\
\cos \angle a'b'c' &\geq -\frac{(|ab|^2 - w^2) + (|bc|^2 - w^2) - |ac|^2}{2|ab||bc|} \\
\cos \angle a'b'c' &\geq -\frac{|ab|^2 + |bc|^2 - |ac|^2 - 2w^2}{2|ab||bc|} \\
\cos \angle a'b'c' &\geq \cos \angle abc - \frac{w^2}{|ab||bc|}
\end{aligned}$$

We prove an upper bound in a similar fashion.

$$\begin{aligned}
\cos \angle a'b'c' &\leq -\frac{|ab|^2 + |bc|^2 - (|ac|^2 - w^2)}{2\sqrt{(|ab|^2 - w^2)(|bc|^2 - w^2)}} \\
\cos \angle a'b'c' &\leq -\frac{|ab|^2 + |bc|^2 - |ac|^2 + w^2}{2\sqrt{(|ab|^2/2)(|bc|^2/2)}} \\
\cos \angle a'b'c' &\leq \cos \angle abc + \frac{w^2}{|ab||bc|}
\end{aligned}$$

Thus $\cos \angle a'b'c'$ differs from $\cos \angle a'b'c$ by at most $w^2/(|ab||bc|)$. Because both angles are between 0 and π , by lemma 3 if $\cos \angle a'b'c' < \cos \angle abc$, then $\angle a'b'c' \leq \angle abc + w\sqrt{2}/\sqrt{|ab||bc|}$. If $\cos \angle a'b'c' > \cos \angle abc$, then $\angle a'b'c' \leq \angle abc - w\sqrt{2}/\sqrt{|ab||bc|}$. ■

We are ready to prove the following.

Theorem 19. *Maximum Span Between Two Planes is NP-hard.*

Proof. Let $S = \{s_1, s_2, \dots, s_n\}$ be the n integers of a Partition problem, and let σ be the sum $s_1 + s_2 + \dots + s_n$.

As in the proof of theorem 18, we create a chain of $n + 1$ edges (and thus $n + 2$ vertices), where each link is of less than unit length, such that the turning angles at the vertices $v_2, v_3, \dots, v_{n+1}$ are proportional to the integers of S . Let $\theta_2 = s_1/\sigma$, $\theta_3 = s_2/\sigma$, et cetera.

Let w , the separation between the two planes, be less than $1/(16\sqrt{2}n\sigma)$. Consider the orthogonal projection of the chain onto a horizontal plane. If the two end links are within an angle $1/(4\sigma)$ of parallel in three-space, then surely they must also be within $1/(4\sigma)$ of parallel in the projection.

The angles of the chain in the projection are identical to the angles of the chain in three-space, plus or minus an error of at most $w\sqrt{2}$ divided by the length of the shortest link. The tightest upper bound we require on the length of the short links of the chain is that each such edge is at most unit length. If we make each such link at least length $1/2$, then by lemma 21 the difference between any angle of the chain in three-space and its corresponding angle in the projection is at most $2w\sqrt{2} = 1/(8n\sigma)$.

Thus the two end links of projection are parallel if and only if there exists a partition of the angles of the chain, plus or minus an accumulated error of $1/(8\sigma)$. In the proof of theorem 18, we approximated cosines within an error of $1/(32n^2\sigma^2)$ and thus the turning angles θ_i of the chain differed from s_{i-1}/σ by at most $1/(4n\sigma)$. Here we approximate the cosines within an error of $1/(128n^2\sigma^2)$ so that the turning angles differ from s_{i-1}/σ by at most $1/(8n\sigma)$.

Thus the difference between the turning angles of the chain *in the projection* and s_{i-1}/σ differ by at most $1/(8n\sigma) + 1/(8n\sigma) = 1/(4n\sigma)$. This is now the same problem as handled in the proof of theorem 18, and thus the Partition question is equivalent to, “Is the maximum span of the chain between two planes separated by $1/(16\sqrt{2}n\sigma)$ at least $512n\sigma^2 - (8/n)$?” ■

We now return to the Minimum Planar Span problem.

Theorem 20. *Minimum Planar Span is NP-hard.*

Proof. Given a set of integers as the instance to a Partition problem, we create a chain C such that the set has a partition if and only if the minimum planar span of C is less than one. Starting at one endpoint of the chain, we place an edge of length s_1 ,




Figure 6.6: The chain created in the proof of theorem 20.




Figure 6.7: A chain for which the minimum planar span is less than one.

followed by an orthogonal short edge of length $1/n$, followed by an orthogonal edge of length s_2 , followed by an orthogonal short edge of length $1/n$, and so on, until the edge of length s_n . A chain generated by a five-element set is illustrated in figure 6.6. (The chain and the reduction in this proof are similar to those used by Hopcroft, Joseph, and Whitesides [78] to prove that Ruler Folding is NP-complete.)

In any planar configuration of the chain, all edges of length $1/n$ lie parallel to each other and orthogonal to the edges corresponding to elements of the set S . Without loss of generality, assume that the edges of length $1/n$ lie vertically (parallel to the y -axis); thus the longer edges corresponding to elements of the set S lie horizontally (parallel to the x -axis).

If we add an orientation to the chain from one endpoint to the other, then the edges corresponding to elements of the set S will either point left or point right. If both endpoints have the same x -coordinate, then the sum of the lengths of all edges pointing left must equal the sum of the lengths of all edges pointing right. In other words, there must exist a partition of the lengths of the edges, which exists if and only if there is a partition of the integers in S . A chain with such a partition is illustrated in figure 6.7.

If the two endpoints have the same x -coordinate, then the span of the chain is determined solely by the difference in y -coordinate. Since the sum of the lengths of




Figure 6.8: A chain which has no planar span less than one.

all vertical edges is $(n - 1/n)$, the span of the chain is at most this value, which is less than one. Likewise, if there does not exist a partition of S , then the two endpoints must differ in x -coordinate. Since the set S is composed of integers, the difference in x -coordinates of the endpoints, and thus the minimum planar span, must be at least one. An example of such a chain is illustrated in figure 6.8.

Therefore, the question, “Does the chain have a minimum planar span which is less than one?” is equivalent to the question, “Does there exist a partition for S ?” Clearly the chain can be constructed in polynomial time; thus the Minimum Planar Span problem is NP-hard. ■

In the chain created in the proof of theorem 18, the turning angles sum to one radian, so the chain is monotone and cannot self-intersect. In the chain created in the proof of theorem 20, if there exists a configuration whose span is less than one, then it also has a similar configuration which proceeds monotonically upward as in figures 6.7 and 6.8. Therefore the following variants of the problems are also NP-hard.

Simple Maximum Planar Span: *Given a polygonal chain and a number k , is there a simple (non-crossing) planar configuration whose span is at least k ?*

Simple Minimum Planar Span: *Given a polygonal chain and a number k , is there a simple (non-crossing) planar configuration whose span is at most k ?*




Figure 6.9: *Stereogram*: A chain whose maximum span occurs at a nonplanar configuration.

6.3 Span of Three-Dimensional Configurations

When considering the three-dimensional version of a problem, naturally one first wonders if the solution is identical to the two-dimensional case. In other words, is the maximum span of the polygonal chain in three dimensions the same as when it is confined to the plane?

While a positive answer would directly imply the hardness of the three-dimensional problem, it turns out that the configuration which yields the maximum or minimum span of a chain is not necessarily planar. An example of a five-vertex chain whose maximum span occurs when it is not planar is illustrated in figure 6.9. Since the distances between the first and third vertices and the third and fifth vertices are fixed (by the second and fourth vertex-angles), the chain achieves its maximum span when these three vertices are collinear. The maximum planar span of this chain is shown in figure 6.10.

Let us phrase the Maximum Span problem as follows.

Maximum Span: *Given a polygonal chain, find the maximum span over all three-dimensional (possibly self-crossing) configurations.*

Assume v_1 is at the origin. Given any configuration of the chain, it can always





Figure 6.10: The maximum planar span of the chain in figure 6.9.

be rigidly rotated about the origin such that v_n lies on the x -axis. Thus finding the configuration in which v_n is farthest from v_1 is equivalent to finding the configuration in which v_n has the largest x -coordinate. The x -coordinate of v_n is the sum of the horizontal components of the edges of the chain, so we can construct a program to compute such a maximum.

We can characterize the horizontal component of each edge $\overline{v_i v_{i+1}}$ (the x -coordinate of v_{i+1} minus that of v_i) by the angle made by the ray $\overrightarrow{v_i v_{i+1}}$ and the positive x -axis, which we denote by ϕ_i . Thus the horizontal component of edge $\overline{v_i v_{i+1}}$ is $\ell_i \cos \phi_i$. Since the lengths of the edges are fixed, the only variables are the ϕ_i 's.

To maximize the x -coordinate of v_n , we maximize the sum of all components $\ell_i \cos \phi_i$ while respecting the constraints of the chain. The first link of the chain may assume any orientation in space, but due to the fixed turning angles each successive ϕ_i is restricted by the previous ϕ_{i-1} and the turning angle θ_i . Figure 6.11 explains the relationship between these values, which leads directly to the program of figure 6.12.

There is no known general algorithm for maximizing over a nonlinear objective function over linear constraints, especially when the function is not concave separable (able to be expressed as the sum of concave functions). However, in the special case where each $\theta_i \leq \pi/2$, the objective function is maximized when all $\phi_i \geq \pi/2$. (Given any sequence of ϕ_i 's, one could replace all ϕ_i which are less than $\pi/2$ with $\phi_i = \pi/2$ and respect the constraints.) For $\phi_i \geq \pi/2$, the objective function is concave separable.

Figure 6.11: Relationship between ϕ_i , ϕ_{i-1} , and θ_i .

MAXIMIZE:

$$\ell_1 \cos \phi_1 + \ell_2 \cos \phi_2 + \cdots + \ell_{n-1} \cos \phi_{n-1}$$

SUBJECT TO:

$$0 \leq \phi_1 \leq \pi$$

$$\forall i : \phi_i \geq \phi_{i-1} - \theta_i$$

$$\forall i : \phi_i \geq \theta_i - \phi_{i-1}$$

$$\forall i : \phi_i \leq \phi_{i-1} + \theta_i$$

$$\forall i : \phi_1 \leq 2\pi - \phi_{i-1} - \theta_i$$

Figure 6.12: Program to compute the maximum span.

Hochbaum and Shanthikumar [76] have devised an algorithm which can approximate the maximum of such a function, within an error of $1/\varepsilon$, in time polynomial in the input and $\log \varepsilon$. We will later describe an algorithm to approximate the maximum span of all chains, without restriction on the turning angles, in section 6.4.


We will not focus on the problem of computing the minimum span of a chain, except to note that this problem is at least as hard the finding the maximum span.

Minimum Span: *Given a polygonal chain and a number k , is there a (possibly self-crossing) three-dimensional configuration whose span is at most k ?*

Theorem 21. *Minimum Span is at least as hard as Maximum Span.*

Proof. Suppose we have a chain $C = (v_1, v_2, \dots, v_n)$, with edge lengths $\ell_1, \dots, \ell_{n-1}$, for which we wish to find the maximum span. We create a three-edge chain with the vertices $x = (0, 0, 0), y = (1, 0, 0), z = (1, 0, \sum \ell_i), v_1 = (0, 0, \sum \ell_i)$. Now we attach the chain C , such that the vertices v_1 match at $(0, 0, \sum \ell_i)$ and that the vertex-angle at v_1 is $\pi/2$. Let us call this new chain of $n+3$ vertices as chain D . We desire now to compute the minimum span of D , which corresponds to the configuration in which v_n and x are closest together.

It is always possible to perform an edge spin on v_1v_2 such that v_1 and v_n lie in a plane perpendicular to zv_1 , a further edge spin on zv_1 such that x, y, z, v_1 , and v_n are coplanar, and thus x, v_1 , and v_n are collinear. In a sense the chain D mimics the hanging of C from a gallows as shown in figure 6.13. In this position, x and v_n are as close together as possible; therefore, for any fixed configuration of C , the smallest span possible for D is $xv_n = xv_1 - v_1v_n$. Thus the minimum span of D is equal to $\sum \ell_i$ minus the maximum span of C . ■

Figure 6.13: The chain D of the proof of theorem 21.

6.4 Approximating the Maximum Span

Since finding the maximum span is a difficult problem, we present an algorithm to compute an approximation. We can estimate the maximum span by allowing the turning angles to assume only values of the discrete space $\{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$ for some small γ which evenly divides π . Since the possible choices of ϕ_i are dependent solely on constants (θ_i) and the previous angle ϕ_{i-1} , finding the maximum span under this quantized configuration space can be solved by dynamic programming.

We create a series of arrays $x_i[\phi]$ for the discrete value of $\phi \in \{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$, such that $x_i[\phi]$ is the maximum possible x -coordinate of v_{i+1} when $\phi_i = \phi$. Step 3 computes the maximum possible x -coordinate of v_2 given ϕ_1 , and step 4a computes each array $x_i[\phi]$ given $x_{i-1}[\phi]$. At the completion of the algorithm, the largest value of $x_{n-1}[\phi]$ is the largest possible x -coordinate for v_n .

Maximum Span Approximation

1. Compute the cosine for $\phi \in \{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$, such that each value is less than the true $\cos \phi$ by at most γ . Place the values in an array $\widehat{\cos}[\phi]$.
2. Initialize a two-dimensional array $x_i[\phi]$ of reals, indexed for $i \in \{1, \dots, n-1\}$

and $\phi \in \{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$.

3. For each $\phi_1 \in \{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$, let $x_1[\phi_1] \leftarrow \ell_1 \widehat{\cos}[\phi_1]$.

4. For each $i \in \{2, \dots, n-1\}$,

(a) For each $\phi_i \in \{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$, let $x_i[\phi_i]$ be $\ell_i \widehat{\cos}[\phi_i]$ plus the maximum value of $x_{i-1}[\phi_{i-1}]$ over all ϕ_{i-1} which satisfy:

- i. $\forall i : \phi_i \geq \phi_{i-1} - \theta_i$
- ii. $\forall i : \phi_i \geq \theta_i - \phi_{i-1}$
- iii. $\forall i : \phi_i \leq \phi_{i-1} + \theta_i$
- iv. $\forall i : \phi_i \leq 2\pi - \phi_{i-1} - \theta_i$

5. Return the largest value of $x_{n-1}[\cdot]$ as the maximum span.

We now prove that the algorithm correctly computes an approximation of the maximum span of the chain.

Lemma 22. *For any fixed $\phi_{i-1} \in \{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$ and any $\theta_i \leq \pi$, there exists a feasible choice of $\phi_i \in \{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$.*

Proof. We show that for any θ_i the constraints of step 4a can be respected.

If $\theta_i \leq \phi_{i-1} \leq \pi - \theta_i$, then $\phi_i \leftarrow \phi_{i-1}$ is a valid assignment.

If $\phi_{i-1} \leq \theta_i$ and $\phi_{i-1} + \theta_i \leq \pi$, then ϕ_i is determined by constraints 4(a)ii and 4(a)iii, which state that $\theta_i - \phi_{i-1} \leq \phi_i \leq \phi_{i-1} + \theta_i$. Since $\phi_{i-1} \geq \frac{\gamma}{2}$, the interval $\theta_i - \frac{\gamma}{2} \leq \phi_i \leq \theta_i + \frac{\gamma}{2}$ also respects the constraints. This interval must contain one of $\{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$.

If $\phi_{i-1} \leq \theta_i$ and $\phi_{i-1} + \theta_i \geq \pi$, then ϕ_i is determined by constraints 4(a)ii and 4(a)iv, which state that $\theta_i - \phi_{i-1} \leq \phi_i \leq 2\pi - \phi_{i-1} - \theta_i$. Since $\theta_i < \pi$, it follows that $\phi_i \leftarrow (\pi - \phi_{i-1})$ is a valid assignment.

If $\phi_{i-1} \geq \theta_i$ and $\phi_{i-1} + \theta_i \geq \pi$, then ϕ_i is determined by constraints 4(a)i and 4(a)iv, which state that $\phi_{i-1} - \theta_i \leq \phi_i \leq 2\pi - \phi_{i-1} - \theta_i$. Since $\phi_{i-1} \leq \pi - \frac{\gamma}{2}$, the interval $\pi - \frac{\gamma}{2} - \theta_i \leq \phi_i \leq \pi + \frac{\gamma}{2} - \theta_i$ also respects the constraints. This interval must contain one of $\{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$. ■

Lemma 23. *For each sequence of arbitrary values $(\phi_1^*, \phi_2^*, \dots, \phi_{n-1}^*)$ which satisfies the constraints of the continuous Maximum Span program, there exists a valid sequence $(\phi_1, \phi_2, \dots, \phi_{n-1})$ such that for all i , $\phi_i \in \{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$ and $(\phi_i^* - i\gamma) < \phi_i < (\phi_i^* + i\gamma)$.*

Proof. We prove the lemma by induction. We first note that there is a valid choice of ϕ_1 within $\gamma/2$ of ϕ_1^* . Now assume the existence of a valid selection $\phi_1, \dots, \phi_i$ which respects the conditions of the lemma. Then ϕ_i is within $i\gamma$ of ϕ_i^* . Consider the constraints of step 4a. The possible values for ϕ_{i+1}^* given ϕ_i^* differ from the values of the those for ϕ_{i+1} given ϕ_i by at most $i\gamma$ by our inductive hypothesis. Thus for any valid choice of ϕ_{i+1}^* , there is a value v which differs by only $i\gamma$ and which respects the constraints of step 4a for ϕ_{i+1} . Of the subset of valid choices $\{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$ for ϕ_{i+1} , which is nonempty by lemma 22, one value lies within γ of v . Therefore there is a valid choice of ϕ_{i+1} which lies within $(i+1)\gamma$ of ϕ_{i+1}^* . ■

Lemma 24. *Let $0 \leq \phi^* \leq \pi$ and $0 \leq \gamma \leq \pi/2$. Then $\min\{\cos \phi : (\phi^* - \gamma) \leq \phi \leq (\phi^* + \gamma)\} \geq \cos \phi^* - \sqrt{2} \sin \gamma$.*

Proof. The function $\cos \phi$ has only one local minimum (at π) from $0 \leq \phi \leq 3\pi/2$, and in this domain $\cos(\phi^* + \gamma) \leq \cos(\phi^* - \gamma)$. Therefore if $\phi^* + \gamma \leq \pi$, $\min\{\cos \phi\} =$

$\cos(\phi^* + \gamma)$; otherwise $\min\{\cos \phi\} = \cos \pi = -1$. We consider the first case.

$$\cos(\phi^* + \gamma) = \cos \phi^* \cos \gamma - \sin \phi^* \sin \gamma$$

if $\cos \phi^* \geq 0$:

$$\cos(\phi^* + \gamma) \geq \cos \phi^* (1 - \sin \gamma) - \sin \phi^* \sin \gamma$$

$$\cos(\phi^* + \gamma) \geq \cos \phi^* - (\cos \phi^* + \sin \phi^*) \sin \gamma$$

$$\cos(\phi^* + \gamma) \geq \cos \phi^* - \sqrt{2} \sin \gamma$$

if $\cos \phi^* < 0$:

$$\cos(\phi^* + \gamma) \geq \cos \phi^* - \sin \phi^* \sin \gamma$$

$$\cos(\phi^* + \gamma) \geq \cos \phi^* - \sin \gamma$$

$$\cos(\phi^* + \gamma) \geq \cos \phi^* - \sqrt{2} \sin \gamma$$

We have proven the case when $\phi^* + \gamma \leq \pi$, where $\min\{\cos \phi\} = \cos(\phi^* + \gamma)$. Now suppose $\phi^* + \gamma > \pi$, and therefore $\min\{\cos \phi\} = \cos \pi$. Then by the first case, $\min\{\cos \phi\} \geq \cos \phi^* - \sqrt{2} \sin(\pi - \phi^*)$. Since $\gamma > \pi - \phi^*$, and therefore $\sin \gamma > \sin(\pi - \phi^*)$, it follows that $\min\{\cos \phi\} \geq \cos \phi^* - \sqrt{2} \sin \gamma$. ■

Theorem 22. *Maximum Span Approximation computes the maximum span of a chain, minus an error of at most $\sqrt{2}n\gamma \sum \ell_i$.*

Proof. Let the configuration M with the maximum span have the sequence $(\phi_1^*, \phi_2^*, \dots, \phi_{n-1}^*)$. By lemma 23 the dynamic program will find some approximation at least as large as a configuration A with the sequence $(\phi_1, \phi_2, \dots, \phi_{n-1})$, where $(\phi_i^* - i\gamma) < \phi_i < (\phi_i^* + i\gamma)$ for all i . The span of A is calculated according to the

following equation.

$$\begin{aligned}
\text{span}(A) &= \sum_{i=1}^{n-1} \ell_i \widehat{\cos}[\phi_i] \\
\text{span}(A) &\geq \sum_{i=1}^{n-1} \ell_i (\cos \phi_i - \gamma) \\
\text{span}(A) &\geq \sum_{i=1}^{n-1} \ell_i (\min\{\cos \phi_i : (\phi_i^* - i\gamma) \leq \phi_i \leq (\phi_i^* + i\gamma)\} - \gamma) \\
\text{span}(A) &\geq \sum_{i=1}^{n-1} \ell_i (\cos \phi_i^* - \sqrt{2} \sin i\gamma - \gamma) \\
\text{span}(A) &\geq \sum_{i=1}^{n-1} \ell_i (\cos \phi_i^* - \sqrt{2} i\gamma - \gamma) \\
\text{span}(A) &\geq \sum_{i=1}^{n-1} \ell_i (\cos \phi_i^* - \sqrt{2}(i+1)\gamma) \\
\text{span}(A) &\geq \left[\sum_{i=1}^{n-1} \ell_i \cos \phi_i^* \right] - \left[\sum_{i=1}^{n-1} \ell_i \sqrt{2}(i+1)\gamma \right] \\
\text{span}(A) &\geq \text{span}(M) - \sqrt{2}n\gamma \sum_{i=1}^{n-1} \ell_i
\end{aligned}$$

Since in step 1 the values $\widehat{\cos}[\phi]$ are computed to be less than the true values $\cos \phi$, we are assured that $\text{span}(M) - \sqrt{2}n\gamma \sum \ell_i \leq \text{span}(A) \leq \text{span}(M)$. $\blacksquare$

Having analyzed the accuracy of the algorithm, we now analyze its complexity. Since the algorithm spends the most time in step 4, we first focus on the computation of each array $x_i[]$ given its predecessor, $x_{i-1}[]$.

Lemma 25. *Step 4a, the computation of the values of $x_i[]$, can be performed in $O(1/\gamma)$ time.*

Proof. Given the array $x_{i-1}[]$ and a value θ_i , we wish to compute for each $x_i[\phi_i]$, the maximum value of $x_{i-1}[\phi_{i-1}]$ under the constraints of step 4a. (We also wish to add $\ell_i \cos \phi_i$ to each, but this can be done afterward.)

Instead of considering which ϕ_{i-1} are valid for each ϕ_i , let us consider the valid values of ϕ_i for each ϕ_{i-1} according to the constraints.

1. If $\theta_i \leq \pi/2$,

(a) If $\phi_{i-1} \leq \theta_i$ then $\theta_i - \phi_{i-1} \leq \phi_i \leq \phi_{i-1} + \theta_i$.

(b) If $\phi_{i-1} \geq \pi - \theta_i$ then $\phi_{i-1} - \theta_i \leq \phi_i \leq 2\pi - \phi_{i-1} - \theta_i$.

(c) If $\theta_i \leq \phi_{i-1} \leq \pi - \theta_i$ then $\phi_{i-1} - \theta_i \leq \phi_i \leq \phi_{i-1} + \theta_i$.

2. If $\theta_i \geq \pi/2$,

(a) If $\phi_{i-1} \leq \pi - \theta_i$ then $\theta_i - \phi_{i-1} \leq \phi_i \leq \phi_{i-1} + \theta_i$.

(b) If $\phi_{i-1} \geq \pi - \theta_i$ then $\phi_{i-1} - \theta_i \leq \phi_i \leq 2\pi - \phi_{i-1} - \theta_i$.

(c) If $\pi - \theta_i \leq \phi_{i-1} \leq \theta_i$ then $\theta_i - \phi_{i-1} \leq \phi_i \leq 2\pi - \phi_{i-1} - \theta_i$.

We consider the computation of $x_i[\cdot]$ when case 1 holds; the argument is similar should case 2 be true. The maximum possible value of each $x_i[\phi_i]$ is determined by the maximum value over all $x_{i-1}[\phi_{i-1}]$ for which ϕ_i is in the intervals of subcases 1a, 1b, and 1c. Thus we can word the problem of finding the maximum possible value of each $x_i[\phi_i]$ as follows.

Let Λ_a be a set of weighted intervals where for each $\phi_{i-1} \leq \theta_i$, there is an interval $\lambda = [(\theta_i - \phi_{i-1}), (\phi_{i-1} + \theta_i)]$ with weight $w_\lambda = x_{i-1}[\phi_{i-1}]$.

Let Λ_b be a set of weighted intervals where for each $\phi_{i-1} \geq \pi - \theta_i$, there is an interval $\lambda = [(\phi_{i-1} - \theta_i), (2\pi - \phi_{i-1} - \theta_i)]$ with weight $w_\lambda = x_{i-1}[\phi_{i-1}]$.

Let Λ_c be a set of weighted intervals where for each $\theta_i \leq \phi_{i-1} \leq \pi - \theta_i$, there is an interval $\lambda = [(\phi_{i-1} - \theta_i), (\phi_{i-1} + \theta_i)]$ with weight $w_\lambda = x_{i-1}[\phi_{i-1}]$.

Then $x_i[\phi_i]$ is determined by $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_a \cup \Lambda_b \cup \Lambda_c\}$.

Computing $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_a\}$

We first solve the problem of computing $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_a\}$. The set Λ_a is composed of a series of intervals $\lambda = [(\theta_i - \phi_{i-1}), (\phi_{i-1} + \theta_i)]$ with $w_\lambda = x_{i-1}[\phi_{i-1}]$. Thus for any $\phi_i < \theta_i$, $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_a\}$ is equivalent to $\max\{x_{i-1}[\phi_{i-1}] : \phi_i \leq \theta_i - \phi_{i-1}\}$. This leads to the following method which runs in $O(\theta_i/\gamma)$ time and space.

1. Let θ^- be the largest value of $\{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots\}$ less than or equal to θ_i .
2. Initialize an array $m_1[\phi_i]$ to store $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_a\}$ for each value of $\phi_i \in \{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, 2\theta^- + \frac{\gamma}{2}\}$.
3. $\text{MAX} \leftarrow -\infty$
4. For each $\phi_{i-1} \in \{\theta^-, \theta^- - \gamma, \theta^- - 2\gamma, \dots, \frac{\gamma}{2}\}$,
 - (a) If $\text{MAX} < x_{i-1}[\phi_{i-1}]$, then $\text{MAX} \leftarrow x_{i-1}[\phi_{i-1}]$.
 - (b) $\phi_i \leftarrow \theta^- - \phi_{i-1} + \frac{\gamma}{2}$
 - (c) $m_1[\phi_i] \leftarrow \text{MAX}$
5. Since the intervals of Λ_a are symmetric about θ_i , reflect the array to compute $m_1[\phi_i]$ for $\phi_i > \theta_i$.
 - (a) Let θ^+ be the smallest value of $\{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots\}$ greater than or equal to θ_i .
 - (b) For each $\phi_i \in \{\theta^-, \theta^- - \gamma, \theta^- - 2\gamma, \dots, \frac{\gamma}{2}\}$, let $m_1[\theta^+ + \theta^- - \phi_i] \leftarrow m_1[\phi_i]$.

Computing $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_b\}$

The problem of computing $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_b\}$ is slightly more complicated. Let θ^+ be the smallest value of $\{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots\}$ greater than or equal to θ_i , and let θ^- be the largest value of $\{\gamma, 2\gamma, \dots\}$ less than or equal to θ_i . Then Λ_b is the set of intervals $\{[(\theta^+ - \theta^-), (\theta^+ + \theta^-)], [(\theta^+ - \theta^- + \gamma), (\theta^+ + \theta^- + \gamma)], \dots, [(\pi - \theta^+ - \theta^-), (\pi - \theta^+ + \theta^-)]\}$.

We compute $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_b\}$ by first removing all redundant intervals. We say an interval λ of weight w_λ is *redundant* if there exist two intervals λ_L (lying left of λ) and λ_R (lying right of λ) whose weights are at least as large than w_λ and which cover λ . We say that two intervals λ_L (lying left of λ) and λ_R *cover* λ if all values of $\{\frac{\gamma}{2}, \frac{3\gamma}{2}, \pi - \frac{\gamma}{2}\}$ in λ are also in $\lambda_L \cup \lambda_R$. This is similar to saying that $\lambda \subset \lambda_L \cup \lambda_R$ but refers only the possible values of ϕ in its discrete domain. An interval can clearly not determine $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_b\}$ for any value of ϕ_i if it is redundant. Likewise, it must determine at least one such value if it is not redundant.

The notion of redundancy is ill-defined if we have a long series of intervals with identical weights, since it becomes unclear which of the intervals is redundant. To resolve the ambiguity we force the weight of λ_R to be strictly larger than that of λ , and thus an interval is redundant if $w_{\lambda_L} \geq w_\lambda$ and $w_{\lambda_R} > w_\lambda$.

Because all intervals are the same size, our set Λ_b has the property that if some interval λ is redundant, then there must be a choice of λ_L and λ_R which are *nonredundant*. Suppose that λ_L and λ_R , whose weights are larger than w_λ , cover λ , but that λ_L is itself also redundant. Then there are two intervals which cover λ_L , an interval λ_L^- to the left of λ_L , and an interval λ_L^+ to the right, both of whose weights are greater than w_{λ_L} (and thus greater than that of λ). If λ_L^+ is between λ_L and λ , then λ_L^+ and λ_R cover λ . If λ_L^+ is to the right of λ , then λ_L^- and λ_L^+ cover all intervals in between, including λ . Thus if we remove redundant intervals one by one, there is no possibility that a redundant interval will become nonredundant.

This leads to the following algorithm to remove redundant intervals of Λ_b .

1. Let θ^+ be the smallest value of $\{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots\}$ greater than or equal to θ_i .
2. Let θ^- be the largest value of $\{\gamma, 2\gamma, \dots\}$ less than or equal to θ_i .
3. Initialize a stack of intervals.
4. For each $i \in \{0, 1, \dots, \frac{\pi - 2\theta^+}{\gamma}\}$,

- (a) $\lambda \leftarrow [(\theta^+ - \theta^- + i\gamma), (\theta^+ + \theta^- + i\gamma)]$
- (b) $w_\lambda \leftarrow x_{i-1}[\theta^+ + i\gamma]$
- (c) If S has at least two elements, then
 - i. $t \leftarrow$ top element of the stack
 - ii. $u \leftarrow$ element directly beneath t
 - iii. While S has at least two elements, $w_u \geq w_t$, $w_\lambda > w_t$, and t is covered by λ and u ,
 - A. mark t as redundant and pop it off the stack.
 - B. $t \leftarrow u$ (the new top element of the stack)
 - C. $u \leftarrow$ element directly beneath t if it exists
- (d) Push λ on the stack.

It is clear by step 4(c)iii that only redundant intervals are so marked, so to prove correctness we must demonstrate that all redundant intervals are marked as such. Suppose the algorithm leaves a contiguous section of redundant intervals $\lambda_1, \lambda_2, \dots, \lambda_k$ in between two nonredundant intervals of greater weight, λ_L (to the left) and λ_R (to the right). Since $\lambda_1, \lambda_2, \dots, \lambda_k$ are covered by λ_L and λ_R , it follows that the union of any two of $\lambda_1, \lambda_2, \dots, \lambda_k$ cover the intervals in between. Thus $w_{\lambda_1} \geq w_{\lambda_2} \geq \dots \geq w_{\lambda_{k-1}} \geq w_{\lambda_k}$, or else one of these intervals would have been popped from the stack. When λ_R is considered in step 4(c)iii, $w_{\lambda_R} > w_{\lambda_k}$, and $w_{\lambda_{k-1}} \geq w_{\lambda_k}$. If λ_k is covered by λ_L and λ_R , then certainly λ_k is covered by $\lambda_{k-1} \cup \lambda_R$. The redundant λ_k is popped from the stack, a contradiction.

Once all redundant intervals are removed, we begin popping intervals off the stack and computing an array $m_2[\phi_i] = \max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_b\}$ from $\phi_i = \pi - \frac{\gamma}{2}$ down to $\phi_i = \frac{\gamma}{2}$. Each interval determines at least one value of $m_2[\cdot]$. Therefore we only need to examine the top intervals on the stack, since if an interval were skipped it would be redundant.

This leads to the following algorithm to compute $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \text{nonredundant intervals of } \Lambda_b\}$.

1. Initialize an array $m_2[\phi_i]$ to store $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_b\}$ for each value of $\phi_i \in \{\frac{\gamma}{2}, \frac{3\gamma}{2}, \dots, \pi - \frac{\gamma}{2}\}$.
2. $t \leftarrow$ top element of the stack (preserved from the previous algorithm)
3. $u \leftarrow$ element directly beneath t (if it exists)
4. For each $\phi_i \in \{\pi - \frac{\gamma}{2}, \pi - \frac{3\gamma}{2}, \dots, \frac{\gamma}{2}\}$,
 - (a) If $(\phi_i \notin t)$ or $(\phi_i \in t \cap u \text{ and } w_u > w_t)$ then
 - i. Pop t off the stack.
 - ii. $t \leftarrow u$ (the new top element of the stack)
 - iii. $u \leftarrow$ element directly beneath t (if it exists)
 - (b) $m_2[\phi_i] \leftarrow w_t$

Both algorithms above run in time $O(|\Lambda_b|)$, or $O(\pi - 2\theta_i/\gamma)$.

Computing $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_c\}$

The problem of computing $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_c\}$ is symmetric to the that of Λ_a and can be solved analogously in time $O(\theta_i/\gamma)$.

Merging the three subproblems to compute $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda_a \cup \Lambda_b \cup \Lambda_c\}$

The three arrays $m_1[], m_2[], m_3[]$ can thus be computed in total time $O(1/\gamma)$ and then merged to find $\max\{w_\lambda : \phi_i \in \lambda; \lambda \in \Lambda\}$. Adding $\ell_i \cos \phi_i$ to each value yields the array $x_i[]$. ■

With an efficient method of performing step 4a, we are ready to prove the final theorem of this chapter.

Theorem 23. *The maximum span of a chain of n vertices can be computed within an error of $\varepsilon \sum \ell_i$, for any $\varepsilon > 0$, in $O(n^2/\varepsilon)$ time and $O(n/\varepsilon)$ space.*

Proof. Let γ be ε/n .

The array $\widehat{\cos}[]$ in step 1 has $O(1/\gamma) = O(n/\varepsilon)$ entries, and therefore uses a like amount of space. Computing the cosine of each value within an error of at most γ can be performed in time $O(\log(1/\gamma)) = O(\log(n/\varepsilon))$ by lemma 2. Thus the entire array can be built in $O((n/\varepsilon) \log(n/\varepsilon))$ time and $O(n/\varepsilon)$ space.

Each iteration of step 4a can be performed in time $O(1/\gamma)$ by lemma 25. Since the algorithm is dominated by n iterations, the algorithm finishes in time $O(n/\gamma)$, or $O(n^2/\varepsilon)$. By theorem 22, the algorithm computes the maximum span within an error of at most $\sqrt{2}n\gamma \sum \ell_i$, or $\sqrt{2}\varepsilon \sum \ell_i$.

The Maximum Span Approximation algorithm was described using $n - 1$ arrays $x_i[]$ of size (π/γ) each for the sake of simplicity. However, at any iteration of step 4a, only the arrays $x_{i-1}[]$ and $x_i[]$ are necessary; the arrays $x_1[], x_2[], \dots, x_{i-2}[]$ are never used again and can be discarded. Thus the algorithm requires only two arrays of size (π/γ) each, for a total space constraint of $O(1/\gamma) = O(n/\varepsilon)$. ■

6.5 Flattening

The overall goal of most computational studies of linkages is to determine whether one can reconfigure a linkage from a given configuration to a certain target configuration. As discussed in section 1.1, the most often used approach is to ask if both can be reconfigured into some canonical configuration. If so, then because the motions are reversible, one can reconfigure the chain from the given configuration to the canonical configuration, and then from the canonical configuration to the target. When the angles between edges are free to change, the canonical configuration for chains is a line segment. Obviously this is not achievable for chains with fixed vertex-angles

between links. Instead of reconfiguring the chain into a one-dimensional segment, we ask if one can reconfigure such a chain to assume a two-dimensional configuration, that is, to lie in the plane without self-intersections.

Unlike the case of reconfiguring a linkage from a given configuration to a target via a straight segment, moving two configurations of a chain into the plane does not guarantee that one can reconfigure between the two. Rather, one must also determine whether a chain can be reconfigured between any two of its planar configurations. This remains an open problem that we will discuss in chapter 8.

We say a chain with fixed vertex-angles can be *flattened* if it can be moved into a planar configuration. There exist, of course, chains which cannot be flattened, such as the knitting needles chain of figure 1.12 (on page 10). This leads us to the following problem.

Flattening: *Given a polygonal chain with fixed vertex-angles in three dimensions, is there a sequence of motions which place the chain into a non-crossing planar configuration?*

While an algorithm for Flattening would be an invaluable tool for computing a motion between two configurations, we demonstrate it to be an NP-hard problem using a method similar to the NP-hardness proof for Minimum Planar Span.

Theorem 24. *Flattening is NP-hard.*

Proof. Given a set $S = \{s_1, s_2, \dots, s_n\}$ of positive integers, we show that in polynomial time we can create a chain which can be reconfigured into the plane if and only if the set has a partition.

We begin by creating, in the plane, either of the subchains in figure 6.14, where σ is the sum of the elements of S . An enumeration will show that these are the only planar configurations of this chain.

We continue the chain from vertex a by building a subchain similar to the one




Figure 6.14: The two possible planar configurations of the starting subchain constructed in the proof of theorem 24.




Figure 6.15: The subchain starting at vertex a .

in the proof of theorem 20, but in a vertical plane (orthogonal to the subchain of figure 6.14). We start at vertex a , and for every element $s_i \in S$, place an edge moving upward of length $1/n$, followed by a horizontal edge of length s_i , and repeat. We finish with an edge of length 5 extending upward. The subchain from a is illustrated in figure 6.15; the entire chain is illustrated in figure 6.16.

Now we flatten the chain into the plane. If the configuration is to be non-crossing, the only place for the edge of length 5 is inside the shaded triangle of figure 6.14, resulting in the configuration of figure 6.17. Therefore this edge must have the same




Figure 6.16: The entire chain constructed in the proof of theorem 24.




Figure 6.17: An instance that can be flattened.

x -coordinate as vertex a . As in the proof of theorem 20, this implies a partition of the edges which correspond to the elements of the set S . ■

Because Flattening is NP-hard, one might wonder if would be easier to answer more restrictive versions of the problem. Instead of seeking an arbitrary planar configuration, we may also wish to consider whether a chain has a monotone planar configuration. Recall the definition of monotonicity from chapter 3 on page 41. A chain is *monotone* with respect to a line ℓ if the intersection of the chain with every line perpendicular to ℓ is connected. When ℓ is the x -axis, then the intersection of the chain with any vertical line is connected, and we call the chain x -monotone. For our purposes, we consider all monotone chains to be x -monotone, since one can always rotate it until it is monotone with respect to the x -axis.

In other words, as one traverses an x -monotone chain edge by edge, the chain always progresses to the right. Examples of a monotone chain and a non-monotone chain are shown in figure 6.18.

Monotone Flattening: *Given a polygonal chain with fixed vertex-angles in three dimensions, is there a sequence of motions which place the chain into a non-crossing planar monotone configuration?*



Figure 6.18: *Left:* An x -monotone chain. *Right:* A chain that is not x -monotone.




Figure 6.19: A chain with a monotone planar configuration.

Despite its apparent simplicity, this problem is also NP-hard. We proved Flattening to be NP-hard via the same reduction as used for Minimum Planar Span; we now prove Monotone Flattening to be NP-hard via the same reduction as used for Maximum Planar Span.

Theorem 25. *Monotone Flattening is NP-hard.*

Proof. Suppose we have a Partition problem on some set $S = \{s_1, s_2, \dots, s_n\}$. We create a chain similar to the one used in the proof of theorem 18, where the turning angles θ_{i+1} are equal to s_i/σ , where σ is the sum of the elements of S .

Instead of extending the extreme links to some large length, we add three links to each to obtain the chain of figure 6.19. The first is $1/2\sigma$ radians within being perpendicular to the end link (in blue), followed by two links $1/2\sigma$ within being parallel to the first (plus or minus trigonometric errors of less than $1/12\sigma$).

The sets of three links at the end of the chain can lie monotonic with respect to the horizontal if and only if the blue links lie within $1/4\sigma$ of being parallel, which

indicates a partition of the angles. (In the figure, $\theta_4 = \theta_2 + \theta_3 + \theta_5$.) ■

Chapter 7

Chirality

The most general problem in reconfiguring chains is to determine whether it is possible to move from a starting configuration to an arbitrary target. In this chapter we prove the difficulty of this problem by demonstrating that a more specific question is intractable.

In section 1.3.3 we introduced the notion of chirality and discussed its importance in biology and pharmacology. An object is *chiral* if it cannot be superimposed on its mirror image. For rigid objects, this determination is quite simple; one needs only to compute the mirror image and check whether it is identical to the original object. A number of algorithms for finding symmetries already exist; a survey paper by Eades [46] provides a thorough coverage. However, from a chemical point of view, it is far more important to discover whether an object can be reconfigured into its mirror image. This is the problem we seek to answer in this chapter.

Recall that an object is chiral if it is *not* possible to reconfigure the chain to its mirror image, and is *achiral* if it is possible. To avoid the confusing situation of determining chirality and answering *yes* when reconfiguration is *not* possible, we phrase the question as the *Achirality* problem.

Achirality: *Given a polygonal chain in three dimensions, can it be reconfigured to*

its mirror image while maintaining its vertex-angles?

In the first chapter, we also presented an example of the achiral molecule synthesized by Mislow and Bolstad [109], shown in figure 1.30 on page 26. The rigid representation of their molecule is chiral but it can easily be reconfigured into its mirror image. The central bond is rigid, but if middle portion is rotated about the red bonds by 90 degrees, the mirror image is attained. It is rather interesting that despite this reconfiguration, the molecule cannot be brought to a single configuration which, when regarded as a rigid object, is achiral. Mislow and Bolstad called this type of molecule a *rubber glove*. One can turn a left-handed glove inside out to make a right-handed glove, but at no time is the glove in a symmetric configuration.

Like many of the problems we have encountered in the past two chapters, determining whether a chain is a rubber glove is intractable. We call a chain of n vertices a *palindrome* if its i^{th} edge length and bond angle is the same as the $n - i + 1^{\text{st}}$. A chain which is not a palindrome can be in a symmetric configuration only if it lies in the plane. This leads us to the following lemma.

Lemma 26. *A chain which is not a palindrome can be a rubber glove only if it cannot be flattened.*

The chain in figure 7.1, the one used in the reduction for proving Flattening to be NP-hard, is achiral. Since determining whether this chain can be flattened is NP-hard, determining whether the chain *cannot* be flattened is co-NP-hard¹. Thus we have proven theorem 26.

¹A problem is co-NP-hard if its negation is NP-hard. *Unsatisfiability*, *Graph-Uncolorability*, and determining if there does *not* exist a Hamiltonian Path in a given graph are all examples of problems which are in co-NP. NP is the class of problems for which the certificate of a *yes* instance (a solution) can be verified in polynomial time; co-NP is the class of problems whose certificate of a *no* instance (a counter-example) can be verified in polynomial time. The textbooks by Garey and Johnson [62], Hopcroft and Ullman [79], and Lewis and Papadimitriou [97] provide a thorough coverage.




Figure 7.1: A rubber glove.

Theorem 26. *Determining whether a chain is a rubber glove is co-NP-hard.*

The remainder of this chapter is dedicated to proving that Achirality is NP-hard by a reduction from Partition. (This in turn would imply that Chirality is co-NP-hard.) Because the chain built in the reduction is highly complex, we first describe the structure as a general gadget using planes, polyhedra, and vertices of high degree. We then show how to approximate the gadget using a single chain.

The gadget is illustrated in figures 7.2 and 7.3. In the center of the object are two blue planes which are very close together, which we will call a *double spatula*. The double spatula is connected to two auxiliary chains. The right auxiliary chain is connected to the top small blue triangle, and the left auxiliary chain is connected to the bottom blue triangle. Therefore, to bring the object to its mirror image, the two auxiliary chains must switch places. The small blue triangles prevent the chains from switching places by rotating around the outside of the double spatula, so if the chains are to switch places, they must pass between the two large blue triangles. Since the double spatula is very narrow, the auxiliary chains must be in a configuration close




Figure 7.2: The gadget to decide Achirality.

to planar, that is, constrained by two close planes, before passing through.

The caveat is that the cap is a tetrahedron and cannot be flattened. If an auxiliary chain passes through the double spatula, its cap must stretch around the endpoints marked in red. The cap can extend past the endpoints if and only if there exists a configuration of the chain, while it is constrained by two close planes of the double spatula, whose span is long enough. If we make the auxiliary chain the same chain that we built to prove that Maximum Span Between Two Planes is NP-hard, then determining if the two auxiliary chains can switch sides is NP-hard.

The remainder of this chapter explains how we can build the object of figure 7.2 using a single chain. We first define a characteristic of a chain, *quasi-rigidity*, that describes the amount of freedom available in a chain's configuration. We then describe the numerous gadgets that will be used in the reduction before completing the proof in section 7.4.




Figure 7.3: *Stereogram*: Side view of the gadget to decide Achirality. (Not to scale.)




Figure 7.4: A staple and hook.




Figure 7.5: A linked staple and hook.

7.1 Quasi-rigidity

Consider the two chains in figure 7.4. We refer to the red chain as a *staple*, and the green as a *hook*. This pair of chains has the property that if they are positioned as in figure 7.5, they cannot be separated. We refer to this special configuration of the two chains as a *linked staple and hook*.

Lemma 27. *The linked staple and hook cannot be separated.*

Proof. Consider the labeling of the vertices in figure 7.5. We assume without loss of generality that the edges $\overline{ab}$ and $\overline{bc}$ are fixed in space, since any motion of these two edges can be realized by moving the rest of the chains.

We first discuss the properties of the hook. The vertex d of the hook is as close as possible to the edge $\overline{ab}$ when the hook is planar, so moving d only serves to increase the distance between d and $\overline{ab}$. Due to the sharp angle at d , both c and e are considerably

farther from $\overline{ab}$ than is d .

The staple is wrapped around $\overline{ab}$ at x , and around d at y . Because the distance between x and y is fixed, the staple places an upper bound on the distance allowed between d and $\overline{ab}$. Because c and e are farther from $\overline{ab}$ than is d , one cannot move y down away from d without moving y farther from $\overline{ab}$ in the downward direction (as defined by the figure). This implies that x must move downward, but x is obstructed by $\overline{ab}$. Similarly one cannot move x away or along $\overline{ab}$ without moving y upward through d .

Since y cannot be moved away from d except for an arbitrarily small amount, d cannot be moved by any more than an arbitrarily small amount. The hook cannot be reconfigured until the staple is removed, and the staple cannot be removed until the hook is reconfigured. Thus neither is possible. ■

The linked staple and hook can be used to make a variety of locked geometric shapes. For example, consider the triangle-like chain in figure 7.6, which we call a *staple-and-hook triangle*. With a linked staple and hook, the four vertices $abcd$ behave like a rigid triangular cycle.

We wish to claim that although there is some freedom in the structure, it is more or less locked. To formalize this notion, we define a distance metric between configurations of a chain so that it is possible to discuss how close together or far apart they are. This will allow us, for example, to claim that the chain in figure 7.6 is locked because it can only be moved to nearby configurations.

Let χ and χ' be two configurations of some chain C . We define the *discrepancy in position of any point p* in χ to be the distance between p and its corresponding point p' in χ' . We define the *discrepancy in position of the two configurations* to be the largest distance between the position of any point in χ and the position of its corresponding point in χ' . This scheme is illustrated in figure 7.7, where the dotted line indicates the




Figure 7.6: A staple-and-hook triangle.

discrepancy. A useful property of this definition is that the largest distance between any two points of a configuration always occurs between corresponding vertices.

Lemma 28. *The largest discrepancy between two configurations is determined by the positions of the vertices of the chain.*

Proof. Let corresponding points p and p' have the largest discrepancy in position, and suppose they each lie interior to an edge. Let o, o' and q, q' be two pairs of corresponding points on the same edges, such that p lies between o and q (and thus p' lies between o' and q'). If $|oo'| = |pp'|$, then the endpoints of the edges have the same discrepancy in position as p and p' . Because edges are linear segments, if $|oo'| < |pp'|$ then $|pp'| < |qq'|$, a contradiction. ■

To obtain a distance metric which is invariant to affine motions of a configuration, we define the *distance between two configurations* of the chain to be minimum discrepancy in position modulo rigid rotations and translations. Given two configurations, we rotate and translate one of them until the largest discrepancy in the




Figure 7.7: The discrepancy in position of the two configurations is the largest discrepancy in position of any point, the distance between d and d' .

position of any point is minimized. This discrepancy is the distance between the two configurations, as demonstrated in figure 7.8.

This metric gives us the power to describe formally the amount of freedom in a given configuration. We will prove that the chain in figure 7.6 is locked because it cannot be moved into any distant configuration. We can further quantify this statement by defining a parameter of any configuration which we shall call the *allowed perturbation*. For any configuration χ , its allowed perturbation is the largest distance from χ to any other reachable configuration.

All the complex structures in this chapter exploit linked staples and hooks. We define a structure as *quasi-rigid* if the allowed perturbation of the chain approaches zero as the length of the staple-and-hook edges also approaches zero (while the remaining lengths of the chain remain fixed). In some structures, we add the constraint that the vertex-angles of the staples approach π .

For example, we would say that the staple-and-hook triangle of figure 7.6 is quasi-rigid since as the staple and hook are made smaller and smaller, the chain becomes more and more locked as in figure 7.9. We prove this formally in the following lemma.



Figure 7.8: The distance between χ and χ' is the largest discrepancy in position of any point on the chain, minimized over all global rotations and translations. From its position in figure 7.7, χ' has been rotated through the third dimension (about the edge $\overline{bc}$ in figure 7.7), rotated slightly clockwise, then translated onto χ .




Figure 7.9: A tightly locked staple-and-hook triangle.

Lemma 29. *In a staple-and-hook triangle, the allowed perturbation is at most twice the total length of the staple and hook.*

Proof. Let χ be the configuration of the staple-and-hook triangle as shown in figure 7.6, and let χ' be any configuration reachable from χ . Because the distance metric is invariant of global rotation and translation, we can assume without loss of generality that the vertices a, b , and c are coincident in both χ and χ' . Thus the distance between the two configurations is at the discrepancy in position of the other points on the chain by lemma 28. Except for b and c , all vertices of the chain are part of the staple or the hook.

Let K be the total lengths of the staple-and-hook edges. Since the staple and hook are not separable, no vertex therein can be more than distance K away from a in either χ or χ' . Because a is coincident in χ and χ' , all vertices other than b and c are within a sphere of diameter $2K$ around a . ■

Thus the staple-and-hook triangle is quasi-rigid, since the allowed perturbation is at most twice the total length of the staple-and-hook edges. Another interpretation, more useful for our purposes, is that for any desired allowed perturbation ε , it suffices for the total size of the staple-and-hook edges to be less than $\varepsilon/2$. Since the sufficient total length is polynomial in the desired allowed perturbation, we say the structure is *polynomially quasi-rigid*. In our NP-hardness proof for Achirality, it will be important that all edges are only polynomially small to achieve the desired quasi-rigidity. We cannot afford to use arbitrarily large space to describe the size of the edges if we hope to obtain a polynomial-time reduction.

The bulk of the NP-hardness proof for Achirality consists of using a single chain to build complex structures which mimic the function of the gadget of figure 7.2. We then show that their allowed perturbations are polynomially small with respect to the size of an instance of Partition. In this way we will be able to build the rigid

components of the gadget such as the double spatula but with a single chain.

If we prove that the structures are polynomially quasi-rigid, then we are part of the way toward proving that the allowed perturbation can be made polynomially small with respect to the size of the Partition instance. If a function is polynomial, it implies that its asymptotic behavior is polynomial. In other words, if a structure is polynomially quasi-rigid, then there exists some $\varepsilon > 0$ such that for all total lengths K of the staple-and-hook edges less than ε , the allowed perturbation is polynomial in K . For a polynomial reduction, we will also require that the space required to store the value ε is polynomial in the size of the Partition instance. For our proof, it will suffice to show that ε is polynomial in the length of the longest edge of the structure.

7.2 Fundamental Quasi-rigid Structures

We will build the components of the gadget out of four fundamental structures, *graduated rulers*, *planar nets*, *caps*, and *staple-edges*.

Graduated rulers

In everyday use, a ruler is a straight line segment marked with graduations at regular intervals. The *graduated ruler* we design below is an object with hooks at regular intervals, used for a similar purpose.

Consider the structure of figure 7.10. We begin with the uppermost drawing, a horizontal edge connecting two hooks. Building from the right hook, we draw a long edge back to the left hook and connect it with a staple. In the final drawing, we continue the chain to the right, alternating between short edges and hooks, until we connect the chain to the rightmost hook with a staple.

If the lengths of the short edges (drawn in the bottom image of figure 7.10) add up to the length of the middle edge (drawn in the middle image), then as the staple-




Figure 7.10: Step-by-step construction of a graduated ruler.



Figure 7.11: A graduated ruler with smaller staple-and-hook edges.

and-hook edges get smaller and smaller, the chain becomes more and more rigid as in figure 7.11.

Lemma 30. *The graduated ruler is polynomially quasi-rigid.*

Proof. We consider the labeling indicated in figure 7.12. Assume the edge $\overline{cd}$ is fixed in space. Let κ_i be the total length of the edges of the i^{th} hook (or staple and hook) as indicated. We demonstrate that as the sum $K = \kappa_0 + \kappa_1 + \cdots + \kappa_n$ approaches zero, the allowed perturbation is $O(\sqrt{K})$, and thus that K is quadratic in the desired allowed perturbation.

We assume K to be smaller than $|cd|$. This will allow us to prove that no point on the graduated ruler can be displaced more than a small distance.





Figure 7.12: Labeling of the chain in the proof of lemma 30.

Figure 7.13: Illustration of x and h .

Let x be a point of the chain not on $\overline{cd}$, as shown in figure 7.13. Then x must be on a path which connects the first and last hooks, which are distance $|cd|$ apart. Let $\widetilde{dx}$ be the total arc length of the chain from the left staple and hook to x , and $\widetilde{cx}$ be the total arc length of the chain from the right staple and hook to x . Then x must be within spheres of radius $\widetilde{dx}$ around d and of radius $\widetilde{cx}$ around c .

Since the sum of the lengths of the edges e_i equals $|cd|$, the total arc length $\widetilde{dx} + \widetilde{cx} \leq |cd| + K$. On the line connecting the centers of the spheres, c and d , the overlap of the two spheres can therefore be at most K . We now compute how far x can lie from the edge $\overline{cd}$, indicated in the figure by the distance h and the position of x . If we follow the path from c to x to d , the chain traverses a total horizontal distance $|cd|$, and a total vertical distance $2h$. By the Pythagorean theorem, this path has total length at least $\sqrt{|cd|^2 + (2h)^2}$, which must be less than $|cd| + K$.

$$\sqrt{|cd|^2 + (2h)^2} \leq |cd| + K$$

$$|cd|^2 + 4h^2 \leq |cd|^2 + 2|cd|K + K^2$$

$$4h^2 \leq 2|cd|K + K^2$$




Figure 7.14: Legend for graduated rulers.




Figure 7.15: Building the frame of the planar net.

We assumed that $K \leq |cd|$, so $4h^2 \leq 3|cd|K$ and thus $h \leq \sqrt{3|cd|K}/2$. The intersection of the spheres has width K along the edge $\overline{cd}$, and width $\sqrt{3|cd|K}/2$ perpendicular to $\overline{cd}$. Thus the total possible discrepancy in position for any point x is at most $2K + \sqrt{3|cd|K}$ for $K < |cd|$ as per our original assumption. ■

The graduated ruler will be the building block to create more complex quasi-rigid structures. To simplify future diagrams, we will draw graduated rulers as shown in figure 7.14. Circles will be drawn to signify staples or hooks. When more than one line meets at a circle, it is to signify linked staples and hooks. For example, in figure 7.15, the three graduated rulers are linked at their endpoints.

Planar nets

A *planar net* is a flat grid-like mesh, much like the seat of a wicker chair or a shelf of a refrigerator. We begin by building an isocles right triangle with three graduated rulers linked at their endpoints, with hooks at regular intervals, as in figure 7.15. We refer to this structure as the frame of the planar net.




Figure 7.16: Labeling used in the proof of lemma 31.

Lemma 31. *The frame of the planar net is polynomially quasi-rigid.*

Proof. Let χ be the configuration of the frame as described, and let χ' be a configuration reachable from χ .

Let a, b, c be the three corners of the triangle (endpoints of the graduated rulers) in the configuration χ as illustrated in figure 7.16, where b is at the right angle. (Since the triangle is isocles, this implies that $|ab| = |bc|$.) Let a', b', c' be the corresponding points in χ' . Let Π be the largest allowed perturbation of any of the individual graduated rulers. Since graduated rulers are polynomially rigid by lemma 30, the lemma is proven if we demonstrate the allowed perturbation of the frame to be a multiple of Π .

We need only show that the discrepancy in position of the three endpoints a, b, c is a multiple of Π . Because the remainder of the vertices belong to graduated rulers, their discrepancy could be at most Π more than that of the endpoints. We now show the allowed perturbation of the frame to be at most $(1 + \sqrt{65})\Pi$.

Because the distance metric is invariant to global rotation and translation of χ and χ' , we assume the points a and a' to be coincident in the two configurations. Furthermore, we rotate the two configuration such that a, b, b' are collinear on the y -axis and that a', b', c' lie in the same plane as a, b, c . We assume this plane to be the horizontal xy -plane.

The distance between a and a' is zero, and the distance between b and b' is at

most Π since they both lie on the y -axis. We compute the maximum possible distance between c and c' by considering the triangle $\triangle a'b'c'$. Because the frame is constructed of graduated rulers, we know that

$$|ab| - \Pi \leq |a'b'| \leq |ab| + \Pi, \quad (7.1)$$

$$|ac| - \Pi \leq |a'c'| \leq |ac| + \Pi, \text{ and} \quad (7.2)$$

$$|bc| - \Pi \leq |b'c'| \leq |bc| + \Pi. \quad (7.3)$$

We assume that Π is smaller than any of the graduated rulers, so that the above values are all positive.

Let ϕ be the angle $\angle a'b'c'$. By the law of cosines, we can express ϕ by

$$\cos \phi = \frac{|a'b'|^2 + |b'c'|^2 - |a'c'|^2}{2|a'b'||b'c'|}. \quad (7.4)$$

Substituting the values of inequalities 7.1, 7.2, and 7.3 into equation 7.4 yields the range of possible values for ϕ . We first compute a lower bound, which we will see is negative. Because the denominator of equation 7.4 is a positive product of distances, our lower bound will be achieved by finding the minimum possible value for the numerator, and the minimum possible value for the denominator.

$$\begin{aligned} \cos \phi &\geq \frac{(|ab| - \Pi)^2 + (|bc| - \Pi)^2 - (|ac| + \Pi)^2}{2(|ab| - \Pi)(|bc| - \Pi)} \\ \cos \phi &\geq \frac{|ab|^2 + |bc|^2 - |ac|^2 + 2\Pi(|ac| - |ab| - |bc|) + 3\Pi^2}{2(|ab||bc| - \Pi(|ab| + |bc|) + \Pi^2)} \end{aligned}$$

Because $\triangle abc$ is a right triangle, $|ab|^2 + |bc|^2 - |ac|^2 = 0$. We can further simplify the inequality to

$$\cos \phi \geq \frac{2\Pi(|ac| - |ab| - |bc|) + 3\Pi^2}{2(|ab||bc| - \Pi(|ab| + |bc|) + \Pi^2)}.$$

Since $|ab| = |bc|$, we know that $|ac| - |ab| - |bc| \geq -|ab|$ by the triangle inequality. We further note that Π^2 is necessarily a positive value. Thus we obtain

$$\cos \phi \geq \frac{-2\Pi|ab|}{2|ab|^2 - 4\Pi|ab| + 2\Pi^2}.$$

Since we are interested in the allowed perturbation when Π approaches zero, we may assume that $\Pi \leq |ab|/4$. Then $4\Pi|ab|$ is at most $|ab|^2$, so

$$\cos \phi \geq \frac{-2\Pi|ab|}{2|ab|^2 - |ab|^2 + 2\Pi^2}.$$

Since $2\Pi^2$ is necessarily a positive value, $\cos \phi \geq -2\Pi/|ab|$.

We find an upper bound for $\cos \phi$ with a similar argument.

$$\begin{aligned} \cos \phi &\leq \frac{(|ab| + \Pi)^2 + (|bc| + \Pi)^2 - (|ac| - \Pi)^2}{2(|ab| - \Pi)(|bc| - \Pi)} \\ \cos \phi &\leq \frac{|ab|^2 + |bc|^2 - |ac|^2 + 2\Pi(|ab| + |bc| - |ac|) + 3\Pi^2}{2(|ab||bc| - \Pi(|ab| + |bc|) + \Pi^2)} \\ \cos \phi &\leq \frac{2\Pi(|ab| + |bc| - |ac|) + 3\Pi^2}{2(|ab||bc| - \Pi(|ab| + |bc|) + \Pi^2)} \\ \cos \phi &\leq \frac{2\Pi|ab| + 3\Pi^2}{2|ab||bc| - |ab||bc|} \\ \cos \phi &\leq \frac{2\Pi|ab| + \Pi|ab|}{|ab|^2} \\ \cos \phi &\leq \frac{3\Pi}{|ab|} \end{aligned}$$

Therefore, we have that

$$\frac{-2\Pi}{|ab|} = \frac{-2\Pi}{|bc|} \leq \cos \phi \leq \frac{3\Pi}{|ab|} = \frac{3\Pi}{|bc|}.$$

Now we compute the maximum distance between c and c' . If we assume that $a = a'$ is at the origin of the plane, and that a and b' lie on the y -axis, then the

y -coordinate of c' is $|ab'| + |b'c'| \cos \phi$. This expression is bounded from below by the following.

$$|ab'| + |b'c'| \cos \phi \geq |ab| - \Pi - (|bc| + \Pi)(2\Pi/|bc|)$$

$$|ab'| + |b'c'| \cos \phi \geq |ab| - \Pi - (2|bc|)(2\Pi/|bc|)$$

$$|ab'| + |b'c'| \cos \phi \geq |ab| - \Pi - 4\Pi$$

$$|ab'| + |b'c'| \cos \phi \geq |ab| - 5\Pi$$

The y -coordinate of c' is likewise bounded from above as follows.

$$|ab'| + |b'c'| \cos \phi \leq |ab| + \Pi + (|bc| + \Pi)(3\Pi/|bc|)$$

$$|ab'| + |b'c'| \cos \phi \leq |ab| + \Pi + (2|bc|)(3\Pi/|bc|)$$

$$|ab'| + |b'c'| \cos \phi \leq |ab| + \Pi + 6\Pi$$

$$|ab'| + |b'c'| \cos \phi \leq |ab| + 7\Pi$$

Thus the y -coordinates of c and of c' differ by at most 7Π .

The x -coordinate of c' is $|b'c'| \sin \phi$. Because $\sin^2 \phi + \cos^2 \phi = 1$, we obtain that $|\sin \phi| \geq 1 - |\cos \phi|$, and thus that $\sin \phi \geq 1 - 3\Pi/|bc|$. The x -coordinate of c is bounded from below by the following.

$$|b'c'| \sin \phi \geq (|bc| - \Pi)(1 - 3\Pi/|bc|)$$

$$|b'c'| \sin \phi \geq |bc| - \Pi - (|bc| - \Pi)(3\Pi/|bc|)$$

$$|b'c'| \sin \phi \geq |bc| - \Pi - (|bc|)(3\Pi/|bc|)$$

$$|b'c'| \sin \phi \geq |bc| - \Pi - 3\Pi$$

$$|b'c'| \sin \phi \geq |bc| - 4\Pi$$



Figure 7.17: A spacer.

The x -coordinate is maximized at $|bc| + \Pi$, when ϕ is $\pi/2$. Thus the x -coordinates of c and of c' differ by at most 4Π . Since we assumed that abc and $a'b'c'$ lie in the same horizontal plane, the discrepancy in position between the c and c' is at most $(\sqrt{7^2 + 4^2})\Pi = \sqrt{65}\Pi$.

Thus the maximum discrepancy in any of the three endpoints is at most $\sqrt{65}\Pi$. Since all other points in the frame belong to graduated rulers whose allowed perturbations are at most Π , the allowed perturbation of the entire frame is at most $(1 + \sqrt{65})\Pi$. ■

Now that we have proven that the frame of the planar net is polynomially quasi-rigid, we add bars to the frame of the net. We use a structure which we call a *spacer*, shown in figure 7.17. A spacer is simply a bar with two staples at the ends. (We will always draw spacers in blue, and will often abbreviate them by leaving off the red staples such as in figure 7.19.) By linking the staples of the spacer with the hooks of graduated rulers, the spacer forms a distance constraint as performed in figure 7.18. As an example, we demonstrate how one could create the structure in the figure using a single chain by connecting the components together by the dotted brown line. It is not difficult to connect two or more simple structures by simply joining the endpoints. If both structures are meant to be rigid, adding edges between them can only serve as additional constraints to possible motions. We will adopt the convention of using dotted brown lines to connect fundamental structures throughout the remainder of this chapter.

We connect the opposing hooks of the planar net frame with spacers as shown in




Figure 7.18: Two rulers joined by three spacers.





Figure 7.19: Stapling the spacers to complete the planar net.

figure 7.19 to obtain a structure similar to a flat plane. Because the hooks are at regular intervals, we join the bottom hooks of both $\overline{ab}$ and $\overline{bc}$, a short spacer to get to the second hook, then a spacer joining the second hooks of $\overline{ab}$ and $\overline{bc}$, and so on. We can make the net out of a single chain by a short edge connecting the frame to the first spacer and short edges connecting each spacer.

Since each vertex of the spacers is adjacent to a staple, and each staple is connected to a hook on the frame, the total allowed perturbation of the planar net is only the allowed perturbation of the frame plus the total length of the staples. This proves the following lemma.

Figure 7.20: *Stereogram: A cap.*

Lemma 32. *A planar net is polynomially quasi-rigid.*

Caps

Unlike graduated rulers and planar nets, the *cap* is not flat but rather is a tetrahedral frame as illustrated in figure 7.20. To simplify the proof that a cap is polynomially quasi-rigid, we assume that abc and adc are congruent planar nets (of which b and d are right angles), and that bd is a spacer.

Lemma 33. *A cap is polynomially quasi-rigid.*

Proof. Let χ be the configuration of the cap where the following hold: (1) abc lies in the horizontal xy -plane, (2) adc lies in the vertical xz -plane, and (3) the two nets are in congruent configurations, differing only by a rotation about $\overleftrightarrow{ac}$.

Let χ' be a configuration (with points a', b', c', d') reachable from χ . We assume that a and a' are coincident, that a, c, c' lie on the x -axis, and that d' lies in the xz -plane with acd . Let Π be the maximum allowed perturbation of the two planar nets and of the spacer. Then a and a' are distance zero apart, c and c' are at most distance Π apart, and d and d' are at most distance Π apart. We must therefore only demonstrate that b and b' are apart by a distance polynomial in Π to show the cap to be polynomially quasi-rigid. Since the planar net abc is polynomially rigid, we need to prove that the z -coordinate of b' is polynomial in Π .



Figure 7.21: A staple-edge.




Figure 7.22: Labeling of the staple-edge in the proof of lemma 34.

Let e be a point on the graduated ruler ac such that $\angle bed$ forms a right angle. Because e is on both planar nets, the distance $|be| = |de|$. The distance between e and e' is at most Π , so in essence bed is the frame of a planar net, since the three distances are rigid, each within an allowed perturbation of Π . Thus, as in the proof of lemma 30, $(-2\Pi/|be|) \leq \cos \angle bed \leq (3\Pi/|be|)$.

Thus the z -coordinate of b' is bounded between -2Π and 3Π . ■

Staple-edges

A *staple-edge* is simply a long straight edge which contains one or more staples, shown in figure 7.21. For simplicity, we assume that all vertex-angles of the chain are congruent. We first prove that the staple-edge is polynomially quasi-rigid, and then discuss its purpose. All of our proofs to date have been based on the size of the staple-and-hook edges approaching zero; here we also require that the vertex-angles of the chain approach π . This will ensure that the staple-edge approaches a line segment.

Lemma 34. *The staple-edge is polynomially quasi-rigid.*

Proof. We show that the staple-edge approaches a line segment as the vertex-angles of the staples approach π , which implies that the black edges approach collinearity. Let $\overline{ab}$ be the first edge of the staple-edge, and $\overline{cd}$ be the second non-staple edge, as shown in figure 7.22.

Assume that $\overrightarrow{ab}$ points in the direction of the positive x -axis. Let $\pi - \theta$ be the measure of each vertex-angle of the chain. Then the second edge of the chain, that of the staple, is at an orientation of $\pi - \theta$ with respect to the positive x -axis. The orientation for the third edge lies in the range given by the constraints of the optimization program of figure 6.12 on page 113, and so the edge must lie at an angle between 0 and 2θ from the positive x -axis. The same argument holds for the next two edges, and thus $\overrightarrow{cd}$ lies at an orientation at most 4θ from the positive x -axis. If there are k staples in the staple-edge, then the first and last edges differ in orientation by at most $4k\theta$.

Suppose the total length of the black edges is ℓ , that the total length of the staples is K , and that $4k\theta < \pi$. The angle that each black edge makes with the x -axis is at most $4k\theta$, so the x -coordinate of the right endpoint of the chain is at least $\ell \cos 4k\theta$ minus the length of the staples, K . This value is at least $\ell[1 - (4k\theta)^2] - K$, which tends to ℓ quadratically with θ and linearly with K . Likewise, the maximum x -coordinate is at most the total length of the chain, $\ell + K$, which also tends polynomially to ℓ . The right endpoint can also be a distance $\ell \sin 4k\theta + K$ from the positive x -axis, which is less than $4\ell k\theta + K$. Thus the right endpoint has discrepancy polynomial in K and θ . This implies that every point on the staple-edge has discrepancy polynomial in K and θ , since cutting the chain at any vertex produces a smaller staple-edge with an identical θ and smaller ℓ and K . ■

Since the sum of the turning angles in the chain is $4K\theta$, we have also proven the following.

Lemma 35. *The two extreme edges of the staple-edge differ in orientation by at most the sum of its turning angles.*

The use of a staple-edge is to form an edge which is attached to the hooks of one or more graduated rulers as in figure 7.23. (We can build both the staple-edge and




Figure 7.23: A staple-edge (magenta) attached to two hooks of a graduated ruler.

the ruler out of a single chain by connecting them with the dotted brown line.) In a sense, a staple-edge is like a set of spacers with the added feature that the spacers are collinear.

7.3 Complex Quasi-rigid Structures

Using the above fundamental structures we can begin building the complex chains of the gadget. The first such structure is a spatula, which is half of the double spatula of figure 7.2.

Spatulas

A *spatula* is a combination of a staple-edge and a planar net, shown in figure 7.24. The two are bound together by the linking of two staples of the staple-edge and two hooks of the planar net. We can create the spatula out of a single chain by connecting the end of the staple-edge to the planar net as indicated by the dotted brown line.

Lemma 36. *The spatula is polynomially quasi-rigid.*

Proof. Both the planar net and the staple-edge are polynomially quasi-rigid. Let Π be the maximum allowed perturbation of either. Assume that the planar net is fixed in space, with c and d on the x -axis, modulo the allowed perturbation of Π . Then c and d can differ from their positions by at most Π . Consider figure 7.25. The staple-edge is connected to c and d , which are at least $|cd| - 2\Pi$ apart and are each at most distance Π from the x -axis.




Figure 7.24: A spatula formed by a planar net and a staple-edge.




Figure 7.25: The maximum slope of the staple-edge.




Figure 7.26: A double spatula.




Figure 7.27: Simplified drawing of a double spatula.

Therefore the slope of the staple-edge with respect to the x -axis, is $2\Pi/(|cd| - 2\Pi)$, so the maximum distance of either endpoint to the x -axis is $2\ell\Pi/(|cd| - 2\Pi)$, where ℓ is the length of the staple-edge.

The x -coordinates of the endpoints of the staple-edge are displaced by at most $2\Pi/(|cd| - 2\Pi)$ due to the perturbation in slope, and an additional 2Π due to the allowed perturbations of the planar net and the staple-edge. Thus the allowed perturbation of the spatula is polynomial with respect to ℓ and Π . ■

Double spatulas

Consider two spatulas joined together as in figure 7.26. If the angle at the leftmost vertex is very sharp, the two planar nets are brought very close together. This creates the double spatula portion of the gadget of figure 7.2. A simplified drawing of the double spatula is shown in figure 7.27.

Each half of the double spatula is polynomially quasi-rigid, but our construction leaves the possibility that one of the planar nets may rotate about the axis of the staple edge as the other stays fixed, as shown in the first two images of figure 7.28.




Figure 7.28: A double spatula with one planar net rotated.

This will not be a problem in our proof as long as the planes cannot pass through each other.

Assume the bottom plane is horizontal, and that the top plane is rotated through it as seen in the right image of figure 7.28. We can limit the extent to which the two planes pass through each other by placing a spacer at approximately distance 1 from the base of each net, as in the left image of figure 7.29. This creates the shaded region, a strip of width 1 between the frame of the planar net and the spacer.

If the two planes are close together, then the two staple-edges of the double spatula are nearly parallel. This means that the two planar nets rotate about near parallel axes. Consider the center image of figure 7.29. When the top planar net rotates through the bottom net, the lower left vertex of the top net must pass through the shaded region. A portion of the top net can only pass through the strip as long as its width does not exceed 1. Since planar nets are isosceles triangles, only the red right triangles in the figure, with leg lengths 1, can pass through. Therefore the two planes cannot pass through one another with the exception of a small piece of constant size. If the two staple-edges of the spatula are separated by a distance h , then the two planar nets can be separated by at most distance $2h$ (excepting the small red triangles), as indicated in the right image of the figure.

Thus if we wish the planar nets to be within a distance h apart (except for small pieces of size 1 at the ends of the nets), we simply need to make the staple-edges of the double spatula distance $h/2$ apart.




Figure 7.29: Computing the maximum distance between the two nets when one is rotated.




Figure 7.30: *Stereogram*: A capped edge.

Capped edges

A *capped edge* is a staple-edge with a cap at the end, like in figure 7.30. Its use will be to place caps at the end of the auxiliary chains. The proof that a capped edge is polynomially quasi-rigid is identical to that of lemma 36 for the spatula.

Lemma 37. *A capped edge is polynomially quasi-rigid.*




Figure 7.31: The auxiliary chain and cap.

Auxiliary chains

The auxiliary chain of the gadget will be identical to that of the chain used in the proof of theorem 19 for Maximum Span Between Two Planes, with the exception that the chain ends in a cap.

We will be reducing Partition to Achirality. Let $S = \{s_1, s_2, \dots, s_n\}$ be the n integers of a Partition problem, and let σ be the sum $s_1 + s_2 + \dots + s_n$. We begin with the auxiliary chain with an edge of length $256n\sigma^2$, followed by $n - 1$ edges (and thus n vertices) of length between $1/2$ and 1 , such that the turning angles at the vertices $v_2, v_3, \dots, v_{n+1}$ are proportional to the integers of S . Let $\theta_2 = s_1/\sigma$, $\theta_3 = s_2/\sigma$, et cetera. We then finish with a capped edge, where the length of the edge is $256n\sigma^2$ plus the length of the cap.

7.4 Achirality is NP-hard

We connect the auxiliary chain to the double spatula as shown in figures 7.32 and 7.33. An additional planar net (the small blue triangle) has been placed above the auxiliary chain to prevent it from moving above the double spatula construct. The




Figure 7.32: The connection of the top half of the double spatula and one auxiliary chain. (Not to scale.)

chain is forced underneath by connecting a graduated ruler (in magenta) to one of the staples under the small planar net. The only way to move it to the other side of the double spatula is to pass it through the middle.

Since the cap will not fit inside the double spatula, it must pass outside the red vertices of figure 7.33. If the auxiliary chain is to pass through the double spatula, it must have a configuration with a long enough span. We refer to the distance from one endpoint of the auxiliary chain to the beginning of the cap as its span.

We know from chapter 6 that determining the maximum possible span of the auxiliary chain when constrained by two parallel planes is NP-hard. By the proof of theorem 19, if there exists a partition of the elements of the set S , then there exists a configuration constrained by two parallel planes distance $1/(16\sqrt{2}n\sigma)$ apart whose span of at least $512n\sigma^2 - (8/n)$. If such a partition does not exist, then the largest span possible is at most $512n\sigma^2 - (8/n) - 1$. If the span is greater than $512n\sigma^2 - (8/n)$, then the two long links of the auxiliary chain must be within an angle $1/\sigma$ of parallel with each other.

Let the distance from the attached endpoint of the auxiliary chain to the red endpoints of the double spatula be $512n\sigma^2 - (8/n) - (1/2)$ as in figure 7.33. Consider




Figure 7.33: The double spatula connected to one auxiliary chain. (Not to scale.)




Figure 7.34: View from above of the long staple-edge of the double spatula and one auxiliary chain. The planar nets of the double spatula are not shown.

the moment the cap passes beyond the red endpoints. The long links of the auxiliary chain must be within $1/\sigma$ of parallel. This situation is illustrated in figure 7.34, which shows the view of figure 7.33 from directly above without the planar nets of the double spatula.

If the two long links of the auxiliary chain are within $1/\sigma$ of parallel, then the angle between the first link of the auxiliary chain and the double spatula must be less than $1/\sigma$. This constrains the short edges of the chain to lie within the region shaded in brown.

Our goal is to ensure that when the cap passes the red endpoints of the double spatula, the short edges of the chain are constrained to lie between the planar nets of the double spatula. Let the spacers of the double spatula be placed at intervals of




Figure 7.35: Placement of the spacers of the double spatula.

at most $1/8$, covering the entire brown region as in figure 7.35.

Assume the double spatula lies along the positive x -axis as in the figure. Since the total curvature of the auxiliary chain is at most 1 radian, when the short links are within the brown region, each short link makes an angle of at most $1 + (1/\sigma)$ radians with the double spatula. Since the short links are of length between $1/2$ and 1 , the difference in x -coordinates of the endpoints of each link is least than $1/4$. If the spacers are $1/8$ apart, then each short link is below at least two spacers of the upper planar net and above at least two spacers in the lower planar net, as illustrated in figure 7.36. Thus at least a quarter of each such link is trapped by the spacers and so lies between the two planar nets. Suppose the two nets of the double spatula are separated by a small distance y . If three-quarters of each link can lie outside the two planes, then the links lie up to $3y$ outside the two planes. Then each link is effectively trapped between two planes separated by $7y$.

Let the planar nets of the double spatula be positioned such that they can be opened to at most a distance $1/224\sqrt{2}n\sigma$ apart, except for the small portions within distance 1 of the extremities of the net as discussed in section 7.3. (These small portions are accounted for in figure 7.35.) If its allowed perturbation is also $1/224\sqrt{2}n\sigma$, then the farthest that the two planes could be separated is $1/112\sqrt{2}n\sigma$. The auxiliary




Figure 7.36: Each link of the central chain is trapped from above by at least two spacers of the double spatula.

chain will then be constrained to lie between two planes distance $1/16\sqrt{2}n\sigma$ apart.

We now determine the size of the brown region. Its leftmost point is at distance $256n\sigma^2 \cos(1/\sigma)$ from the attachment of the auxiliary chain; its rightmost point is at distance $256n\sigma^2 + n$ from the attachment, as shown in figure 7.35. As long as the length of the spacers is at least $512n\sigma^2 \sin(1/\sigma) + 2$, the brown region is covered. We can bound these trigonometric expressions as follows.

$$\begin{aligned}
 256n\sigma^2 \cos(1/\sigma) &\geq 256n\sigma^2(1 - \frac{1}{2\sigma^2}) \\
 256n\sigma^2 \cos(1/\sigma) &\geq 256n\sigma^2 - 128n \\
 512n\sigma^2 \sin(1/\sigma) + 2 &\leq 512n\sigma^2(1/\sigma) + 2 \\
 512n\sigma^2 \sin(1/\sigma) + 2 &\leq 512n\sigma + 2
 \end{aligned}$$

Therefore the region is contained within a rectangle of width $128n$ and length

$512n\sigma + 2$. If we place each spacer of the two planar nets at intervals of $1/16$ instead of at $1/8$, we require $2048n$ spacers, each of length about $512n\sigma + 2$. If the allowed perturbation of each planar net is less than $1/16$, then we are guaranteed to have one spacer every distance $1/8$ as desired.

At the moment the cap passes beyond the red endpoints, the short links are between the planar nets, and the two long links must also pass through the double spatula. Thus the entire auxiliary chain is bounded between the two close planes. The span of the auxiliary chain will be at least $512n\sigma^2 - (8/n)$ if it can pass through, and at most $512n\sigma^2 - (8/n) - 1$ if it cannot. Thus if the allowed perturbation of the chain is less than $1/4$, perturbations will not affect the possibility of the auxiliary chain's passage.

We now prove that the reduction can be performed in polynomial complexity. First, note that we have proven all substructures are polynomially quasi-rigid. Since we require quasi-rigidity only up to $1/(224\sqrt{2}n\sigma)$, we need only polynomial precision in computing the tiny sizes of staples, hooks, and vertex-angles. Thus no coordinate value of any vertex requires more than polynomially many significant digits with respect to the input size of the Partition instance. Furthermore, no edge of the structure is longer than about $512n\sigma^2$, so no coordinate value of any vertex is larger than polynomial in the values of the Partition instance.

In total, the structure contains a double spatula containing two staple edges with $O(1)$ edges and two planar nets with $O(n)$ edges, and an auxiliary chain containing $O(n)$ edges. Thus the complexity of the chain is polynomial. Therefore it is NP-hard to decide whether the auxiliary chain can pass through the double spatula.

We are at last ready to prove the following theorem.

Theorem 27. *Achirality is NP-hard.*

Proof. We connect the double spatula to two identical auxiliary chains as in figure 7.37. The chain can be divided exactly in half by cutting it at the brown vertex




Figure 7.37: The chain constructed to prove that Achirality is NP-hard. (Not to scale.)

at the bottom of the double spatula. If one creates the chain as a palindrome and builds the top half in the mirror image of the bottom half, the chain is achiral if and only if the two auxiliary chains can switch places! This can be performed only if the cap can fit outside the endpoints of the spatula, and determining if this is possible is NP-hard. ■

Chapter 8

Conclusion

We began our research by convexifying polygons in two and three dimensions. We proved that in the plane, all convex configurations of a given polygon are reachable from any other. We duplicated this result using pivots and disproved a conjecture by Wegner concerning the number of deflations admissible by a planar polygon. We further provided efficient algorithms for convexifying planar monotone polygons and three-dimensional polygons that admit simple projections.

In chapters 5 through 7, we considered a model of a molecule often used by the chemistry and physics communities, that of a polygonal chain with fixed vertex-angles. We described an algorithm to determine the feasibility of a single dihedral rotation and proved lower bounds on the computational complexity of preprocessing a chain to determine the feasibility of multiple rotations. We continued by proving the intractability of several questions: (1) computing the maximum possible and minimum possible distance between the endpoints of the chain, (2) determining whether certain canonical configurations are in any given component of the configuration space, and (3) determining achirality, whether a chain can be reconfigured into its mirror image. This last problem also implies the intractability of the most general problem of reconfiguring chains with fixed vertex-angles, that of determining whether a chain

can be moved from a starting configuration to a target.

We have only touched the surface of the numerous applications of geometry in this vast field. Several questions remain unanswered, and we conclude with a list of open problems suggested by this discussion.

Open Problem 1. *The algorithm for Dihedral Rotation presented in section 5.1 has a worst-case performance of $\Theta(n^2)$ time, but the lower bound provided is $\Omega(n \log n)$. Close the gap, either with a stronger lower bound or a more efficient algorithm.*

Open Problem 2. *Design an approximation algorithm for the Minimum Span problem.*

We proved the intractability of determining if a chain with fixed vertex-angles could be flattened in chapter 6. Consider two configurations A and B of a chain. If both could be flattened, then it would not directly follow that A could be reconfigured into B . It would remain to show that one can reconfigure a chain from a planar configuration to any other. Even small chains such as in figure 8.1 do not admit simple motions between any two of its configurations. In the two configurations shown, moving from one to the other could be achieved by simultaneous dihedral rotations at each of the red edges. However, neither rotation can be completed before the other starts; otherwise the chain would self-intersect. Thus we begin to have a more complex motion planning problem than it originally seems. Alternatively, one could start by moving half of the chain with a rotation at the bottom edge as in figure 8.2, but here we introduce a non-local maneuver for such a minor difference in configuration.

It can be shown that if all turning angles of the chain are at most $\pi/2$, or if they are equal, then the chain can be reconfigured between any two of its planar configurations, but the general case remains open.

Open Problem 3. *Can all chains be reconfigured between any two of their planar configurations?*




Figure 8.1: A chain and two planar configurations.




Figure 8.2: Moving between the two configurations.

Open Problem 4. *Consider only those chains with unit-length links and/or uniform vertex-angles. Are the problems of chapters 6 and 7 still intractable in this restricted domain?*

Where fixed vertex-angles are concerned, we have dealt solely with chains. Introducing a cycle in such a structure creates complex constraints on the possible motions, as demonstrated by the regular hexagons of figure 1.27 on page 22.

Open Problem 5. *Explore the above problems for polygons or geometric graphs with fixed angles between adjacent edges.*

Bibliography

- [1] Pankaj K. Agarwal and Micha Sharir. Red-blue intersection detection algorithms, with applications to motion planning and collision detection. *SIAM Journal of Computing*, 19(2):297–321, 1990.
- [2] Hee-Kap Ahn, Prosenjit Bose, Jurek Czyzowicz, Nicolas Hanusse, Evangelos Kranakis, and Pat Morin. Flipping your lid. *Geombinatorics*, 10(2):57–63, 2000.
- [3] Oswin Aichholzer, David Bremner, Erik Demaine, Henk Meijer, Vera Sácristan, and Michael Soss. Long proteins with unique optimal foldings in the H-P model. In *Proceedings of the Seventeenth European Workshop on Computational Geometry*, March 2001.
- [4] Oswin Aichholzer, Carmen Cortés, Erik D. Demaine, Vida Dujmović, Jeff Erickson, Henk Meijer, Mark Overmars, Belén Palop, Suneeta Ramaswami, and Godfried T. Toussaint. Flipping polygons. *Discrete and Computational Geometry*, to appear.
- [5] T. Alper, W. A. Cramp, D. A. Haig, and M. C. Clarke. Does the agent of scrapie replicate without nucleic acid? *Nature*, 214:764–766, 1967.

- [6] T. Alper, D. A. Haig, and M. C. Clarke. The exceptionally small size of the scrapie agent. *Biochemical and Biophysical Research Communications*, 22:278–284, 1966.
- [7] C. B. Anfinsen. Principles that govern the folding of protein chains. *Science*, 181:223–230, 1973.
- [8] Esther Arkin, Michael Bender, Erik Demaine, Martin Demaine, Joseph Mitchell, Saurabh Sethia, and Steven Skiena. When can you fold a map? In *Proceedings of the Seventh Workshop on Algorithms and Data Structures*, Providence, Rhode Island, August 2001.
- [9] Esther Arkin, Sándor Fekete, Joseph Mitchell, and Steve Skiena. On the manufacturability of paperclips and sheet metal structures. Technical report, Department of Applied Mathematics and Statistics, State University of New York at Stony Brook, 1999.
- [10] Boris Aronov, Jacob E. Goodman, and Richard Pollack. Convexification of planar polygons in $\mathbb{R}^3$. Unpublished manuscript, October 1999.
- [11] Ivan J. Balaban. An optimal algorithm for finding segment intersections. In *Proceedings of the Eleventh Annual ACM Symposium on Computational Geometry*, pages 211–219, 1995.
- [12] L. M. Balbes, S. W. Mascarella, and D. B. Boyd. A perspective of modern methods in computer-aided drug design. In K. B. Lipkowitz and D. B. Boyd, editors, *Reviews in Computational Chemistry*, volume 5, pages 337–379. VCH Publishers, New York, 1994.
- [13] Michael A. Baldwin, Keh-Ming Pan, Jack Nguyen, Ziwei Huang, Darlene Groth, Ana Serban, Maria Gasset, Ingrid Mehlhorn, Robert J. Fletterick, Fred E. Cohen, and Stanley B. Prusiner. Spectroscopic characterization of conformational

- differences between PrP^{C} and PrP^{Sc} : an α -helix to β -sheet transition. *Philosophical Transactions of the Royal Society of London, Series B*, 343(1306):435–441, March 1994.
- [14] Gill Barequet, Matthew Dickerson, and David Eppstein. On triangulating three-dimensional polygons. *Computational Geometry: Theory and Applications*, 10(3):155–170, 1998.
- [15] Joseph-Achille Le Bel. Sur les relations qui existent entre les formules atomiques des corps organiques et le pouvoir rotatoire de leurs dissolutions. *Bulletin de la Société chimique de Paris*, 22:337–347, November 1874. Translated into English in Richardson [133], pages 47–59.
- [16] M. Ben-Or. Lower bounds for algebraic computation trees. In *Proceedings of the Fifteenth Annual Symposium on Theory of Computing*, pages 80–86, 1983.
- [17] H. Benoit. Calcul de l'écart quadratique moyen entre les extrémités de diverses chaînes moléculaires de type usuel. *Journal of Polymer Science*, 3(3):376–388, 1948.
- [18] Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, 1997.
- [19] B. Berger and T. Leighton. Protein folding in the hydrophilic-hydrophobic (HP) model is NP-complete. *Journal of Computational Biology*, 5(1):27–40, 1998.
- [20] 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. In *Proceedings of the Tenth ACM-SIAM Symposium on Discrete Algorithms*, pages 866–867, 1999. Expanded version available at <http://www.arXiv.org/ps/cs.CG/9910009>.

- [21] T. Biedl, E. Demaine, M. Demaine, S. Lazard, A. Lubiw, J. O'Rourke, S. Robbins, I. Streinu, G. Toussaint, and S. Whitesides. On reconfiguring tree linkages: trees can lock. In *Proceedings of the Tenth Canadian Conference on Computational Geometry*, pages 4–5, 1998.
- [22] Therese Biedl, Erik Demaine, Sylvain Lazard, Steven Robbins, and Michael Soss. Convexifying monotone polygons. In *Proceedings of the Tenth International Symposium on Algorithms and Computation*, pages 415–424, Chennai, India, December 1999.
- [23] R. H. Bing and N. D. Kazarinoff. On the finiteness of the number of reflections that change a nonconvex plane polygon into a convex one. *Matematicheskoe Prosveshchenie*, 6:205–207, 1961.
- [24] Prosenjit Bose, Francisco Gomez, Pedro Ramos, and Godfried T. Toussaint. Drawing nice projections of objects in space. In *Symposium on Graph Drawing, GD'95*, volume 1027 of *Lecture Notes in Computer Science*, pages 52–63. Springer-Verlag, 1996.
- [25] F. Bueche. The effect of interactions in real polymer chains. *Journal of Chemical Physics*, 21:205–208, 1953.
- [26] H. Büeler, A. Aguzzi, A. Sailer, R. A. Greiner, P. Autenried, M. Aguet, and C. Weissmann. Mice devoid of PrP are resistant to scrapie. *Cell*, 73(7):1339–1347, July 1993.
- [27] Jason Cantarella and Heather Johnston. Nontrivial embeddings of polygonal intervals and unknots in 3-space. *Journal of Knot Theory and its Ramifications*, 7(8):1027–1039, 1998.

- [28] Sergio Caracciolo, Andrea Pelissetto, and Alan D. Sokal. Nonlocal Monte Carlo algorithm for self-avoiding walk with fixed endpoints. *Journal of Statistical Physics*, 60:1–53, 1990.
- [29] Augustin L. Cauchy. Sur les polygones et les polyèdres, seconde mémoire. *Journal de l'École Polytechnique*, 16(9):26–38, 1813.
- [30] H. S. Chan and K. A. Dill. The protein folding problem. *Physics Today*, pages 24–32, February 1993.
- [31] S. Chandrasekhar. Stochastic problems in physics and astronomy. *Reviews of Modern Physics*, 15(1):1–89, 1943.
- [32] Bernard Chazelle, Herbert Edelsbrunner, Leonidas J. Guibas, Micha Sharir, and Jack Snoeyink. Computing a face in an arrangement of line segments and related problems. *SIAM Journal of Computing*, 22:1286–1302, 1993.
- [33] Gustave Choquet. Variétés et corps convexes. In *Association Française pour l'Avancement des Sciences*, pages 91–93, Congrès de la Victoire, 1945.
- [34] D. E. Clark, G. Jones, P. Willett, P. W. Kenny, and R. C. Glen. Pharmacophoric pattern matching in files of three-dimensional chemical structures: comparison of conformational searching algorithms for flexible searching. *Journal of Chemical Information and Computer Sciences*, 34:197–206, 1994.
- [35] Roxana Cocan and Joseph O'Rourke. Polygonal chains cannot lock in 4D. In *Proceedings of the Eleventh Canadian Conference on Computational Geometry*, Vancouver, August 1999.
- [36] Fred E. Cohen and Stanley B. Prusiner. Structural studies of prion diseases. In S. Prusiner, editor, *Prion Biology and Diseases*, chapter 5, pages 191–228. Cold Harbor Laboratory Press, New York, 1999.

- [37] Robert Connelly, Erik D. Demaine, and Günter Rote. Every polygon can be untangled. In *Proceedings of the Sixteenth European Workshop on Computational Geometry*, Eliat, Israel, March 2000.
- [38] P. Crescenzi, D. Goldman, C. Papadimitriou, A. Piccolboni, and M. Yannakakis. On the complexity of protein folding. *Journal of Computational Biology*, 5(3):423–466, 1998.
- [39] J. Cuillé and P. Chelle. La maladie dite tremblante du mouton, est-elle inoculable? *Comptes rendu de l'Academie des Sciences*, 203:1552–1554, 1936.
- [40] J. Cuillé and P. Chelle. Transmission experimentale de la tremblante chez la chèvre. *Comptes rendu de l'Academie des Sciences*, 208:1058–1060, 1939.
- [41] John G. Curro. Computer simulation of multiple chain systems—the effect of density on the average chain dimensions. *Journal of Chemical Physics*, 61(3):1203–1207, 1974.
- [42] John G. Curro. Computer simulation of multiple chain systems—equation of state of hard sphere chains. *Journal of Chemical Physics*, 64(6):2496–2500, 1976.
- [43] J. M. Díaz-Banez, F. Gomez, and G. T. Toussaint. Partial deflation of closed polygons in the plane. Unpublished manuscript, 1999.
- [44] K. A. Dill. Dominant foci in protein folding. *Biochemistry*, 29(31):7133–7155, August 1990.
- [45] C. Domb and M. F. Sykes. Use of series expansions for the Ising model susceptibility and excluded volume problem. *Journal of Mathematical Physics*, 2(1):63–67, 1961.

- [46] Peter Eades. Symmetry finding algorithms. In G. T. Toussaint, editor, *Computational Morphology*, pages 41–51. North-Holland, Amsterdam, Netherlands, 1988.
- [47] Paul van Eikeren. Commercial manufacture of chiral pharmaceuticals. In Sander Ahuja, editor, *Chiral Separations: Applications and Technology*. American Chemical Society, 1997.
- [48] I. Z. Emiris and B. Mourrain. Computer algebra methods for studying and computing molecular conformations. *Algorithmica*, 25:372–402, 1999.
- [49] Paul Erdős. Problem number 3763. *American Mathematical Monthly*, 42:627, 1935.
- [50] Jeff Erickson. Lower bounds for linear satisfiability problems. In *Proceedings of the Sixth ACM-SIAM Symposium on Discrete Algorithms*, pages 388–395, 1995.
- [51] H. Everett, S. Lazard, S. Robbins, H. Schröder, and S. Whitesides. Convexifying star-shaped polygons. In *Proceeding of the Tenth Canadian Conference on Computational Geometry*, pages 2–3, 1998.
- [52] Henry Eyring. The resultant electric moment of complex molecules. *Physical Review*, 39:746–748, 1932.
- [53] Rida T. Farouki. Curves from motion, motion from curves. In P.-J. Laurent, P. Sablonnière, and L. Schumaker, editors, *Curve and Surface Design: Saint-Malo 1999*, pages 63–90. Vanderbilt University Press, 2000.
- [54] P. Finn, D. Halperin, L. Kavraki, J.-C. Latombe, R. Motwani, C. Shelton, and S. Venkatsubramanian. Geometric manipulation of flexible ligands. In M. C. Lin and D. Manocha, editors, *Applied Computational Geometry: Towards*

- Geometric Engineering*, volume 1148 of *Lecture Notes in Computer Science*, pages 67–78. Springer-Verlag, 1996.
- [55] Erica Flapan. Topological techniques to detect chirality. In P. Mezey, editor, *New Developments in Molecular Chirality*, pages 209–239. Kluwer Academic Publishers, 1991.
- [56] Erica Flapan. *When Topology Meets Chemistry: A Topological Look at Molecular Chirality*. Cambridge University Press, Cambridge, 2000.
- [57] P. J. Flory. The configuration of a real polymer chain. *Journal of Chemical Physics*, 17:303–310, 1949.
- [58] Aviezri S. Fraenkel. Complexity of protein folding. *Bulletin of Mathematical Biology*, 55(6):1199–1210, 1993.
- [59] Maxim D. Frank-Kamenetskii. *Unravelling DNA*. Addison-Wesley, 1997.
- [60] J. J. Freire and A. Horta. Mean reciprocal distances of short polymethylene chains. Calculation of the translational diffusion coefficient of *n*-alkanes. *Journal of Chemical Physics*, 65:4049–4054, 1976.
- [61] Anka Gajentaan and Mark H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom. Theory Appl.*, 5:165–185, 1995.
- [62] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, New York, NY, 1979.
- [63] R. A. Gibbons and G. D. Hunter. Nature of the scrapie agent. *Nature*, 215:1041–1043, September 1967.
- [64] N. Gö and H. A. Scheraga. Ring closure and local conformational deformations of chain molecules. *Macromolecules*, 3(2):178–187, 1980.

- [65] Jack Graver, Brigitte Servatius, and Herman Servatius. *Combinatorial Rigidity*. American Mathematical Society, Providence, 1993.
- [66] U. Grenander, Y. Chow, and D. M. Keenan. *Hands: A Pattern Theoretic Study of Biological Shapes*. Springer-Verlag, 1991.
- [67] Ulf Grenander. Can one understand shapes in nature? Seminar talk, Indiana University, March 1987.
- [68] J. S. Griffith. Self-replication and scrapie. *Nature*, 215:1043–1044, September 1967.
- [69] T. B. Grimley. High-polymer solutions. *Proceedings of the Royal Society of London, Series A*, 212:339–361, 1952.
- [70] Branko Grünbaum. How to convexify a polygon. *Geombinatorics*, 5:24–30, 1995.
- [71] Leonidas J. Guibas, Micha Sharir, and Shmuel Sifrony. On the general motion planning problem with two degrees of freedom. *Discrete and Computational Geometry*, 4:491–521, 1989.
- [72] Brian Hayes. Fast protein folding in the Hydrophilic-Hydrophobic model within three-eighths of optimal. In *Proceedings of the Twenty-seventh Annual ACM Symposium on the Theory of Computing*, pages 157–168, Las Vegas, May 1995.
- [73] Brian Hayes. Prototeins. *American Scientist*, 86(3):216–221, May 1998.
- [74] Sir Thomas Little Heath. *The thirteen books of Euclid's Elements translated from the text of Heiberg with introduction and commentary*. Dover Publications, New York, 1956.
- [75] J. J. Hermans. An estimate of the volume effect in coiling long chain molecules. *Recueil des travaux chimiques des Pays-Bas*, 69:220–224, 1950.

- [76] Dorit S. Hochbaum and J. George Shanthikumar. Convex separable optimization is not much harder than linear optimization. *Journal of the Association of Computing Machinery (ACM)*, 37(4):843–862, October 1990.
- [77] Jacobus Hendrikus van’t Hoff. A suggestion looking to the extension into space of the structural formulas at present used in chemistry. And a note upon the relation between the optical activity and the chemical constitution of organic compounds. Utrecht, 1874. Translated into English in Richardson [133], pages 35–45.
- [78] John Hopcroft, Deborah Joseph, and Sue Whitesides. On the movement of robot arms in 2-dimensional bounded regions. *SIAM Journal of Computing*, 14(2):315–333, May 1985.
- [79] John E. Hopcroft and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.
- [80] A. E. Howard and P. A. Kollman. An analysis of current methodologies for conformational searching of complex molecules. *Journal of Medical Chemistry*, 31(9):1669–1675, 1988.
- [81] Z. Huang, J.-M. Gabriel, M. A. Baldwin, R. J. Fletterick, S. B. Prusiner, and F. E. Cohen. Proposed three-dimensional structure for the cellular prion protein. *Proceedings of the National Academy of Science of the United States of America*, 91(15):7139–7143, July 1994.
- [82] Ziwei Huang, Stanley B. Prusiner, and Fred E. Cohen. Scrapie prions: A three-dimensional model of an infectious fragment. *Folding and Design*, 1:13–19, December 1995.

- [83] G. D. Hunter and G. C. Millson. Studies on the heat stability and chromatographic behaviour of the scrapie agent. *Journal General Microbiology*, 37:251–259, 1964.
- [84] G. D. Hunter and G. C. Millson. Attempts to release the scrapie agent from tissue debris. *Journal of Comparative Pathology*, 77:301–307, 1967.
- [85] Donald J. Jacobs, Leslie A. Kuhn, and Michael F. Thorpe. Flexible and rigid regions in proteins. In M. F. Thorpe and P. M. Duxbury, editors, *Rigidity Theory and Applications*, pages 357–384. Kluwer Academic/Plenum Publishers, 1999.
- [86] T. Kaluza. Problem 2: Konvexieren von Polygonen. *Mathematische Semesterberichte*, 28:153–154, 1981.
- [87] N. D. Kazarinoff. *Analytic Inequalities*. Holt, Rinehart and Winston, 1961.
- [88] N. D. Kazarinoff and R. H. Bing. A finite number of reflections render a non-convex plane polygon convex. *Notices of the American Mathematical Society*, 6:834, 1959.
- [89] A. K. Kron. A Monte Carlo method in the statistical calculations of macromolecules. *Vysokomolekulyarnye Soedineniya*, 7:1228–1234, 1965.
- [90] Werner Kuhn. Über die Gestalt fadenförmiger Moleküle in Lösungen. *Kolloid Zeitschrift*, 68(1):2–15, 1934.
- [91] Werner Kuhn and Hans Kuhn. Die Frage nach der Aufrollung von Fadenmolekeln in strömenden Lösungen. *Helvetica Chimica Acta*, 26:1394–1465, 1943.
- [92] Moti Lal. ‘Monte Carlo’ computer simulation of chain molecules. I. *Molecular Physics*, 17(1):57–64, 1969.

- [93] A. R. Leach and I. D. Kuntz. Conformational analysis of flexible ligands in macromolecular receptor sites. *Journal of Computational Chemistry*, 13(6):730–748, 1992.
- [94] William Lenhart and Sue Whitesides. Turning a polygon inside-out. In *Proceedings of the Third Canadian Conference on Computational Geometry*, pages 66–69, Vancouver, Canada, August 1991.
- [95] William Lenhart and Sue Whitesides. Reconfiguring closed polygonal chains in Euclidean d -space. *Discrete and Computational Geometry*, 13:123–140, 1995.
- [96] Cyrus Levinthal. How to fold gracefully. In J. T. P. DeBrunner and E. Munck, editors, *Mossbauer Spectroscopy in Biological Systems: Proceedings of a meeting held at Allerton House, Monticello, Illinois*, pages 22–24. University of Illinois Press, 1969.
- [97] Harry R. Lewis and Christos H. Papadimitriou. *Elements of the Theory of Computation*. Prentice-Hall, Englewood Cliffs, New Jersey, second edition, 1998.
- [98] L. A. Lyusternik. *Convex Figures and Polyhedra*. D. C. Heath and Company, Boston, 1966. Translated and adapted from the first Russian edition (1956) by Donald L. Barnett.
- [99] Bruce MacDonald, Naeem Jan, D. L. Hunter, and M. O. Steinitz. Polymer conformations through ‘wiggling’. *Journal of Physics A: Mathematical and General*, 18:2627–2631, 1985.
- [100] 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.

- [101] Neal Madras and Gordon Slade. *The Self-Avoiding Walk*. Birkhäuser, Boston, 1993.
- [102] Neal Madras and Alan D. Sokal. The pivot algorithm: A highly efficient Monte Carlo method for the self-avoiding walk. *Journal of Statistical Physics*, 50:109–186, 1988.
- [103] J. Michael McCarthy. *Geometric Design of Linkages*. Springer-Verlag, New York, 2000.
- [104] D. S. McKenzie. Polymers and scaling. *Physics Reports (Section C of Physics Letters)*, 27(2):35–88, 1976.
- [105] F. Meyer auf der Heide. A polynomial time linear search algorithm for the n -dimensional knapsack problem. *Journal of the Association of Computing Machinery (ACM)*, 31:668–676, 1984.
- [106] Paul Mezey. *Shape in Chemistry: An Introduction to Molecular Shape and Topology*. VCH Publishers, New York, 1993.
- [107] Kenneth Millett. Algebraic topological indices of molecular chirality. In P. Mezey, editor, *New Developments in Molecular Chirality*, pages 165–207. Kluwer Academic Publishers, 1991.
- [108] Kenneth Millett. Knotting of regular polygons in 3-space. *Journal of Knot Theory and its Ramifications*, 3:263–278, 1994.
- [109] K. Mislow and R. Bolstad. Molecular dissymmetry and optical inactivity. *Journal of the American Chemical Society*, 77:6712–6713, 1955.
- [110] Joseph Mitchell. A counterexample(?) to the carpenter’s ruler conjecture. Technical report, SUNY Stony Brook Department of Applied Mathematics and Statistics, 1999.

- [111] Béla de Szökefalvi Nagy. Solution of problem 3763. *American Mathematical Monthly*, 46:176–177, 1939.
- [112] Arnold Neumaier. Molecular modeling of proteins and mathematical prediction of protein structure. *SIAM Review*, 39(3):407–460, 1997.
- [113] O. F. Olaj and K. H. Pelinka. Pair distribution function and pair potential of lattice models under theta condition. I. Numerical evaluation. *Makromolekulare Chemie*, 177:3413–3425, 1976.
- [114] Joseph O’Rourke. *Computational Geometry in C*. Cambridge University Press, second edition, 1998.
- [115] Joseph O’Rourke. Folding and unfolding in computational geometry. In *Proceedings of the Japan Conference on Discrete and Computational Geometry*, pages 142–147, December 1998.
- [116] Joseph O’Rourke. On the development of the intersection of a plane with a polytope. Technical Report 068, Smith College, June 2000.
- [117] Keh-Ming Pan, Michael Baldwin, Jack Nguyen, Maria Gasset, Ana Serban, Darlene Groth, Ingrid Mehlhorn, Ziwei Huang, Robert J. Fletterick, Fred E. Cohen, and Stanley B. Prusiner. Conversion of α -helices into β -sheets features in the formation of the scrapie prion proteins. *Proceedings of the National Academy of Science of the United States of America*, 90:10962–10966, December 1993.
- [118] Louis Pasteur. Recherches sur la dissymétrie moléculaire des produits organiques naturels: leçons professées à la Société chimique de Paris le 20 janvier et le 3 février 1860. Paris, 1860. Translated into English in Richardson [133], pages 1–33.

- [119] I. Pattison, W. Gordon, and G. Millson. Experimental production of scrapie in goats. *Journal of Comparative Pathology*, 69:300–312, 1959.
- [120] Franco P. Preparata and Michael I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, third edition, October 1990.
- [121] S. B. Prusiner, D. C. Bolton, D. F. Groth, K. A. Bowman, S. P. Cochran, and M. P. McKinley. Further purification and characterization of scrapie prions. *Biochemistry*, 21:6942–6950, 1982.
- [122] S. B. Prusiner, M. P. McKinley, D. F. Groth, K. A. Bowman, N. I. Mock, S. P. Cochran, and F. R. Masiaarz. Scrapie agent contains a hydrophobic protein. *Proceedings of the National Academy of Science of the United States of America*, 78(11):6675–6679, November 1981.
- [123] Stanley B. Prusiner. Novel proteinaceous infectious particles cause scrapie. *Science*, 216(4542):136–144, April 1982.
- [124] Stanley B. Prusiner. Molecular biology and genetics of prion diseases. *Philosophical Transactions of the Royal Society of London, Series B*, 343(1306):447–463, March 1994.
- [125] Stanley B. Prusiner. An introduction to prion biology and diseases. In S. Prusiner, editor, *Prion Biology and Diseases*, chapter 1, pages 1–66. Cold Harbor Laboratory Press, New York, 1999.
- [126] Richard Randell. Conformation spaces of molecular rings. In *Studies in Physical and Theoretical Chemistry*, pages 141–156. Elsevier Science Publishers, Amsterdam, 1988.

- [127] Richard Randell. A molecular conformation space. In *Studies in Physical and Theoretical Chemistry*, pages 125–140. Elsevier Science Publishers, Amsterdam, 1988.
- [128] John H. Reif. Complexity of the mover’s problem and generalizations. In *Proceedings of the Twentieth Annual IEEE Symposium on Foundations of Computer Science*, pages 421–427, 1979.
- [129] Johannes Reiter. Monte Carlo simulations of linear and cyclic chains on cubic and quadratic lattices. *Macromolecules*, 23:3811–3816, 1990.
- [130] E. J. Janse van Rensburg, S. G. Whittington, and N. Madras. The pivot algorithm and polygons: results on the FCC lattice. *Journal of Physics A: Mathematical and General Physics*, 23:1589–1612, 1990.
- [131] Yu. G. Reshetnyak. On a method of transforming a nonconvex polygonal line into a convex one. *Uspehi Matematicheskikh Nauk*, 12(3):189–191, 1957.
- [132] Franz Reuleaux. *The Kinematics of Machinery*. Dover Publications, New York, 1963. Translated by Alexander B. W. Kennedy in 1876 from the original, *Theoretische Kinematik: Grundzüge einer Theorie des Maschinenwesens*, first published in 1874.
- [133] G. M. Richardson, editor. *The Foundations of Stereo Chemistry: Memoirs by Pasteur, van’t Hoff, LeBel, and Wislicenus*. American Book Company, New York, 1901.
- [134] S. Roberts. Three-bar motion in plane space. *Proceedings of the London Mathematical Society*, 7:14–23, 1875.
- [135] S. A. Robertson. Inflation of plane curves. In *Geometry and Topology of Submanifolds III*, pages 264–275. World Scientific, 1991.

- [136] S. A. Robertson and B. Wegner. Full and partial inflation of plane curves. In *Intuitive Geometry, Colloquia Mathematica Societatis János-Bolyai*, pages 389–401. North-Holland, Amsterdam, 1994.
- [137] M. N. Rosenbluth and A. W. Rosenbluth. Monte carlo calculation of the average extension of molecular chains. *Journal of Chemical Physics*, 23:356–359, 1955.
- [138] H. Sachse. Über die Konfigurationen der Polymethylenringe. *Zeitschrift für physikalische Chemie*, 10:202–241, 1892.
- [139] A. Sailer, H. Büeler, M. Fischer, A. Aguzzi, and C. Weissmann. No propagation of prions in mice devoid of PrP. *Cell*, 77(7):967–968, July 1994.
- [140] G. T. Sallee. Stretching chords of space curves. *Geometriae Dedicata*, 2:311–315, 1973.
- [141] B. Sandak, R. Nussinov, and H. J. Wolfson. 3-D flexible docking of molecules. In *Proceedings of the First IEEE Workshop on Shape and Pattern Recognition in Computational Biology*, 1994.
- [142] I. J. Schoenberg and S. K. Zaremba. On Cauchy’s lemma concerning convex polygons. *Canadian Journal of Mathematics*, 19(4):1062–1071, 1967.
- [143] Georg E. Schulz and R. Heiner Schirmer. *Principles of Protein Structure*. Springer-Verlag, New York, 1979.
- [144] Jacob T. Schwartz and Micha Sharir. On the “piano movers” problem II: General techniques for computing topological properties of real algebraic manifolds. *Advances in Applied Mathematics*, 4:298–351, 1983.
- [145] Rajgopal Srinivasan and George D. Rose. LINUS: A hierarchic procedure to predict the fold of a protein. *Proteins: Structure, Function, and Genetics*, 22:81–99, 1995.

- [146] E. Steinitz and H. Rademacher. *Vorlesungen über die Theorie der Polyeder*. Julius Springer, Berlin, 1934.
- [147] Steven D. Stellman and Paul J. Gans. Computer simulation of polymer conformation. II. Distribution function for polymers with excluded volume. *Macromolecules*, 5(6):720–729, 1972.
- [148] Steven D. Stellman and Paul J. Gans. Efficient computer simulation of polymer conformation. I. Geometric properties of the hard-sphere model. *Macromolecules*, 5(4):516–526, 1972.
- [149] Ileana Streinu. A combinatorial approach to planar non-colliding robot arm motion planning. In *Proceedings of the Forty-First Annual Symposium on Foundations of Computer Science*, Redondo Beach, California, November 2000.
- [150] M. F. Sykes. Some counting theorems in the theory of the Ising model and the excluded volume problem. *Journal of Mathematical Physics*, 2(1):52–62, 1961.
- [151] William J. Taylor. Average length and radius of normal paraffin hydrocarbon molecules. *Journal of Chemical Physics*, 16(4):257–267, 1948.
- [152] Bernard Testa. Definitions and concepts in biochirality. In B. Holmstedt, H. Frank, and B. Testa, editors, *Chirality and Biological Activity*, pages 15–32. Alan R. Liss, Inc., 1990.
- [153] Edwin Thall. When drug molecules look in the mirror. *Journal of Chemical Education*, 73(6):481–484, 1996.
- [154] M. Thorpe and P. Duxbury, editors. *Rigidity Theory and Applications*. Kluwer Academic Publishers, 1999.
- [155] Michael F. Thorpe, Brandon M. Hespenheide, Yi Yang, and Leslie A. Kuhn. Flexibility and critical hydrogen bonds in cytochrome *C*. In R. B. Altman,

- A. K. Dunker, L. Hunter, K. Lauderdale, and T. E. Klein, editors, *Proceedings of the Pacific Symposium on Biocomputing*, pages 191–202. World Scientific, 1999.
- [156] Godfried T. Toussaint. Computational polygonal entanglement theory. In *VIII Encuentros de Geometría Computacional*, Castellon, Spain, July 1999.
- [157] Godfried T. Toussaint. The Erdős-Nagy theorem and its ramifications. In *Eleventh Canadian Conference on Computational Geometry*, Vancouver, Canada, August 1999.
- [158] Godfried T. Toussaint. A new class of stuck unknots in Pol_6 . *Beiträge zur Algebra und Geometrie (Contributions to Algebra and Geometry)*, 42(2):301–306, 2001.
- [159] Ron Unger and John Moulton. Finding the lowest free energy conformation of a protein is an NP-hard problem: proof and implications. *Bulletin of Mathematical Biology*, 55(6):1183–1198, 1993.
- [160] Peter H. Verdier. A simulation model for the study of the motion of random-coil polymer chains. *Journal of Computational Physics*, 4:204–210, 1969.
- [161] David M. Walba. A topological hierarchy of molecular chirality and other tidbits in topological stereochemistry. In P. Mezey, editor, *New Developments in Molecular Chirality*, pages 119–129. Kluwer Academic Publishers, 1991.
- [162] F. T. Wall and F. Mandel. Macromolecular dimensions obtained by an efficient Monte Carlo method without sample attrition. *Journal of Chemical Physics*, 63:4592–4595, 1975.
- [163] Frederick T. Wall, Stanley Windwer, and Paul J. Gans. Monte Carlo methods applied to configurations of flexible polymer molecules. In *Methods in Com-*

- putational Physics: Advances in Research and Applications*, volume 1, pages 217–243. Academic Press, 1963.
- [164] Bernd Wegner. Partial inflation of closed polygons in the plane. *Beiträge zur Algebra und Geometrie (Contributions to Algebra and Geometry)*, 34(1):77–85, 1993.
- [165] Bernd Wegner. Chord-stretching convexifications of spherical polygons. In *Proceedings of the Topology and Geometry Day, Textos de Matematica (Texts in Mathematics)*, Series B, Coimbra, May 1996.
- [166] Bernd Wegner. Convexification of spherical curves. In *Proceedings of the Fourth Congress on Geometry*, pages 424–430, Thessaloniki, 1996.
- [167] C. Weissman, H. Büeler, M. Fischer, A. Sauer, and M. Aguet. Susceptibility to scrapie in mice is dependent on PrP^C. *Philosophical Transactions of the Royal Society of London, Series B*, 343(1306):431–433, March 1994.
- [168] Walter Whiteley. Rigidity of molecular structures: Generic and geometric analysis. In M. Thorpe and P. Duxbury, editors, *Rigidity Theory and Applications*, pages 21–46. Kluwer Academic Publishers, 1999.
- [169] Sue Whitesides. Algorithmic issues in the geometry of planar linkage movement. *The Australian Computer Journal*, 24(2):42–50, May 1992.
- [170] Johannes Wislicenus. The lactic acids. *Ann. Alen. Chem.*, 167:302–346, 1873.
- [171] Johannes Wislicenus. Über die räumliche Anordnung der Atome in organischen Molekülen und ihre Bestimmung in geometrisch-isomeren ungesättigten Verbindungen (On the spatial arrangement of the atoms in organic molecules and the resulting geometrical isomerism in unsaturated compounds). *Abhandlungen der mathematischen-physikalischen Classe der königlichen sächsischen*

- Gesellschaft der Wissenschaften*, 14:1, 1887. Translated into English in Richardson [133], pages 47–59.
- [172] A. Ya. Yusupov. A property of simply-connected nonconvex polygons. *Uchen. Zap. Buharsk. God. Pedagog.*, pages 101–103, 1957.
- [173] B. H. Zimm, W. H. Stockmayer, and M. Fixman. Excluded volume in polymer chains. *Journal of Chemical Physics*, 21:1716–1723, 1953.