Reconfiguring Convex Polygons

Oswin Aichholzer* Erik D. Demaine† Jeff Erickson‡
Ferran Hurtado§ Mark Overmars¶ Michael Soss‖
Godfried T. Toussaint‖

Abstract

We prove that there is a motion from any convex polygon to any convex polygon with the same counterclockwise sequence of edge lengths, that preserves the lengths of the edges, and keeps the polygon convex at all times. Furthermore, the motion is "direct" (avoiding any intermediate canonical configuration like a subdivided triangle) in the sense that each angle changes monotonically throughout the motion. In contrast, we show that it is impossible to achieve such a result with each vertex-to-vertex distance changing monotonically. We also demonstrate that there is a motion between any two such polygons using three-dimensional moves known as pivots, although the complexity of the motion cannot be bounded as a function of the number of vertices in the polygon.

1 Introduction

This paper is concerned with *linkages* modeled by polygons (primarily in the plane), whose vertices represent hinges and whose edges represent rigid bars. A fundamental question about such linkages is whether it is possible to reach every polygon with the same sequence of edge lengths by motions that preserve the edge lengths. Several papers have shown that the answer to this question is yes for various types of polygons; we call this a *universality*

*Institut für Grundlagen der Informationsverarbeitung, Technische Universität Graz, Infeldgasse 16b/I, A-8010 Graz, Austria, email: oaich@igi.tu-graz.ac.at. Supported by the Austrian Programme for Advanced Research and Technology (APART).

†Department of Computer Science, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada, email: eddemaine@uwaterloo.ca.

‡Department of Computer Science, University of Illinois, Urbana, IL 61801, USA, email: jeffe@cs.uiuc.edu. Supported by a Sloan Fellowship.

§Departament de Matemàtica Aplicada II, Universitat Politècnica de Catalunya, Pau Gargallo 5, 08028-Barcelona, España, email: hurtado@ma2.upc.es. Supported by Proyectos MEC-DGES-SEUID PB98-0933 and CUR Gen. Cat. SGR1999-00356

¶Department of Computer Science, Utrecht University, Padualaan 14, De Uithof, 3584 CH Utrecht, The Netherlands, email: markov@cs.uu.nl.

‖School of Computer Science, McGill University, 3480 University Street, Montréal, Québec H3A 2A7, Canada, email: {soss, godfried}@cs.mcgill.ca.

result. If edges are allowed to cross each other, then this is true in every dimension [12, 16]. If edges are not allowed to cross, universality does not hold in general for polygons in 3-D [2, 5], but has been shown for polygons in the plane and motions in 3-D [1, 2], for polygons and motions in the plane [9], for polygons in 3-D with simple projections [4], and for all polygons in 4-D and higher dimensions [8].

All of these papers show universality by proving that every polygon can be *convexified*, that is, moved to a convex (planar) polygon while preserving edge lengths. Convex polygons are used as an intermediate state; because motions can be reversed and concatenated, all that remains is to show that a convex polygon can be moved to every other convex polygon with the same counterclockwise sequence of edge lengths. This fact is established in [12] when edges are allowed to cross. No proof has been published for the case in which edges cannot cross.

The basic idea in the proof in [12] of universality of convex polygons is to show how to reconfigure every convex polygon into another intermediate state, a "canonical triangle." In the first half of this paper, we show that this intermediate state can be avoided. Specifically, a convex polygon can be moved into any other convex polygon with the same counterclockwise sequence of edge lengths in such a way that each vertex angle varies monotonically with time (either never increasing or never decreasing). In this sense, the motion goes directly from the source to the destination. Our motion is also of the simplest type possible [3]: it can be decomposed into a linear number of *moves*, each of which changes only four joint angles at once.

In the second half of this paper, we study the same problem of reconfiguring convex polygons, under a more restrictive type of move. Specifically, we study motions consisting of a sequence of *pivots*, which are the simplest kind of motion in three dimensions, changing only two joint angles at once. Such motions are popular in biology and physics circles; see Section 5. It may seem that the freedom to move in three dimensions is a significant advantage, but in fact the limited motions make it difficult to change angles in the plane. Nonetheless, we show that it is possible to simulate our planar motions by a sequence of pivots. Thus we obtain the result that a convex polygon can be pivoted to any other convex polygon with the same counterclockwise sequence of edge lengths.

The paper is organized as follows. In Section 2 we introduce some basic notation that we will use throughout the paper. Section 3 proves the theorem about angle-monotone motions in the plane, using an old lemma of Cauchy and Steinitz. Section 4 shows an example in which a different type of monotonicity cannot be achieved. Finally, Section 5 proves the theorem about pivots in three dimensions.

2 Notation

For a polygon P, we denote its vertices by $v_1, \ldots, v_n$ in counterclockwise order, its edges by $e_i = (v_i, v_{i+1})$, and its edge lengths by $\ell_i = |e_i| = |v_i - v_{i+1}|$. Throughout, index arithmetic is modulo n.

A *convex configuration* of edge lengths (positive real numbers) $\ell_1, \ldots, \ell_n$ is a convex polygon with those edge lengths in counterclockwise order. The convex polygon is allowed

to have "straight" vertices with angle π. A well-known result characterizes the edge lengths for which convex configurations exist:

Lemma 1 (Lemma 3.1 of [12]) *The edge lengths $\ell_1, \ldots, \ell_n$ admit a convex configuration precisely if $\ell_i \leq \sum_{j \neq i} \ell_j$ for all i.*

A *motion* or *reconfiguration* is a continuous function from the unit interval $[0, 1]$ (representing time) to a configuration, where each *configuration* is a polygon with the same counterclockwise sequence of edge lengths. An *angle-monotone motion* is a motion in which each vertex angle is a monotone function in time.

3 Reconfiguring between Two Convex Configurations

Consider two convex configurations C and C' of the same sequence of edge lengths. We think of C as the source configuration and C' as the destination configuration. Label each angle of C by $+$ if it needs to get bigger in order to match the corresponding angle in C', by $-$ if it needs to get smaller, or by 0 if they already match.

This set up is exactly what arises in the proof of Cauchy's theorem about the rigidity of convex polyhedra [6, 10], except that in Cauchy's application the polygon is on the sphere. His key lemma about alternations in such $+, -, 0$ labelings is what we need as well. Cauchy's original proof of this lemma (in 1813) had an error, noticed and corrected over a century later by Steinitz in 1934 [19].

Lemma 2 (Cauchy-Steinitz Lemma) *In a $+, -, 0$ labeling that comes from two distinct convex configurations, there are at least four sign alternations.*

Proof (Sketch): Because the configurations are distinct, not all labels are 0. By circularity, the number of alternations between $+$ and $-$ (ignoring 0's) is even. It cannot be zero, because there is no motion of any polygon that increases or decreases all angles. It cannot be two, because then there is a chain of increasing angles and a chain of decreasing angles; the former chain specifies that the ends of the chain should get further apart, whereas the latter chain specifies the opposite. It is this last part of the argument that needs careful analysis; for details, see [19] for Steinitz's original (complicated) proof, [10] for a simpler proof due to Isaac J. Schoenberg, or [17] for another elementary proof. □

The idea is to take vertices v_i, v_j, v_k, v_l in cyclic order around the polygon, whose angles are labeled $+, -, +, -$ in that order, and flex the quadrangle defined by those vertices until one angle matches the desired value in C'. See Figure 1.

Now we need a lemma about reconfiguring convex quadrangles:

Lemma 3 *Given a convex quadrangle v_1, v_2, v_3, v_4, there is a motion that decreases the angles at v_1 and v_3, and increases the angles at v_2 and v_4. The motion can continue until one of the angles reaches 0 or π.*

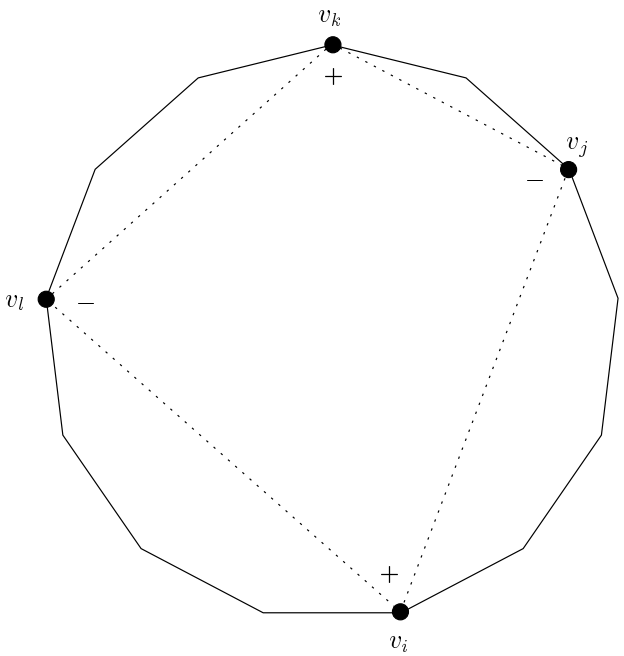

Figure 1: Applying a quadrangle motion to a convex polygon by taking vertices labeled $+, -, +, -$ in that order.

Proof: We consider the following viewpoint: v_1 is pinned to the plane, and v_3 moves along the directed line from v_1 to v_3 (see Figure 2). The motions of v_2 and v_4 are determined by maintaining their distances to v_1 and v_3. Applying Euclid's Proposition I.25[1] [11] to triangle v_1, v_2, v_3, because $|v_1 - v_3|$ is increasing, so is the angle at v_2. Similarly, the angle at v_4 is increasing throughout the motion. Because no angle goes past 0 or π, we maintain a convex quadrangle throughout the motion, so by the Cauchy-Steinitz lemma (Lemma 2), there must be at least four sign alternations when compared to any future quadrangle we will visit. This proves that the angles at v_1 and v_3 are decreasing throughout the motion. □

We are now in the position to prove the main theorem of this section:

Theorem 1 *Given two convex configurations C, C' of the same edge lengths $\ell_1, \ldots, \ell_n$, there is an angle-monotone motion from C to C' that involves $O(n)$ moves, each of which changes only four vertex angles at once.*

Proof: Consider configuration C. By Lemma 2, we can find vertices v_i, v_j, v_k, v_l in cyclic order around the polygon, whose angles are labeled $+, -, +, -$ in that order; see Figure 1. By specifying the subchains between these vertices to move rigidly, we obtain a convex quadrangle. Move this quadrangle according to Lemma 3 until one of the four angles matches the angle in C'. (No angle will ever reach 0 or π because of our stopping condition.) Repeat this process until all angles match. The result is a sequence of motions from C to C'. There

[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." [11, p. 297]

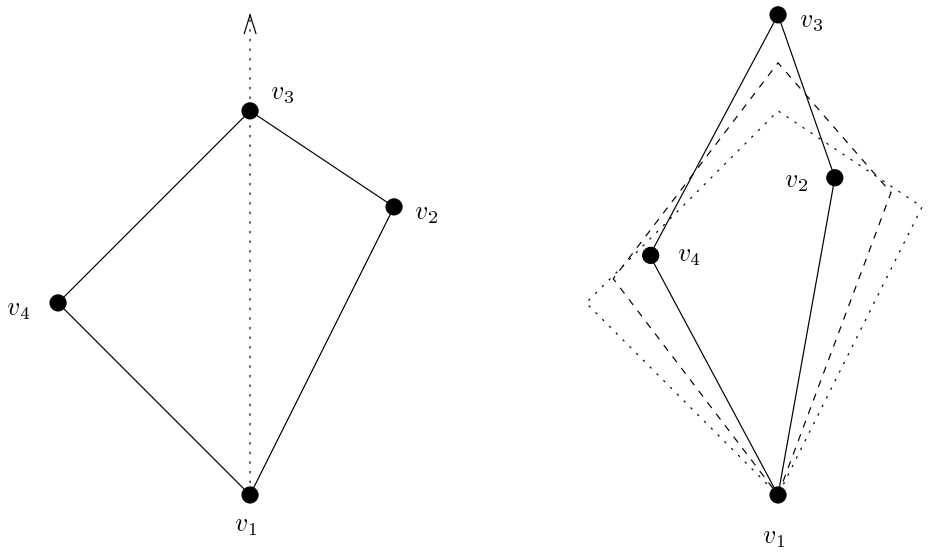

Figure 2: Moving a convex quadrangle as in Lemma 3.

are at most n moves, because each motion changes the label of an angle from $+$ or $-$ to 0, and that label persists. □

Proposition 1 *Computing the motion in Theorem 1 can be done in $O(n)$ time on a pointer machine with real numbers.*

Proof: The first part is to maintain the vertices of the quadrangle, v_i, v_j, v_k, v_l, throughout the motion. We maintain four consecutive blocks I, J, K, L of the same sign; specifically, we maintain the first and last vertex in each block. This can be found initially in linear time by scanning along the polygon's vertices in order. The desired vertices v_i, v_j, v_k, v_l are identified with the first vertex in the corresponding block. When the label of one of them switches to 0, it and the block's first vertex advance to the next element in the block. If this was the last element (the block is empty), we make the following modifications. If I becomes empty, we advance it to the block of $+$'s after L. Similarly, if L becomes empty, it retreats to the block before I. If K becomes empty, it advances to the block after L, the blocks J and L merge to produce a new J, and L advances to the block after K. The case of J becoming empty is symmetric.

The second part is to apply the quadrangle motions from Lemma 3. This involves computing the time at which the quadrangle motion stops, and then updating the coordinates. These computations can be done analogously to Lemma 7 of [3]. Basically, we compute the times at which each angle would match the desired angle in C', and take the minimum of these times. At worst, each time can be computed by solving a degree-four polynomial, which reduces to an arithmetic expression involving square and cube roots. □

4 Distance-Monotone Motions

We have shown that an angle-monotone motion between any two convex configurations of a common sequence of edge lengths can be computed in linear time. An interesting consequence is that any polygon can be moved to a unique *inscribed* configuration [18], in which the vertices lie on a common circle, a natural generalization of regular polygons.

It is interesting to note that we cannot hope for a *distance-monotone* motion between any two convex polygons, in which every distance between a pair of vertices varies monotonically with time. (This is in direct contrast to convexification of a polygon [9], where all distances can be made to increase.) An example is shown in Figure 3. Because the dotted lines are the same length in both configurations, these distances must be preserved throughout the motion; in other words, the chains v_1, v_2, v_3 and v_4, v_5, v_6 must move rigidly. The problem is thus reduced to moving a quadrangle v_1, v_3, v_4, v_6, which can be moved in only two different ways. Only one motion decreases $|v_1 - v_4|$ and increases $|v_3 - v_6|$ as desired, but then the distance $|v_2 - v_5|$ increases and later decreases. Specifically, the distance in the middle configuration is more than 0.6% larger than the (equal) distances in the left and right configurations.

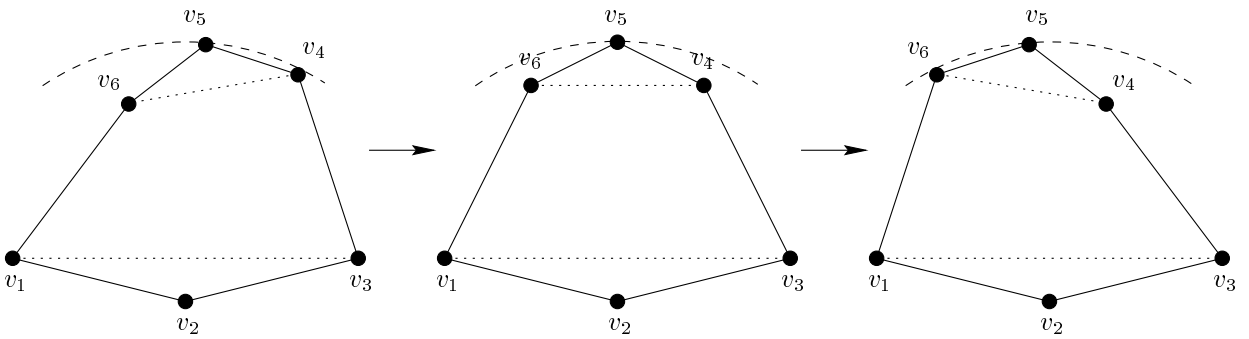

Figure 3: (Left and right) An example for which a distance-monotone motion is impossible. (Middle) The transition between $|v_2 - v_5|$ increasing and decreasing.

5 Reconfiguration with Pivots

In this section, we show that a convex polygon can be reconfigured to any other convex polygon (with the same edge lengths) by the use of three-dimensional motions called pivots. Let v_i and v_j be two (nonadjacent) vertices of a polygon. A *pivot* on $\overline{v_i v_j}$ is a motion whereby the counterclockwise section of the polygon between v_i and v_j (denoted henceforth as $[v_i, v_j]$) is rotated about the diagonal $\overline{v_i v_j}$. Examples of pivots are illustrated in Figures 4 and 5.

Pivots are of great interest to polymer physicists and molecular biologists, who consider polygons as models of large molecules and are interested in the configurations that they can take. This motion has been used in many contexts over the last few decades in both physics and mathematics [13, 14, 15, 21].

Erdős-Nagy flips [20] are special cases of pivots with planar polygons in 3D, in which the pairs of vertices that define the pivots are determined by lines of support of the polygon, and in which each rotation has angle π.

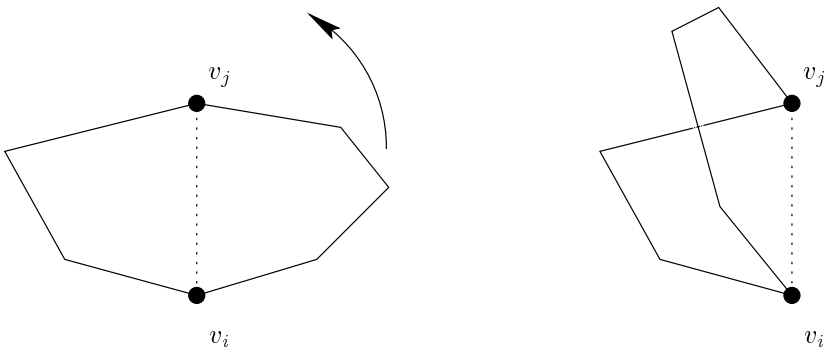

Figure 4: A pivot on $\overline{v_i v_j}$.

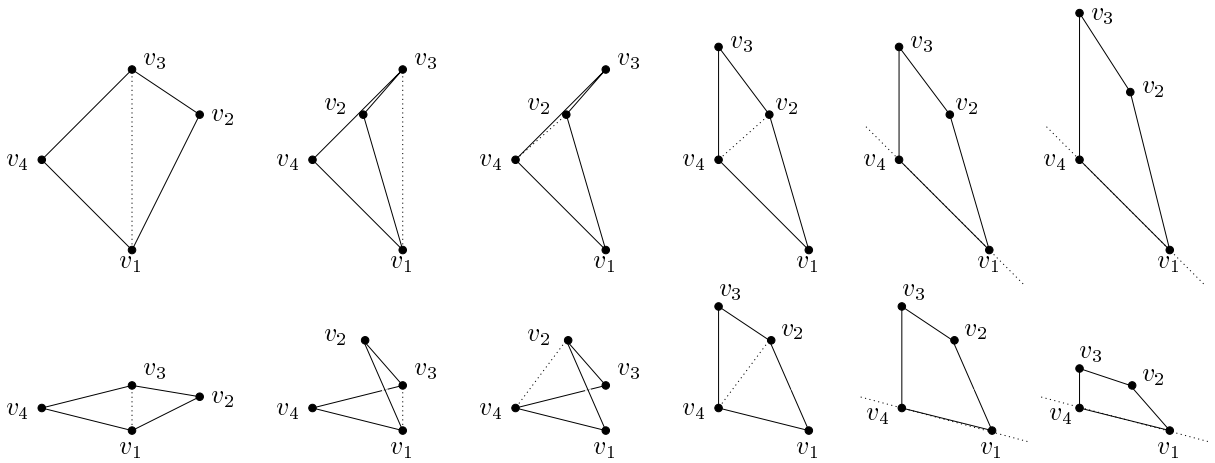

Figure 5: The same transformation as illustrated in Figure 2 but accomplished with three pivots, shown chronologically from left to right. (Top row) Bird's eye view. (Bottom row) Oblique view.

Another special type of pivot is a natural generalization of Erdős-Nagy flips. Let P be a polygon in R^d and let H be a hyperplane supporting the convex hull of P and containing at least two vertices of P. Reflect one of the resulting polygonal chains across H. Let us call such motions *hyperplane flips*. The first person to propose these hyperplane flips appears to be Gustave Choquet [7] in 1945, for applications to curve stretching. He claimed in [7] (but published no proof) that after a suitable choice of a *countable* number of hyperplane flips the polygons generated converge to planar convex polygons. These results were rediscovered in 1973 by Sallee [16].

In 1994 Millett [15], in connection with exploring varieties, proposed a "walk" algorithm consisting of a sequence of pivots to take any *equilateral* polygon (knot) in 3D into any other. (Millett allows self-crossings during the motions.) The interest in equilateral polygons comes from molecular biology where homogeneous macromolecules or polymers such as DNA are modeled by polygons with equal length edges. Here the vertices correspond to the monomers and the edges to the bonding force between them. To establish the walk Millett proposed taking an arbitrary equilateral polygon P in 3D to a planar regular polygon. His algorithm consists of three parts: (1) convert P to a *planar* star-shaped polygon P', (2) convert P' to a convex polygon P'' and (3) convert P'' to a regular polygon. However, his algorithm for part (1) does not always work correctly. His procedure may yield non-simple planar polygons in which all turns are right turns and the winding number is high, invalidating step (2) of the algorithm. Toussaint [20] proposed an alternate walk algorithm to convexify a 3D polygon that generalizes Millett's theorem to polygons in d dimensions with no restrictions on edge lengths.

Millett [15] showed in step (3) of his procedure that any convex planar polygon with equal edge lengths can be taken to any other via a bounded number (as a function of n) of pivots in 3D. In this section we demonstrate that this procedure also works for non-equilateral convex polygons, although in that case an unbounded number of pivots may be required.

We now prove the main theorem of this section.

Theorem 2 *There is a sequence of pivots moving between any two planar convex configurations of the same counterclockwise sequence of edge lengths.*

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 labelled $-,+,-,+$, respectively. This situation and the following motions are demonstrated in Figure 6. First, we pivot on $\overline{v_1v_3}$, rotating the entire subpolygon until v_2 is directly over v_1v_3. This brings us to the second quadrangle of Figure 6. Then we pivot on $\overline{v_2v_4}$ to achieve a planar quadrangle, as shown in the third and fourth pictures of the figure.

We will defer a discussion of the outcome of these pivots until later, and instead now demonstrate that no collisions can occur during these two pivots. During the first pivot, the polygon is folded along one crease like a taco up to an angle of $\pi/2$, so no collisions are possible there. During the second pivot, the subpolygon from v_1 to v_2 (which we will denote as $[v_1, v_2]$) and $[v_4, v_1]$ cannot collide, as these move in concert; likewise for $[v_2, v_3]$ and $[v_3, v_4]$. We can also readily see that the subpolygon $[v_4, v_1]$ remains motionless, and that $[v_3, v_4]$ rotates about $\overline{v_2v_4}$ by at most $\pi/2$ and therefore moves strictly upward. Therefore these two polygonal chains cannot intersect; by a symmetrical argument, neither can $[v_1, v_2]$ and $[v_2, v_3]$. The

only remaining possibility is the collision of $[v_2, v_3]$ and $[v_4, v_1]$ (and symmetrically $[v_3, v_4]$ and $[v_1, v_2]$). Note that after the first pivot the subpolygon $[v_2, v_3]$ points upward from $\overline{v_2v_3}$ in a vertical plane. After the second pivot, which is of at most $\pi/2$, this subpolygon must lie above a horizonal plane through $\overline{v_2v_3}$. Therefore $[v_2, v_3]$ and $[v_4, v_1]$ cannot collide, because the latter still lies below the former in the original plane of the polygon.

Following these first two pivots, the quadrangle, but not the polygon, is planar. We now perform a pivot on $\overline{v_4v_1}$ to bring the quadrangle back into its original plane, and then perform pivots on the remaining three edges of the quadrangle to bring the polygon into the same plane as the quadrilateral. These last motions are also illustrated in Figure 6.

We have shown that no collisions occur during these pivots, but it remains to be shown that any quadrangle desired can be achieved through the repetition of these motions. Consider again the first quadrangle of the top row of Figure 6. Let x be the closest point from v_2 on the line $\overline{v_1v_3}$. By the law of cosines, the distance between v_2 and v_4 is expressed 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 $|2(v_2x)(v_4x)\cos\angle v_2xv_4|$. Thus we always make considerable progress toward our goal configuration, unless our goal configuration is one where either v_2x, v_4x, or $\cos\angle v_2xv_4$ are zero. If the goal has both v_2x and v_4x as zero, then our goal configuration is self-intersecting and therefore invalid, so this case needs no consideration. In the other instances, this implies that in the goal, v_1, v_2 (or v_4, but this case is symmetric to v_2), and v_3 are collinear. Thus we can state that v_2 is the only vertex in between v_1 and v_3, else we violate convexity of the goal configuration. When v_1, v_2, and v_3 are almost collinear, a pivot about v_1v_3 of any angle (even as much as π) will not cause any self-intersections to arise. Therefore, if we pivot until v_2v_4 is the same distance as v_4x, and perform the remaining pivots to restore planarity of the polygon, $v_1v_2v_3$ will be collinear as desired. □

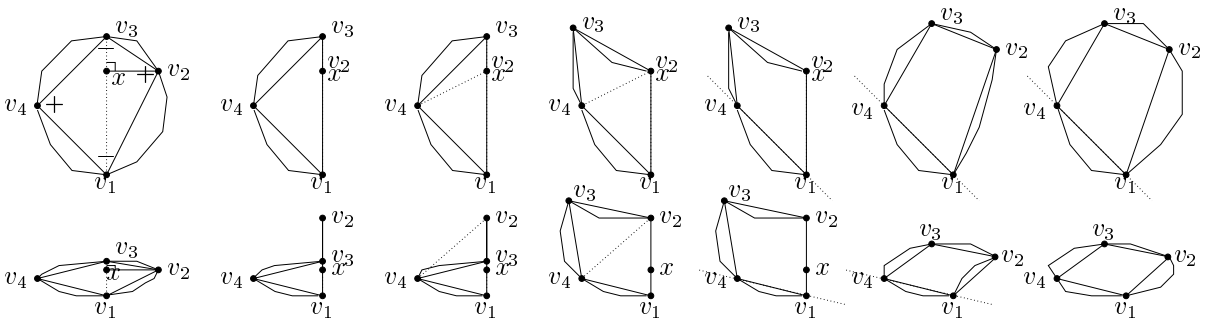

Figure 6: Illustration of the pivots used in Theorem 2. (Top row) Bird's eye view. (Bottom row) Oblique View.

The geometric progression of the above proof hints at the notion that there may exist some polygons for which the number of pivots required to move between any two arbitrary goal configurations may not be bounded by a function of the number of edges in the polygon. In fact, we will soon show this to be the case. Before proving this statement in Theorem 3,

we require the following lemma. We draw the reader's attention to Figure 7 which may serve as a useful visual aid during the course of the proof of Lemma 4.

Lemma 4 *Let $v_1v_2v_3v_4$ be a planar convex quadrangle. After two pivots, suppose the quadrangle is once again planar, resulting in a quadrangle $v_1''v_2''v_3''v_4''$. Then $\angle v_2''v_1''v_4''$ will be at least the original 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 will break the remaining possibilities into two cases; that where the pivot on $\overline{v_1v_3}$ is the first pivot (or both), and the case where it is preceded by a pivot on $\overline{v_2v_4}$.

If the pivot on $\overline{v_1v_3}$ occurs first, 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. Since $\angle v_4v_1v_3$ does not vary during the pivot, the resulting $\angle v_2'v_1'v_4'$ is 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 $\overline{v_1v_3}$ and must 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 the law of sines, $\angle v_2v_1v_4$ must have increased. □

The next theorem follows easily from Lemma 4.

Theorem 3 *There exist polygons which require arbitrarily many pivots to achieve a goal configuration.*

Proof: Examine the leftmost parallelogram in Figure 7. This polygon has the property that because it is a parallelogram, there exist configurations where it is as flat as desired; that is, where $\angle v_1$ is arbitrarily close to zero. Furthermore, because it is not a rhombus, $\angle v_2v_1v_3 \neq \angle v_4v_1v_3$. In fact, due to the law of sines,

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

Furthermore, for small angles $\angle x$, $\sin x \approx x$. Therefore as $\angle v_1$ gets smaller and smaller, for every two pivots $\angle v_1$ is only able to be reduced to $\angle v_1'$ according to the expression

$$\angle v_1' \leq |\frac{v_2v_3 - v_3v_4}{v_2v_3 + v_3v_4}|\angle v_1.$$

Because $\angle v_1$ approaches but cannot attain zero, we can choose a goal configuration with a small enough $\angle v_1$ as to require any number of pivots desired. (We note that although one cannot achieve a configuration where $\angle v_1 = 0$, this is not a valid configuration as the polygon would be flat and therefore self-intersecting.) While this proves the theorem for the case where every two pivots restores the polygon to a planar configuration, we have not directly proven the theorem for arbitrary pivots. However, 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. □

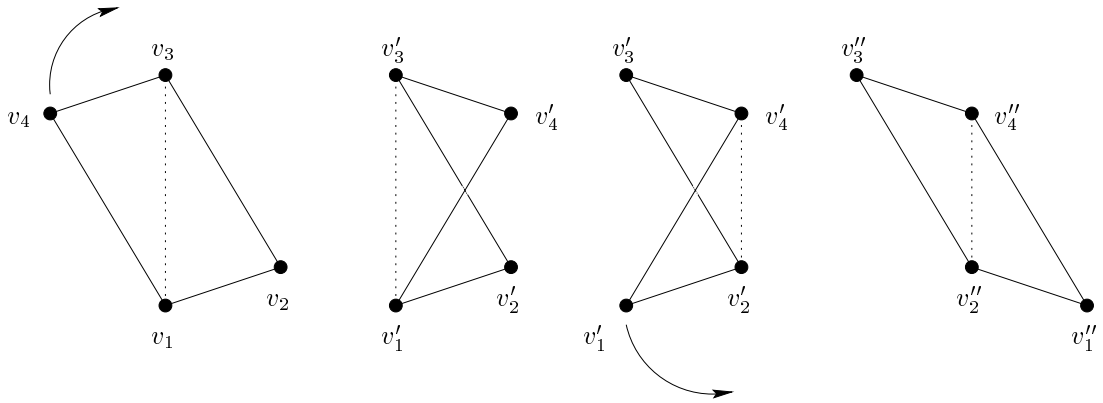

Figure 7: Two pivots performed to reconfigure a parallelogram that is not a rhombus.

Acknowledgments

The second author thanks Anna Lubiw for helpful discussions that inspired some of this research. This work was done at the Workshop on Computational Polygonal Entanglement Theory, organized by Godfried Toussaint on February 4–11, 2000 at the Bellairs Research Institute in Holetown, Barbados. We thank Vida Dujmović for important discussions at that meeting.

References

[1] Boris Aronov, Jacob E. Goodman, and Richard Pollack. Convexification of planar polygons in $\mathbb{R}^3$. Manuscript, October 1999. `http://www.math.nyu.edu/faculty/pollack/convexifyingapolygon10-27-99.ps%`.

[2] 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. Technical Report 060, Smith College, 1999. A preliminary version appeared in the *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms*, Baltimore, Maryland, January 1999, pages 866–867.

[3] Therese C. Biedl, Erik D. Demaine, Sylvain Lazard, Steven M. Robbins, and Michael A. Soss. Convexifying monotone polygons. In *Proceedings of the 10th Annual International Symposium on Algorithms and Computation*, volume 1741 of *Lecture Notes in Computer Science*, pages 415–424, Chennai, India, December 1999. An expanded version is Technical Report CS-99-03, Department of Computer Science, University of Waterloo, 1999.

[4] Jorge Alberto Calvo, Danny Krizanc, Pat Morin, Michael Soss, and Godfried Toussaint. Convexifying polygons with simple projections, 2001.

[5] 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.

[6] A. L. Cauchy. Deuxième mémoire sur les polygones et polyèdres (Second memoire on polygons and polyhedra). *Journal de l'École Polytechnique*, 9:87–98, May 1813.

[7] Gustave Choquet. Variétés et corps convexes. In *Association Française pour l'Avancement des Sciences*, pages 91–93, 1945. Congrès de la Victoire-1945.

[8] Roxana Cocan and Joseph O'Rourke. Polygonal chains cannot lock in 4D. In *Proceedings of the 11th Canadian Conference on Computational Geometry*, Vancouver, Canada, August 1999. `http://www.cs.ubc.ca/conferences/CCCG/elec_proc/c17.ps.gz`.

[9] Robert Connelly, Erik D. Demaine, and Günter Rote. Straightening polygonal arcs and convexifying polygonal cycles. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 432–442, Redondo Beach, California, November 2000.

[10] Peter R. Cromwell. Equality, rigidity, and flexibility. In *Polyhedra*, chapter 6, pages 219–247. Cambridge University Press, Cambridge, 1997.

[11] 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.

[12] W. J. Lenhart and S. H. Whitesides. Reconfiguring closed polygonal chains in Euclidean d-space. *Discrete and Computational Geometry*, 13:123–140, 1995.

[13] Neal Madras and Gordon Slade. *The Self-Avoiding Walk*. Birkhäuser, Boston, 1993.

[14] 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.

[15] Kenneth C. Millett. Knotting of regular polygons in 3-space. *Journal of Knot Theory and its Ramifications*, 3(3):263–278, 1994.

[16] G. T. Sallee. Stretching chords of space curves. *Geometriae Dedicata*, 2:311–315, 1973.

[17] I. J. Schoenberg and S. K. Zaremba. Cauchy's lemma concerning convex polygons. *Canadian Journal of Mathematics*, 19(4):1062–1071, 1967.

[18] Peter Schreiber. On the existence and constructability of inscribed polygons. *Beiträge zur Algebra und Geometrie (Contributions to Algebra and Geometry)*, 34(2):195–199, 1993.

[19] Ernst Steinitz and Hans Rademacher. *Vorlesungen über die Theorie der Polyeder*. Springer-Verlag, Berlin, 1934. Reprinted 1976.

[20] Godfried Toussaint. The Erdős-Nagy theorem and its ramifications. In *Proceedings of the 11th Canadian Conference on Computational Geometry*, Vancouver, Canada, August 1999. `http://www.cs.ubc.ca/conferences/CCCG/elec_proc/fp19.ps.gz`.

[21] 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.