The Erdős-Nagy Theorem and its Ramifications

Godfried Toussaint *
School of Computer Science
McGill University
Montréal, Québec, Canada

June 29, 1999

Abstract

Given a simple polygon in the plane, a *flip* is defined as follows: consider the convex hull of the polygon. If there are no pockets do not perform a flip. If there are pockets then reflect one pocket across its line of support of the polygon to obtain a new simple polygon. In 1934 Paul Erdős introduced the problem of repeatedly flipping all the pockets of a simple polygon simultaneously and he conjectured that the polygon would become convex after a finite number of flips. In 1939 Béla Nagy pointed out that flipping several pockets simultaneously may result in a non-simple polygon. Modifying the problem slightly he then proved that if at each step only one pocket is flipped the polygon will become convex after a finite number of flips. We call this result the Erdős-Nagy Theorem. Since then this theorem has been rediscovered many times in different contexts, apparently, with none of the authors aware of each other's work. One purpose of this paper is to bring to light this "hidden" work. We review the history of this problem, provide a simple elementary proof of a stronger version of the theorem and consider variants, generalizations and applications of interest in computational knot theory, polymer physics and molecular biology. We also improve several results in the literature with the application of the Erdős-Nagy theorem. We close with a list of open problems.

*This research was supported by NSERC and FCAR. e-mail: godfried@cs.mcgill.ca

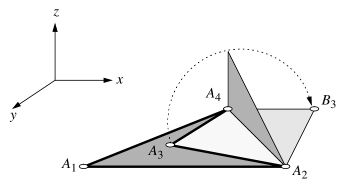

Figure 1: Rotating a reflex vertex using the third dimension.

1 Introduction

Let $A = A_1A_2A_3A_4$ be a nonconvex quadrilateral in the two-dimensional xy-plane with A_3 as its reflex vertex (refer to Figure 1). Furthermore, assume that the quadrilateral (although planar) is embedded in the 3D space with axes x, y and z, that the vertices are *ball-joints* which allow rotations in all directions in 3D. Finally, assume the links (edges) are rigid line segments with $A_1A_2 = A_1A_4$ and $A_2A_3 = A_3A_4$. If we lift vertex A_3 off the xy-plane into the third dimension z (leaving the other three vertices fixed) by rotating it about the line through A_2 and A_4 until it returns to the xy-plane at position B_3, then the quadrilateral has been convexified with one simple motion. This rotation motion in 3D is equivalent to a reflection transformation in the xy-plane: B_3 is the reflection of A_3 across the line through A_2 and A_4.

A generalization of this problem has been discovered and re-discovered independently by several

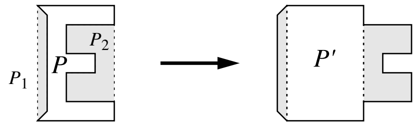

Figure 2: Flipping the pockets of a polygon.

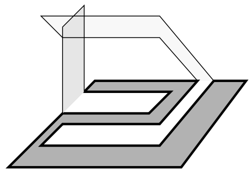

Figure 3: Flipping two pockets simultaneously may lead to a crossing polygon.

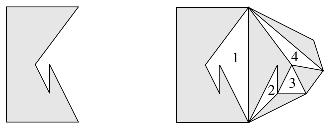

Figure 4: The polygon on the left is convexified after four flips.

mathematicians, biologists, physicists and computer scientists dating back to 1935. Computer scientists are motivated by practical robotics problems with linkages. Molecular biologists and polymer physicists are interested in unravelling large molecules (modeled as polygons) such as circular DNA [10]. Mathematicians are curious about the geometric properties of polygons and simple closed curves.

The first person to propose this problem appears to have been Paul Erdős in 1935 [8] in the context of planar polygons. Consider the simple polygon P in Figure 2 (a). If we subtract this polygon from its convex hull we obtain the convex deficiency: a collection of open connected regions. Each such region together with its boundary is itself a polygon, often called a *pocket* of P. The polygon P in Figure 2 (a) has two pockets P_1 and P_2. Each pocket has an edge which coincides with a convex hull edge of P (shown in the figure by dotted lines). Such an edge is called the pocket *lid*.

Erdős defined a reflection operation on P as a simultaneous reflection of all the pockets of P about their corresponding pocket lids. Applying a reflection operation to polygon P in Figure 2 (a) yields the new polygon P' in Figure 2 (b). In 1935 Erdős conjectured that given any simple polygon, a *finite* number of such reflection steps will convexify it. The first proof of Erdős' conjecture was provided in 1939 by Béla de Sz.-Nagy [6]. First Nagy observed that reflecting *all* the pockets in one step can lead from a simple polygon to a non-simple one. One such example due to Nagy is shown in Figure 3. Therefore he modified Erdős' problem slightly by defining one step to be the reflection of only *one* pocket. Since a pocket is reflected into a previously empty half-plane, no collisions can occur with such a motion. Let us call such an operation a *flip*.

Figure 4 shows a polygon being convexified after four flips. The pockets at each flip are shown in white before flipping and shaded after the flip is completed. Nagy then proceeded to prove that any simple polygon can be convexified by a *finite* number of flips.

2 Rediscoveries of the Erdős-Nagy Theorem

Branko Grünbaum [12] described some of the strange history of this problem and uncovered several rediscoveries of the theorem. He also provided his own version of a proof which is similar to Nagy's proof with one of the main differences being that at each step he flips the pocket that has maximum area (if more than one pocket exists). Since [12] is rather inaccessible, here we first briefly outline his findings and then add some more rediscoveries and variants to the history of this problem.

As mentioned previously, in 1939 Béla Nagy changed Erdős' problem slightly by reflecting only one pocket of the polygon at each step so that simplicity is maintained during the convexification process. As we shall see later, maintaining simplicity during the process is not necessary if the definition of a flip is suitably modified.

In 1957 there appeared two Russian papers by Reshetnyak [26] and Yusupov [38] proving the theorem with variants of basically the same proof.

In 1959 Kazarinoff and Bing [16] announced the problem with a solution. Two years later a proof appeared in a paper by Bing and Kazarinoff [3] and also in Kazarinoff's book [15]. They also conjectured that every simple polygon will be convex after at most $2n$ flips.

In 1973 two students of Grünbaum at the University of Washington, R. R. Joss and R. W. Shannon worked on this problem but did not publish their results. An account of the unfortunate circumstances surrounding this event is given by Grünbaum [12]. They found a counter-example to the conjecture of Bing and Kazarinoff (unaware of the conjecture of course). They showed that given any positive integer k, there exist simple polygons (indeed quadrilaterals suffice) that cannot be convexified with fewer than k flips.

In 1981 Kaluza [14] posed the problem again and asked if the number of flips could be bounded as a function of the number of vertices of the polygon.

In 1993 Bernd Wegner [36] took up Kaluza's challenge and solved both problems again. His proof of convexification in a finite number of flips is quite different from the others but his example for unboundednes is the same as that of Joss and Shannon.

In 1999 Biedl et al. [2] rediscovered the problem again and obtain the same results as Wegner. Their proofs of convexification are remarkably similar and their unboundedness example is the same quadrilateral.

3 A Proof of the Erdős-Nagy Theorem

Some of the published proofs of the Erdős-Nagy theorem are long and technical, others make references to higher mathematics, and some have gaps. We will also prove several theorems that make use of the Erdős-Nagy theorem as a lemma. Therefore, for both completeness and pedagogical reasons, it is appropriate to borrow the best features of the existing proofs, fill in the gaps, and present a simple, clear, elementary and short proof of the theorem. In this section we present such a proof along the lines of Nagy's reasoning, but first we consider a simple lemma for convex polygons that will be used in the proof. We assume that the convex polygon has no vertices with angle equal to π. If this is not the case it is a simple matter to scan the polygon in $O(n)$ time and delete vertices with angle equal to π by substituting longer adges appropriately.

Lemma 1 *Given a convex polygon, there exists a positive real number ϵ such that if some or all of the vertices are each moved by a distance less than ϵ, then the polygon remains convex.*

Proof: Consider vertex A_i and its two adjacent vertices A_{i-1} and A_{i+1} (refer to Figure 5). Let L_i be the line passing through the midpoints of the two edges A_iA_{i-1} and A_iA_{i+1}, and let r_i denote the minimum distance between A_i and L_i. Note that r_i is also the minimum distance between A_{i-1} and L_i as well as between A_{i+1} and L_i. Now construct disks D_i, D_{i-1} and D_{i+1}, all of the same radius r_i, centered at A_i, A_{i-1} and A_{i+1}, respectively. No matter where the vertices move, as long as each remains in the interior of its corresponding disk, their final positions B_i, B_{i-1} and B_{i+1} will have the property that B_i is separated by the line L_i from B_{i-1} and B_{i+1}. Therefore vertex B_i is convex. If we choose for the radius of our disk for every vertex the value $\epsilon = min\{r_1, r_2, ..., r_n\}$ then all vertices will remain convex and since the polygon is simple it follows from Proposition 5 in [11] that it is convex. ■

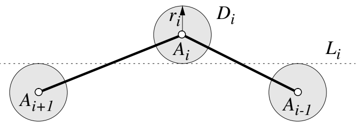

Figure 5: The convexity tolerance disk of a vertex.

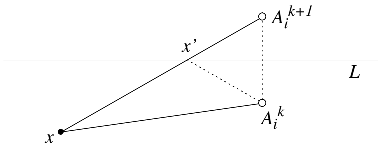

Figure 6: A flip increases the distance from a fixed point to a vertex.

This number ϵ is sometimes called the *convexity-tolerance* of the polygon [1]. It is a measure of how much the vertices of a convex polygon may be perturbed while guaranteeing that the polygon remains convex.

Theorem 1 *Every simple polygon can be convexified with a finite number of flips.*

Proof: Let $A^0 = A_1^0 A_2^0 ... A_n^0$ denote the given polygon before any flips have taken place. After performing k flips we obtain the polygon $A^k = A_1^k A_2^k ... A_n^k$ where vertex A_i^0 is taken to A_i^k for all $i = 1, 2, ..., n$. We will call polygon A^m a *descendant* of A^k if $m > k$. Consider any point x in A^0. Since for all k, A^{k+1} contains A^k, point x remains in all the descendants of A^0. We are interested in the distance betwen point x and a vertex of the k-th descendant of A^0, $d(x, A_i^k)$. After the next flip A_i^k either remains fixed or is reflected across a line of support L (refer to Figure 6). In the latter case this line is the perpendicular bisector of the segment $A_i^k A_i^{k+1}$. Let x' denote the intersection of line L with segment xA_i^{k+1}. Then we have

$$d(x, A_i^{k+1}) = d(x, x') + d(x', A_i^{k+1}).$$

Since x' is equidistant from A_i^k and A_i^{k+1} we obtain

$$d(x, A_i^{k+1}) = d(x, x') + d(x', A_i^k).$$

It follows from the triangle inequality that

$$d(x, A_i^{k+1}) \geq d(x, A_i^k).$$

Therefore the distance function $d(x, A_i^k)$ is a monotonically non-decreasing function of k. Furthermore, since the edges are rigid, the perimeter of every descendant of A^0 remains constant after every flip. Therefore the distance $d(x, A_i^k)$ is bounded from above by half the perimeter of A^0. From these two observations it follows that the sequence $\{A_i^0 A_i^1 A_i^2 ...\}$ has a limit. Let us denote the limit of A_i^k, as k goes to infinity, by A_i^* and let $A^* = A_1^* A_2^* ... A_n^*$ denote the limit polygon.

Firstly we remark that the limit polygon A^* must be a simple polygon. In other words, different vertices cannot converge to one and the same limit vertex. This follows from the observation above that $d(x, A_i^k)$ is a monotonically non-decreasing function of k, where the role of x is now played by another vertex A_j^k where $j \neq i$. If both A_j^k and A_i^k move with the next flip then

$$d(A_j^{k+1}, A_i^{k+1}) = d(A_j^k, A_i^k).$$

If only A_i^k moves, then

$$d(A_j^{k+1}, A_i^{k+1}) \geq d(A_j^k, A_i^k).$$

Therefore two vertices of A^k cannot move closer together when we flip A^k.

Secondly we note that the limit polygon A^* must be convex, for otherwise, being a simple polygon, another flip would alter its shape contradicting that it is the limit polygon.

Thirdly, some vertices of A^* will have interior angles equal to π and others less than π. Note also that whenever a vertex A_i^k becomes straight it remains straight for all descendants of A^k. Therefore we may ignore straight vertices in the analysis.

It remains to show that the sequence $\{A^0, A^1, ..., A^k\}$ where $A^k = A^*$ is finite. To this end let us now construct around each vertex

A_i^* whose interior angle is less than π a disk D_i of radius ϵ, the convexity tolerance of A^*. Consider the sequence of positions of the i-th vertex $\{A_i^0 A_i^1 A_i^2 ...\}$. Since A_i^m converges to A_i^* as m approaches infinity, there must exist a finite number c_i of flips after which $A_i^{c_i}$ first enters disk D_i. Furthermore, once it enters D_i it stays there. This follows from the fact that A^{c_i} is contained in A^* (A^* contains all previous polygons) and L_i separates $A_i^{c_i}$ from $A_{i+1}^{c_i}$ and $A_{i-1}^{c_i}$, thus preventing $A_i^{c_i}$ from being contained in the interior of any pockets of subsequent descendants. Therefore not only does A^{c_i} stay in D_i but it is in fact immobilized. If we let $c^* = max\{c_1, c_2, ..., c_n\}$, then after c^* flips every vertex has entered its corresponding limit disk and since the vertices must then remain in their respective disks it follows from Lemma 5 that A^{c^*} must be convex. Hence A^k is convex for $k = c^*$ flips. ■

To conclude this section we mention that since the convex hull of a simple polygon may be computed in $O(n)$ time [21], [23] it follows that each flip may be done in $O(n)$ time.

4 A Stronger Version of the Erdős-Nagy Theorem

Even in a non-simple (self-crossing) polygon a line of support of the polygon may contain two vertices, say A and B which divide the polygon into two chains connecting A and B. One may wonder if repeatedly flipping one of these chains across the line of support will also convexify a non-simple polygon in a finite number of flips. This is indeed the case and was recently proved by Grünbaum and Zaks [13]. All the published proofs of the Erdős-Nagy theorem (for simple polygons) are based on increasing area in that they depend on the fact that after every flip the new polygon contains the previous one in its interior. For crossing polygons the notions of increasing area and interior lose their meaning. Therefore Grünbaum and Zaks [13] use a different approach in their proof for crossing polygons. At each step they select from all the possible candidate flips determined by lines of support, the flip that maximizes the sum of the distances between all pairs of vertices of the polygon. This function has the desired property that it strictly increases after each flip. Unfortunately it requires $O(n^2)$ time to compute for each flip. We will show that there is a simple proof of a stronger version of this theorem that follows directly from the Erdős-Nagy theorem and that leads to a simpler algorithm in which each flip may be computed in $O(n)$ time. First we show for the case of simple polygons that the Erdős-Nagy theorem can be strengthened by requiring that during the entire convexification procedure a specified edge of the polygon remain fixed. This result will then be used to prove the theorem for crossing polygons.

4.1 Mirror-Flips

To strengthen the Erdős-Nagy theorem for simple polygons in the plane we introduce a new reflection operation we call a *mirror-flip*. Consider the polygon P in Figure 7 (a) and let L be a line of support of P that contains the non-adjacent vertices A and B. These vertices divide the polygon into two chains: the *inner* and the *outer* chains. Each of these two chains together with segment AB define a new polygon. The inner chain is the chain that defines the polygon contained in the polygon determined by the other (outer) chain. The standard flip employed in the Erdős-Nagy theorem flips the inner chain about L as illustrated in Figure 7 (b) to obtain the shaded polygon P'. On the other hand, the mirror-flip operation reflects the outer chain as illustrated in Figure 7 (c) to obtain the shaded polygon P''. We now prove a stronger version of the Erdős-Nagy theorem for simple polygons.

Theorem 2 *Every simple polygon can be convexified with a finite number of* flips *or* mirror-flips *while keeping a specified edge fixed.*

Proof: First observe that if instead of performing a flip we do a corresponding mirror-flip the two resulting polygons are mirror images of each other. It follows from the Erdős-Nagy theorem that if at each step a flip is to be performed we choose at random to either perform a flip or its corresponding

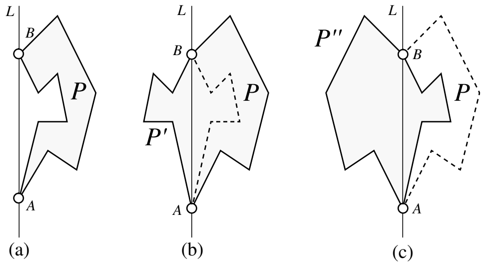

Figure 7: A simple polygon (left), a *flip* (center) and its corresponding *mirror-flip* on the right.

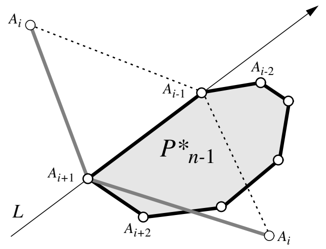

Figure 8: Illustrating the induction proof for flipping a crossing polygon.

mirror-flip (say by the flip of a coin), then such a procedure must result in a convex polygon after a finite number of flips and mirror-flips. Secondly, instead of using a coin to decide whether to flip or mirror-flip, we could just as well decide by requiring a specified edge of the polygon to remain fixed in the plane during the entire process. If at any step the specified edge is part of the flipping (inner) chain we perform the mirror-flip (outer chain) instead and vice-versa. Thus this procedure will convexify the polygon with the additional constraint that a preselected edge remain fixed in the plane at all times. Finally, note that this procedure may yield a convex polygon with the opposite orientation to that of the initially given polygon. If this is the case it suffices to flip the convex polygon once more about the line containing the fixed edge. ■

4.2 Crossing Polygons

We will now use Theorem 2 to prove that a self-crossing polygon P_n with vertices $A_1, A_2, ..., A_n$ may be convexified with a finite number of flips and mirror-flips while maintaining a specified edge of P fixed in the plane, thereby strengthening the result of Grünbaum and Zaks [13]. For simplicity of exposition we assume the polygon is in general position in the sense that its vertices are distinct and no three of them lie on a line.

Theorem 3 *Every crossing polygon can be convexified with a finite number of* flips *and* mirror-flips, *while keeping a specified edge fixed, in $O(n)$ time per flip and mirror-flip.*

Proof: (by induction on n) We begin by observing that the result is trivial if $n = 3$. For $n > 3$ assume that the assertion is already known for polygons with $n-1$ vertices. We want to prove the result for polygons of n vertices. Let P_n denote the given polygon. Let us assume that edge A_iA_{i+1} has been chosen to remain fixed during convexification. Replace vertex A_i and its incident edges in P_n with an edge joining A_{i-1} and A_{i+1}, resulting in a polygon P'_{n-1} of $n-1$ edges. By the induction hypothesis P'_{n-1} may be convexified leaving any edge fixed. Therefore let us choose edge $A_{i-1}A_{i+1}$ as the fixed edge and let P^*_{n-1} denote the convexified version of P'_{n-1} (refer to Figure 8). Now delete edge $A_{i-1}A_{i+1}$ from P^*_{n-1} and replace it with A_i and its incident edges in P_n to obtain polygon P^*_n.

It remains to show that P^*_n can be convexified with a finite number of flips and mirror-flips leaving A_iA_{i+1} fixed. Consider where A_i may lie with respect to line L, the line that contains A_{i-1} and A_{i+1}. Note that A_i cannot lie on line L for it would mean A_i, A_{i-1} and A_{i+1} are collinear contradicting our assumption of non-collinearity. Without loss of generality let the line L be directed from A_{i+1} to A_{i-1} and assume P^*_{n-1} lies to the right of L. We have two cases.

Case 1: If A_i lies to the left of L then P_n^* is a simple polygon and the result follows from Theorem 2.

Case 2: If A_i lies to the right of L then we can reflect the chain connecting A_{i+1} to A_{i-1}, namely $A_{i+1}, A_{i+2}, ..., A_{i-2}, A_{i-1}$, across line L to reduce it to case 1. This completes the correctness part of the proof.

Let us turn to the complexity. The induction proof suggests the following algorithm. Let $P_n = A_1, A_2, ..., A_n$ denote the given polygon and without loss of generality assume that edge A_1A_2 is selected to remain fixed. Initially select the first three vertices $A_1A_2A_3$ as the convex polygon with edge A_1A_2 fixed. Advance one vertex on P_n to A_4 and consider the quadrilateral determined by these four vertices and the "phantom" edge (diagonal) A_1A_4. Next convexify this quadrilateral using the rules specified above to ensure edge A_1A_2 remains fixed. If triangle $A_3A_4A_1$ is to be flipped about the line containing A_1 and A_3 then flip the *entire* chain $A_3A_4, ..., A_1$. Proceed in this fashion until A_n is reached.

Note that this convexification procedure may yield a convex polygon with orientation opposite to that of the original given polygon P_n. If a convexification with the same orientation is desired simply flip the chain connecting the fixed edge A_1A_2 about the line containing this edge. Each flip or mirror flip only requires precomputing the convex hull of a simple polygon (for flips only) and recomputing the coordinates of the polygon's vertices. Therefore $O(n)$ time suffices per flip or mirror-flip. ∎

Note that the proof of Theorem 3 carries through even if A_i lies on L. If it lies on the segment $A_{i+1}A_{i-1}$ the resulting polygon is convex and there is nothing to do. If it lies on L and, say, above A_{i-1} then the line of support L' of A_i and some other vertex A_k determines a subchain that can be reflected across L' to yield the desired simple polygon. Therefore, this approach applies to more general polygons where collinearities of vertices are allowed. Indeed, Grünbaum and Zaks [13] prove convexification for a general class of non-simple polygons called *exposed* polygons.

To conclude, our proof has two features not present in [13]. The first is that we can "freeze" a given edge of the polygon during convexification thus preventing the unfolding from "running away". The second is the reduced computational complexity carried out per flip. The algorithm in [13] uses an $O(n^2)$ test to decide which flip to perform, namely the computation of the sum of the distances between all the pairs of vertices of the polygon. Furthermore, before this distance computation is performed the convex hull of the crossing polygon is needed, adding another $O(n \log n)$ time per flip [25]. For the algorithm given here $O(1)$ time suffices to decide which flip is to be performed after the convex hull is computed. A mirror flip only involves recomputation of the coordinates of the vertices and takes $O(n)$ time. Furthermore, a convexification flip involves the computation of the convex hull of a *simple* polygon and is therefore only $O(n)$ time per flip [21], [23].

5 Variants and Generalizations

5.1 Mouth Flips

Knot theorists are interested in polygons in 3D (knots). In particular, for the computer analysis of knot spaces (or exploring the respective variety) they are interested in "walk" algorithms that will take one knot into another. Millett [24] rediscovered a special case of the Erdős-Nagy theorem when the polygons satisfy all of the following properties: (1) they are star-shaped, (2) they are equilateral (all edges have equal length) and (3) a flip is made not on a complete pocket of the polygon but only on a reflex vertex reflected across the line joining its adjacent vertices. We will call such a flip a *mouth-flip*. Millet proves that ultimately enough mouth-flips convexify the polygon. However, one can prove with an argument similar to that in [12] that not only will the polygon be convexified after a finite number of mouth-flips but this number can be bounded as a function of n because the polygon is equilateral. To see this note that the *before-after* positions of a mouth form a parallelogram. Therefore no new slopes (aside from the slopes of the edges of the original polygon) are ever introduced by mouth-flipping. But the area strictly increases after each mouth-flip.

Therefore each new polygon generated on the path towards convexity is composed of a new permutation of the edges (no permutation is revisited during this walk). Therefore the number of mouth-flips is bounded by the number of permutations. We therefore have the following theorem.

Theorem 4 *A star-shaped equilateral polygon with n vertices can be convexified with at most $(n-1)!$ mouth-flips.*

5.2 Pivots and Hyperplane Flips

One way to generalize the original Erdős-Nagy flip in the plane is to consider any two vertices of the polygon and to reflect one of the polygonal chains they determine across the line they define. An additional generalization is obtained if the selected chain is not reflected but rotated (about the line as axis) by some angle (assuming the polygon is embedded in 3D). Finally, a third further generalization is to polygons in d dimensions. Combining all three ideas leads to a motion which in knot theory and physics is called a *pivot* [24], [19], [20], [35]. Erdős-Nagy flips may be considered as special cases of pivots with planar polygons in 3D where the pairs of vertices that define the pivots are determined by the lines of support of the polygon that detrmine pockets and each rotation has an angle of π. For the results obtained in this paper that refer to pivots, a pivot will mean a *rotation* pivot. Two vertices of the polygon partition the polygon into two chains. A rotation pivot rotates one of these two chains about the line containing the two vertices as the axis of rotation.

Another special type of pivot which is a natural generalization of Erdős-Nagy flips is as follows. 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 [5] in 1945 for applications to curve stretching, a topic to be discussed below. He claimed in [5] (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 [29].

In 1994 Millett [24], in connection with exploring varieties, proposed a "walk" algorithm (consisting mainly of a sequence of pivots) to take any equilateral polygon (knot) in 3D into any other. The interest in equilateral polygons comes from molecular biology where homogeneous macromolecules or polymers such as DNA are modelled by polygons with equal length edges. Here the vertices correspond to the mers and the edges to the bonding force between them. To establish the walk Millet proposed taking an arbitrary 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. Part (2) is done using the mouth-flips discussed above on the reflex vertices of P'. 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 thus invalidating step (2) of the algorithm (no reflex vertices are found). However, we can obtain a walk algorithm by modifying (1) and applying the Erdős-Nagy theorem for (2). Furthermore, this modification generalizes Millet's theorem to polygons in d dimensions with no restrictions on edge lengths. Assume P_n is a polygon of n vertices in R^d such that all vertices are distinct and no three are collinear. Consider the first four vertices of P_n. They determine a possibly skew quadrilateral. Rotate one of the triangles (if necessary) so that the quadrilateral is planar (one pivot). If the quadrilateral is not convex apply Erdős flips (pivots) to it until it is convex. Note that some of these pivots may carry the remaining polygon with them as in the case of crossing planar polygons. Now advance to the next vertex of P_n, pivot this triangle so it is co-planar with the convex quadrilateral and again apply flips to the pentagon if it is not convex. Continuing this process leads to convexification with pivots only. Furthermore, if we desire to keep one segment fixed in space at all times we can apply Theorem 3 to incorporate mirror-flips when necessary on the planar portion of the polygon. Therefore we have the following result.

Theorem 5 *In dimensions higher than two any polygon can be convexified with a finite number of pivots, while keeping a specified edge fixed in $O(n)$ time per pivot.*

Of course it follows from our previous discussion that the number of pivots in Theorem 5 cannot be bounded as a function of n. However, convexification is possible in a polynomial number of moves if we are willing to use more complicated motions. For example in 1995 Lenhart and Whitesides [17] showed that (in any dimension greater than two) a polygon may be convexified in $O(n)$ time with $O(n)$ *5-joint* line-tracking motions. Each such motion rotates five joints with two cooperating "elbows". In 1973 Sallee [29] proved that this can be accomplished with pivots and *4-joint* line-tracking motions. He gives no complexity analysis in [29] but examination of his algorithm reveals that it can be accomplished in $O(n)$ time with $O(n)$ such motions.

5.3 Curve Inflation

A generalized version of Erdős' problem for the case of arbitrary simple curves has also been discovered independently. In this context the operation is referred to as *inflation*. Flipping several arcs simultaneously as originally proposed by Erdős is called *full inflation* and flipping only one arc is called *partial inflation*. For sufficiently smooth curves Robertson [27] proves that they converge to a convex curve after a suitable infinite sequence of flips. Robertson and Wegner [28] investigate the degree of smoothness of the limit curves obtained by flipping.

5.4 Stretching

Let $A = A_1, A_2, ..., A_n$ be a polygon that is reconfigured to another $B = B_1, B_2, ..., B_n$. In other words the corresponding line segments have the same length and to each point on A there corresponds a point on B in the obvious way. If for every two points on A their corresponding points on B are further (or the same distance) apart then we say that B is a *chord-stretched* version of A. In 1973 Sallee [29] proved that for every polygon in d dimensions there exists a planar convex stretched version. Furthermore he gives an algorithm for carrying out the reconfiguration. Therefore these are stronger results than the convexification results mentioned earlier. The same results were aparently obtained as early as 1945 by Choquet [5]. Strantzen and Brooks [32] prove a conjecture of Yang Lu that if B is a chord-stretched version of A and if A is convex, then A and B are congruent.

5.5 Flipturns

Consider a planar polygon with a pocket determined by vertices A_i and A_j and refer to Figure 9. Another generalization of the Erdős flip was considered in 1973 by Joss and Shannon [12] where instead of flipping the pocket we rotate it by 180 degrees about the center of the convex hull edge that determines the pocket. The effect of this kind of flip which they called a *flipturn* is that no new slopes are introduced after a flipturn. What was automatically obtained in the case of mouthflips for star-shaped equilateral polygons is obtained here for any simple polygon by flipping and "turning" in a vertical plane. Joss and Shannon proved that any simple polygon with n sides can be convexified by a sequence of at most $(n-1)!$ flipturns. This bound is very loose and they conjectured that $(n^2)/4$ flipturns are always sufficient. Grünbaum and Zaks [13] showed that even crossing polygons could be convexified with a finite number of flipturns. In 1999 Therese Biedl discovered a polygon such that a bad sequence of flipturns leads to convexification only after $\Theta(n^2)$ flipturns.

A related "cutting" operation is used in physics for self-avoiding walks where the polygonal chain connecting the two vertices in question is just inverted with respect to these vertices. Such a "pivot" is called a *diagonal reflection* [20]. Even more relevant is the work of Dubins et al., [7] on planar simple polygons in Z^2, the square lattice. Physicists call the flipturn an *inversion*. In [7] it is shown that any simple lattice polygon of n vertices may be convexified with at most $n-4$ flipturns. They are not concerned with computational complexity but clearly each flipturn can be done in $O(n)$ time with any of several convex hull algorithms [21], [23]. Therefore we can state the

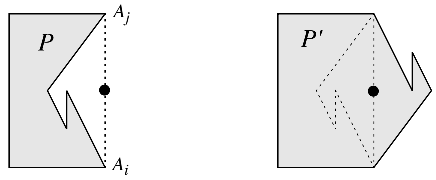

Figure 9: The polygon P' results from performing a flipturn on polyogn P.

following theorem.

Theorem 6 *Any simple n-vertex polygon in Z^2 can be convexified in $O(n^2)$ time with at most $n-4$ flipturns.*

5.6 Non-Crossing Linkages

None of the work discussed above, apart from the original Erdős-Nagy theorem, is concerned with whether or not edges of the polygon cross each other during reconfiguration. However, in some applications such as linkage analysis in robotics, exploring varieties in knot theory and molecular biology problems, the edges are to be considered as physical barriers so that no crossings are allowed. Biedl et al., [2] explore the area of convexifying polygons under these constraints. Among other things they show that a planar simple polygon in 3D may be convexified in $O(n^2)$ time with $O(n)$ pivots and $O(n^2)$ 4-joint line-tracking motions. Jeffrey Erickson pointed out that with a slight modification and an amortized complexity analysis the algorithm in [2] runs in $O(n)$ time. A survey of this area can be found in [33]. In three dimensions unknotted polygons that cannot be convexified have been discovered independently by Biedl et al., [2] (with ten edges) and Cantarella and Johnston [4] (with six edges). Toussaint [34] discovered an additional class of stuck unknotted hexagons.

5.7 Self-Avoiding Walks

To simplify Monte-Carlo simulations most work related to the problems discussed in this paper that is done in physics is restricted to Z^2 and Z^3, i.e., square and cubic lattices in two and three dimensions, respectively. Almost no results are available for the continuum (also off-lattice) model. One notable exception is the work of Stellman and Gans [31] which concerns open polygonal chains in 3D and considers a motion they call a *dihedral rotation* which selects a randon edge of the chain and rotates the smaller of the two chains incident to that edge, about the line through that edge as an axis. Some work has also been done on the FCC lattice [35]. Furthermore, in Z^2 and Z^3 not only must vertices be situated on lattice points but the edges are parallel to the coordinate axes and their lengths are all equal. Like the robotics research on linkages, the problems of interest to physicists involve closed simple polygons [7], open simple polygonal chains [22] and simple polygonal trees [9], i.e., polygons, chains and trees that do not intersect themselves; hence the term *self-avoiding walks* for the case of polygons and chains. Generating a random walk that does not self-intersect, especially if it must return to its starting point as in the case of polygons, is difficult (the waiting time is too long due to attrition, i.e., if a random walk crosses itself at any point other than its starting point it must be discarded and a new walk started). Therefore an efficient method frequently used to generate the chains or polygons is to modify one such object into another by means of a pivot for various definitions of pivot. Unlike the work in linkages however, here we do not care if intersections happen *during* the pivot as long as when the pivot is complete we end up with a simple polygon or chain. In other words, simplicity is required only at certain "snapshots" during the process. In general the pivots used are selected from a variety of transformations such as reflections and rotations of the sub-chain in question. Such transformations even include "cut-and-paste" operations. The reader is referred to a multitude of such problems and results contained in [19]. For example, Madras and Sokal [20] have shown that in Z^d for $d \geq 2$ every simple lattice polygonal chain of

n edges can be straightened by some sequence of at most $2n-1$ suitable pivots while maintaining simplicity *after* each pivot. The pivots used here are either reflections through coordinate hyperplanes or rotations by $\pm\pi/2$.

In order to prove the ergodicity of self-avoiding walks, polymer physicists are also interested in convexifying (and straightening) open polygonal chains under various geometric constraints [37], [30]. Such constraints, which include remaining in between two parallel lines or fixing the two endpoints of the chain, have application to polymer adsorption, steric stabilization of colloids and surface magnetism [37]. Madras, Orlitsky and Shepp [18] showed that any lattice polygon in Z^d, with endpoints fixed, can be convexified with $O(n^{6(d-2)})$ *generalized* pivots. As a corollary of Theorem 5 we obtain a continuum version of this fixed endpoint theorem for the lattice. Let $C_n = A_1, A_2, ..., A_n$ be an n-vertex open polygonal chain in R^d which we want to convexify while holding A_1 and A_n fixed. Here convexification of the chain means that inserting edge A_1A_n makes the chain a convex polygon. The fixed endpoint restriction on an open chain is equivalent to having an edge between the endpoints. Therefore we may consider C_n to be a closed, possibly self-crossing, polygon P_n which is to be convexified while keeping edge A_1A_n fixed. Thus Theorem 5 immediately implies the following result.

Corollary 1 *In dimensions higher than two any open polygonal chain can be convexified with a finite number of pivots while holding its endpoints fixed in* $O(n)$ *time per pivot.*

6 Conclusion and Open Problems

We conclude by mentioning several open problems in this area.

1. Wegner [36] proposed a very interesting variant of Erdős flips which can be considered the inverse problem which he called *deflation*. Given a simple polygon P in the plane, if there exists a pair of non-adjacent vertices A_i and A_j such that the line through A_i and A_j is not a line of support of P, the line intersects the boundary of the polygon only at A_i and A_j, and the polygonal chain $A_i, A_{i+1}, ..., A_j$ can be reflected about this line to lie inside the polygon then this reflection operation is called a *deflation*. If this cannot be done the polygon is called *deflated*. Wegner conjectured that every simple polygon can be deflated with a finite number of deflations.

2. Wegner also introduced two measures of convexity for simple polygons that are functions of the number of flips that will convexify the polygon. He called these the *maximal* and *minimal* inflation complexities. The former is the maximum number of flips that will convexify a polygon. The latter is the minimum number of flips. There are polygons (quadrilaterals) for which these two numbers are the same. What is the computational complexity of computing these numbers?

3. The Joss-Shannon conjecture that every simple polygon can be convexified with at most $(n^2)/4$ flipturns is still open. In fact, no upper bound lower than $(n-1)!$ is known!

4. A planar lattice simple polygon on the other hand can be convexified in $O(n^2)$ time with $n-4$ flipturns. Can this complexity be reduced?

5. The results concerning stuck unknotted hexagons in [4] and [34] show that there exist at least five classes of nontrivial embeddings of the hexagon in 3D. It is conjectured that there are no more than five such classes.

7 Acknowledgements

I would like to thank Jorge Calvo, Erik Demaine and Branko Grünbaum for reading the manuscript and providing constructive comments.

References

[1] Manuel Abellanas, Ferran Hurtado, and Pedro Ramos. Tolerance of geometric structures. In *Proc. 6th Canadian Conference on Computational Geometry*, pages 250–255, 1994.

[2] T. Biedl, E. Demaine, M. Demaine, S. Lazard, A. Lubiw, J. O'Rourke, M. Overmars, S. Rob-

bins, I. Streinu, G. T. Toussaint, and S. Whitesides. Locked and unlocked polygonal chains in 3d. In *Proc. 10th ACM-SIAM Symposium on Discrete Algorithms*, pages 866–867, 1999.

[3] 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. In Russian.

[4] 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.

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

[6] Béla de Sz.-Nagy. Solution of problem 3763. *American Mathematical Monthly*, 46:176–177, 1939.

[7] L. E. Dubins, A. Orlitsky, J. A. Reeds, and L. A. Shepp. Self-avoiding random loops. *IEEE Transactions on Information Theory*, 34:1509–1516, 1988.

[8] Paul Erdős. Problem number 3763. *American Mathematical Monthly*, 42:627, 1935.

[9] Paul W. Finn, Dan Halperin, Lydia E. Kavraki, Jean-Claude Latombe, Rajeev Motwani, Christian Shelton, and Suresh Venkatasubramanian. Geometric manipulation of flexible ligands. In *Applied Computational Geometry*, pages 67–78. Springer-Verlag, 1996.

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

[11] Branko Grünbaum. Polygons. In *Lecture Notes in Mathematics*, pages 147–184. Springer-Verlag, 1975.

[12] Branko Grünbaum. How to convexify a polygon. *Geombinatorics*, 5:24–30, 1995.

[13] Branko Grünbaum and Joseph Zaks. Convexification of polygons by flips and by flipturns. Technical Report 6/4/98, Department of Mathematics, University of Washington, Seattle, 1998.

[14] T. Kaluza. Problem 2: Konvexieren von Polygonen. *Math. Semesterber.*, 28:153–154, 1981.

[15] N. D. Kazarinoff. *Analytic Inequalities.* Holt, Rinehart and Winston, 1961.

[16] N. D. Kazarinoff and R. H. Bing. A finite number of reflections render a nonconvex plane polygon convex. *Notices of the American Mathematical Society*, 6:834, 1959.

[17] William J. Lenhart and Sue H. Whitesides. Reconfiguring closed polygonal chains in euclidean d-space. *Discrete and Computational Geometry*, 13:123–140, 1995.

[18] 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.

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

[20] 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.

[21] D. McCallum and D. Avis. A linear algorithm for finding the convex hull of a simple polygon. *Information Processing Letters*, 9:201–206, 1979.

[22] Frank M. McMillan. *The Chain Straighteners.* The MacMillan Press, 1979.

[23] A. Melkman. On-line construction of the convex hull of a simple polyline. *Inform. Process. Lett.*, 25:11–12, 1987.

[24] K. Millett. Knotting of regular polygons in 3-space. *Journal of Knot Theory and its Ramifications*, 3:263–278, 1994.

[25] J. O'Rourke. *Computational Geometry in C (Second Edition)*. Cambridge University Press, 1998.

[26] Yu. G. Reshetnyak. On a method of transforming a nonconvex polygonal line into a convex one. *Uspehi Mat. Nauk*, 12(3):189–191, 1957. In Russian.

[27] S. A. Robertson. Inflation of plane curves. In *Geometry and Topology of Submanifolds-III*, pages 264–275. World Scientific, 1991.

[28] S. A. Robertson and B. Wegner. Full and partial inflation of plane curves. In *Intuitive Geometry, Colloquia Math. Soc. Janos Bolyai*, pages 389–401. North-Holland, Amsterdam, 1994.

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

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

[31] S. D. Stellman and P. J. Gans. Efficient computer simulation of polymer conformation. I. Geometric properties of the hard-sphere model. *Macromolecules*, 5:516–526, 1972.

[32] J. Strantzen and J. Brooks. A chord stretching map of a convex loop is an isometry. *Geometriae Dedicata*, 41:51–62, 1992.

[33] G. T. Toussaint. Computational polygonal entanglement theory. In *VIII Encuentros de Geometría Computacional*, July 7-9 1999. Castellon, Spain.

[34] G. T. Toussaint. A new class of stuck unknots in Pol_6. Technical Report SOCS-99.1, School of Computer Science, McGill University, April 1999.

[35] 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.

[36] Bernd Wegner. Partial inflation of closed polygons in the plane. *Contributions to Algebra and Geometry*, 34(1):77–85, 1993.

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

[38] A. Ya. Yusupov. A property of simply-connected nonconvex polygons. *Uchen. Zapiski Buharsk. Gos. Pedagog.*, pages 101–103, 1957. In Russian.