DISCRETE MATHEMATICS

Discrete Mathematics 241 (2001) 333–342

www.elsevier.com/locate/disc

Convexification of polygons by flips and by flipturns

Branko Grünbaum[a,*], Joseph Zaks[b]

[a] *Department of Mathematics, University of Washington, Seattle, WA 98195-4350, USA*

[b] *University of Haifa, Haifa, Israel 31905*

Dedicated to Helge Tverberg on the occasion of his 65th birthday

Abstract

Simple polygons can be made convex by a finite number of flips, or of flipturns. These results are extended to very general polygons. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Polygon; Convexification; Flip; Flipturn

1. Introduction

Let P be a simple polygon in the plane. For a pair A, B of non-adjacent vertices of P let P^* and P^{**} be the two paths from A to B in P. Non-adjacent vertices A and B are an *exposed pair* of vertices provided that a support line L of the convex hull $\mathrm{conv}(P)$ of P contains A and B, but neither P^* nor P^{**} is contained in L. The *flip image* $f(P; A, B)$ of P with respect to the exposed pair A and B is the polygon $P^* \cup r_L(P^{**})$, where r_L denotes the *flip map* (reflection in the line L); see Fig. 1. Similarly, the *flipturn image* $g(P; A, B)$ of P with respect to A and B is the polygon $P^* \cup h_{AB}(P^{**})$, where h_{AB} denotes the *flipturn map* (halfturn about the midpoint M of the segment $[A, B]$); see Fig. 2. We note that for both flip map and flipturn map, if the roles of P^* and P^{**} were reversed, a polygon congruent to $f(P; A, B)$ or $g(P; A, B)$ would result; hence we can choose the path to be flipped or flipturned as convenient, without affecting the final outcome of the constructions discussed.

The following result was first established by Sz.-Nagy [6] in 1939, as a solution to a modification of a problem posed by Erdős [3] in 1935; see comment (4) in Section 4. Concerning later proofs and developments, and the somewhat chaotic history of the result, see [5].

* Corresponding author.

E-mail addresses: grunbaum@math.washington.edu (B. Grünbaum), jzaks@mathcs2.haifa.ac.il (J. Zaks).

PII: S0012-365X(01)00152-2

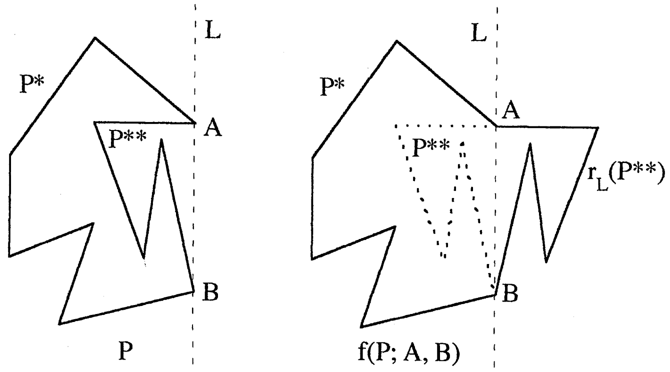

Fig. 1. An illustration of the flip operation.

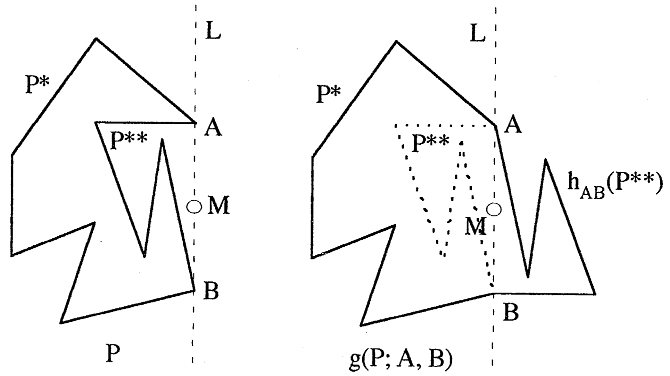

Fig. 2. An illustration of the flipturn operation.

Theorem 1 (Sz.-Nagy [6]). *Every simple polygon in the plane can be transformed into a convex polygon by a finite sequence of flips determined at each step by an exposed pair of vertices.*

An analogous result for flipturns was established in the early 1970s by R. R. Joss and R. W. Shannon, who at that time were graduate students at the University of Washington. Their result was published only in [5], where there is also an account of the unfortunate circumstances that led to the delay in publication. The result is:

Theorem 2 (Joss and Shannon). *Every simple polygon in the plane can be transformed into a convex polygon by a finite sequence of flipturns determined at each step by an exposed pair of vertices. Moreover, if the polygon has n sides, the sequence needs at most $(n-1)!$ flipturns.*

This is in contrast to the situation in Theorem 1, where—as shown by Joss and Shannon—there is no fixed bound on the number of flips needed even in case of quadrangles. It should be noted that in these results "convex polygon" has to be understood in a slightly wider sense than usual, by allowing adjacent edges to be collinear; in Section 2, we shall call such polygons "weakly convex". Polygons P which are convex in the usual sense, that is, for which the only vertices are the extreme points of their convex hull conv(P), will here be called "strictly convex".

Our main aim is to extend the above two theorems to more general plane polygons. It will turn out that in such a more general setting Erdös's original problem has an affirmative solution.

In the next section, we shall give the definitions necessary for the formulation of our results; the proofs will be given in Section 3, while the last section will present various comments and some open problems.

2. Definitions and results

An n-gon $P=[V_1,V_2,\ldots,V_n]$, where $n\geqslant 2$, is a sequence of points V_i (the *vertices* of P) in a plane, and closed segments $[V_i,V_{i+1}]$, for $i=1,2,\ldots,n$, where $V_{n+1}=V_1$ and $V_i\neq V_{i+1}$ for all i (the *edges* of P). If the value of n is not important, instead of n-gon we shall often say *polygon*. Polygons may have coinciding vertices (provided they are not adjacent), selfintersections and multiple selfintersections, overlaps or coincidences of edges. Moreover, it is possible for a polygon to be *subdimensional*, that is, to have a segment as its convex hull.

We need several concepts that specify different classes of polyhedra. They are illustrated in Fig. 3.

A polygon P is said to be *weakly convex* if either

(i) P is a simple polygon, and it is contained in the boundary of its convex hull; in other words, P is obtained from the strictly convex polygon conv(P) by subdividing (that is, inserting additional vertices along) the edges of bd(conv(P)), the boundary of its convex hull; or
(ii) P is subdimensional, and coincides with a subdivision of a segment $[A,B]$ in two (possibly coinciding) ways.

It is obvious that, in general, the "convex polygons" obtained in Theorems 1 and 2 are, in fact, only weakly convex. Neither flips, nor flipturns, can lead to a strictly convex polygon if one starts with a weakly convex one that is not strictly convex; even if the starting polygon has no collinear adjacent edges, such edges may appear after flips or flipturns (see Fig. 4).

A polygon P is said to be *nearly convex* if either

(i) P is contained in bd(conv(P)); or
(ii) P is subdimensional.

It is clear that every weakly convex polygon is nearly convex.

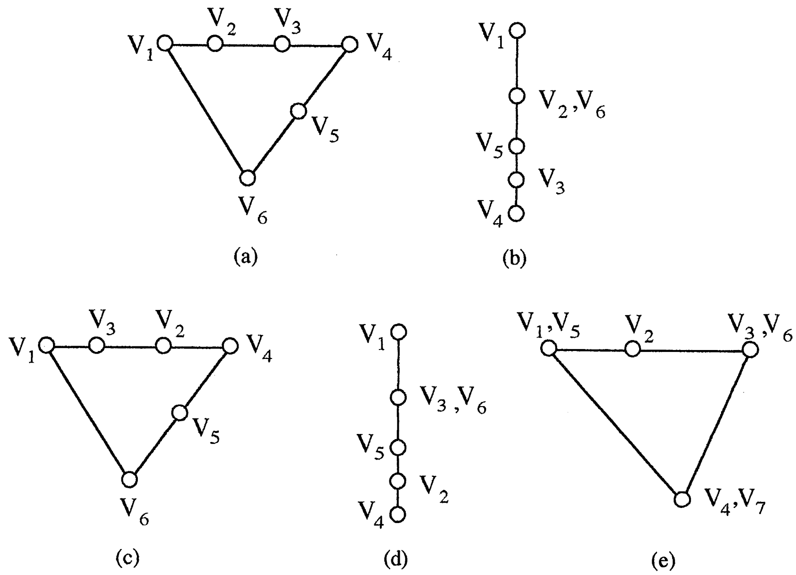

Fig. 3. Nearly convex polygons; for clarity, vertices are indicated by small circles and labelled. Polygons (a) and (b) are weakly convex, polygons (c), (d) and (e) are nearly convex but not weakly convex, Polygons (b) and (d) are subdimensional. All except (e) are exposed.

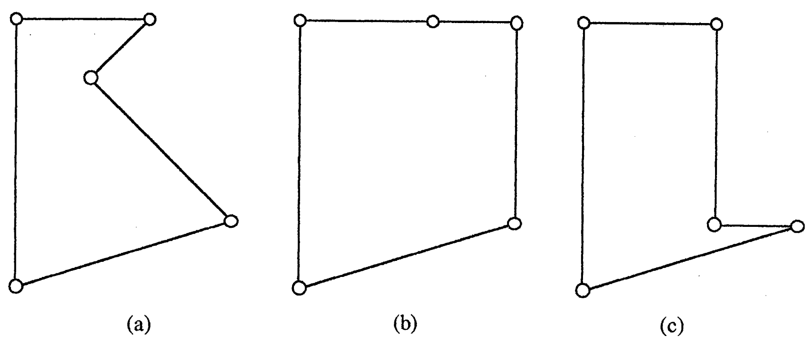

Fig. 4. The polygon in (b) is weakly convex but not strictly convex; it arises from the polygon in (a) by a flip, and from the polygon in (c) by a flipturn.

A polygon P is called *exposed* if it is nearly convex and each vertex of its convex hull coincides with just one vertex of P.

For the more general polygons we consider here, the definition of exposed pairs of vertices has to be modified as well. The modification is quite simple: If A and B are a pair of vertices of P, determining the two paths P^* and P^{**} in P, then A and B are an *exposed pair* provided they are contained in a support line L of conv(P), and neither of the paths P^* and P^{**} is a subdivision of the segment $[A,B]$. Clearly, if P is a simple polygon then the modified definition of exposed pairs coincides with the original one.

It is clear that an exposed polygon is invariant under any flips that may be performed on it. Hence the following is a best possible result:

Theorem 3. *Every polygon in the plane can be transformed into an exposed polygon by a finite sequence of flips, determined at each step by an exposed pair of vertices.*

However, the analogue of Theorem 2 leads to a better result:

Theorem 4. *Every polygon in the plane can be transformed into a weakly convex polygon by a finite sequence of flipturns, determined at each step by an exposed pair of vertices.*

3. Proofs

In order to prove Theorem 3, we first consider the case in which P is subdimensional, hence $C = \mathrm{conv}(P)$ is a segment. If P is not exposed, at least two vertices of P—say A and B—coincide with one of the endpoints of C and form an exposed pair. Taking a supporting line L of C that contains A and B but neither contains C nor is perpendicular to C, and performing the flip with respect to L, leads to a polygon $f(P;A,B)$ which is not subdimensional; see Fig. 5. Since no flip will turn a fulldimensional polygon into a subdimensional one, we can now restrict attention to the case that P is not subdimensional.

We associate with the polygons we are flipping a positive-valued function μ, which strictly increases with every flip. Various choices μ are possible; we shall use the simplest one, which was suggested to us by Ayal Zaks: for any n-gon P, the value of $\mu(P)$ is the sum of the $n(n-1)/2$ distances between pairs of vertices of P. Since we assume that P is full-dimensional, it is obvious that $\mu(P) < \mu(f(P;A,B))$ for every

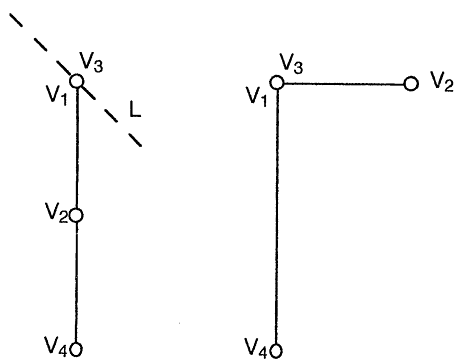

Fig. 5. A subdimensional polygon that is not exposed can be transformed to a full-dimensional polygon by a suitable flip.

flip image of P; see Fig. 1. The existence of such a strictly increasing μ shows that there can be no revisits of any polygon from which we departed by a flip.

We note that if a full-dimensional polygon P is not exposed, then either some vertex of $\text{conv}(P)$ coincides with two (or more) vertices of P, or, failing that, some pair of neighboring vertices of $\text{conv}(P)$ form an exposed pair. In either case a flip is possible.

Let us now define the sequence of flips which, we claim, will lead to an exposed polygon. The choice we make is that, as long as the polygon reached is not exposed, we choose among the applicable flips one which maximizes the increase in μ. If this procedure ends after a finite number of steps, we are done. Hence, we shall assume that the sequence $P_j = [V_{j,1}, V_{j,2}, \dots, V_{j,n}]$ of polygons obtained by successive flips can be continued indefinitely, and we shall show that this leads to a contradiction.

Since the perimeter (sum of lengths of all edges) of a polygon is unchanged under flips, the values of $\mu(P_j)$ are bounded; since they are increasing, they have a limit M. From the uniformly bounded sequence of polygons P_j we can extract a subsequence which converges to a limit-polygon Q, and is such that for each $i = 1, 2, \dots, n$, the sequence of corresponding vertices $V_{j,i}$ also converges to a vertex W_i of $Q = [W_1, W_2, \dots, W_n]$. Now, we first show that Q must be an exposed polygon. Indeed, otherwise there would be a flip that would increase $\mu(Q)$ by a positive δ. Since μ is a continuous function, and forming the convex hull is a continuous operation, every P_j sufficiently far in the subsequence would admit a flip which would increase μ by at least $\delta/2$, thus contradicting the choice of the flips—since sufficiently far polygons P_j have maximal increases of μ which tend to 0.

Now we are almost done. Since Q is exposed, every vertex W_i of $\text{conv}(Q)$ can be strictly separated by a line L from all the other vertices of Q. Let C_i be a circular disk centered at W_i and not meeting L, and such that circles of the same radius centered at all the other vertices of Q also miss L. The vertex W_i is a limit of the $V_{j,i}$'s of the convergent subsequence, hence all but a finite number of them are contained in C_i. By the choice of C_i each such $V_{j,i}$ is a vertex of $\text{conv}(P_j)$, and as such is not moved by any of the following flips. Therefore, all such vertices $V_{j,i}$ coincide with W_i. Since there are at most n vertices W_i that are vertices of $\text{conv}(Q)$, it follows that for all sufficiently large j the polygons $\text{conv}(P_j)$ coincide, hence coincide with $\text{conv}(Q)$. Thus, the sequence P_j cannot be infinite, and Theorem 3 is proved.

Turning now to a proof of Theorem 4, we note that the general idea of the proof is similar to the above, but with two main differences. First, we have to find a different function μ to use in the full-dimensional case, since the one used above does not necessarily increase under flipturns. Second, we shall consider P as having one of the two possible orientations; then the edges of P are vectors, and these vectors are only permuted in the order in which they appear in the polygon when the polygon is flipturned. However, as shown by an example in Section 4, we cannot expect to find a function μ that increases under *every* flipturn. Hence we shall be satisfied with a function μ that increases under every flipturn *we use*; such a μ will show that not more than $(n-1)!$ successive flipturns of this kind are possible if P is an n-gon; it follows that there is no need for any limits, or for convergence arguments.

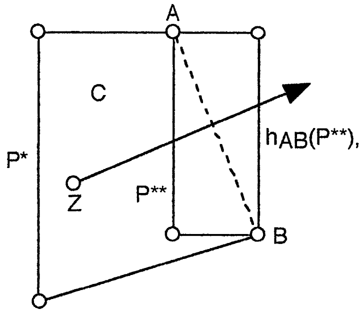

Fig. 6. An illustration of the assertion that the contribution of points of C to α will be the same before and after the flipturn $g(P;A,B)$. By flipturning that path between A and B (in the illustration this is P^{**}) for which $P^{**} \cup [AB]$ is even, a ray through a relatively interior point of the segment $[AQB]$ shows that the parity of the winding number with respect to Z is unchanged, hence odd.

Our choice of μ is the following. Let $\lambda(P)$ be the minimum of the total length of segments in a family that covers all edges of P, and let $\alpha(P)$ be the sum of the areas of all those components of the complement of P in the plane for which the winding number with respect to P is odd. We put $\mu(P) = \lambda(P) + \alpha(P)$, and we shall first justify out claim that the μ strictly increases with every flipturn of the type we use. Clearly, $\lambda(P)$ is not decreasing, but may well be unchanged by a flipturn. Also, $\alpha(P)$ may obviously stay unchanged (for example, if P is subdimensional), but it is less obvious that it cannot decrease under a flipturn.

To show this latter fact, let $\Omega(P)$ be the union of those regions (of the complement of P in the plane) the points of which have an odd winding number. We recall that given a point Z which is not on any edge of P, and given a ray H with endpoint Z which passes through no vertex of P, then the winding number $w(Z,H;P)$ is the number of times H meets an edge of P which crosses H from right to left, less the number of such edges which cross H from left to right. It is easy to show that $w(Z,H;P)$ does not depend on the particular ray H chosen; the common value for all H is the winding number $w(Z;P)$ of the point Z. It is equally simple to show that all points Z in one connected component C of the complement of P in the plane have the same winding number $w(Z;P)$; this common value is the *winding number* $w(C;P)$ of the region C with respect to P. Thus $\Omega(P) = \bigcup\{C\colon\, w(C;P) \text{ is odd}\}$, and $\alpha(P) = \text{area}(\Omega(P))$.

If $\Omega(P) = \emptyset$ then $\alpha(P)$ will not decrease under any flipturn. Hence let $\Omega(P) \neq \emptyset$ and let C be a region which contributes to $\alpha(P)$. For a flipturn h_{AB} determined by extreme vertices A and B of P consider the two paths P^* and P^{**} determined by A and B, and the two polygons $Q^* = P^* \cup [AB]$ and $Q^{**} = P^{**} \cup [AB]$. Then precisely one of the numbers $w(Z;Q^*)$ and $w(Z;Q^{**})$ is odd for every Z in C, hence (see the caption of Fig. 6) the contribution of points of C to α will be the same before and after the flipturn $g(P;A,B)$. On the other hand, points not in $\Omega(P)$ may after a flipturn belong to an Ω-component of the image $g(P;A,B)$, thus leading to an increase in α.

Since both $\lambda(P)$ and $\alpha(P)$ are nondecreasing under any flipturns of P, we would be done if their sum, $\mu(P) = \lambda(P) + \alpha(P)$, were strictly increasing for every flipturn of P. However, this is not always the case, and we need to restrict the types of flipturns which we shall perform and to show that for them there is a strict increase in μ.

If P is a subdimensional n-gon let $\mathrm{conv}(P) = Q = [R, S]$ be the segment determined by P. If P is not weakly convex, then either one of the endpoints of Q corresponds to at least two vertices of P, or, failing that, at least one of the paths P^* and P^{**} determined by the (unique) vertices V_i at R and V_j at S contains overlapping edges. In the first case, we perform a flipturn with respect to the two coinciding vertices; the resulting polygon determines a segment which is longer than Q hence yields a strict increase in $\lambda(P) = \mu(P)$. In the second case, let V_k be that vertex which is closest to V_i among those vertices of P for which both incident edges lead in the direction away from V_i. Then we perform a flipturn with respect to V_i and V_k. (The first case could be interpreted as a special case of the second.) Again the segment determined by the resulting polygon has increased in length, thus increasing $\lambda(P) = \mu(P)$. Since flipturn images of subdimensional polygons are subdimensional, we are done in case of such polygons.

Turning now to the case of full-dimensional P, let $Q = \mathrm{conv}(P)$. We first consider the case in which A and B are exposed vertices of P such that A is a vertex of Q and B satisfies one of the following conditions:

(i) B coincides with A;
(ii) B does not coincide with A, but B is a point of a supporting line L of Q that passes through A, and two edges of P incident with B overlap but contain no relatively interior point of the segment $[A, B]$;
(iii) B does not coincide with A, but B is a point of a supporting line L of Q that passes through A, two edges of P incident with B do not overlap but neither contains a relatively interior point of the segment $[A, B]$.

In cases (i) and (ii) the value of $\mu(P)$ increases under the flipturn $g(P; A, B)$ since $\lambda(P)$ clearly increases. In case (iii), the flipturn increases $\lambda(P)$ if one of the paths P^* and P^{**} determined by A and B is contained in L, and it increases $\alpha(P)$ if neither of these paths is contained in L.

If no such A and B exist, let L be a support line of Q determined by an edge E of Q. Then, since the former case is assumed not to occur, the boundary of Q contains no overlapping edges of P. Hence either the part of P contained in L is just a subdivision of E, or else there are exposed points A and B of P contained in L such that neither of the paths P^* and P^{**} is contained in L, and the edges incident with A or B contain no relatively interior points of the segment $[A, B]$. Then the flipturn $g(P; A, B)$ increases $\mu(P)$ since points near $[A, B]$ now contribute to $\alpha(P)$.

Since the possibility of applying flipturns of the kinds described can be absent only if the polygon is weakly convex, this completes the proof of Theorem 4.

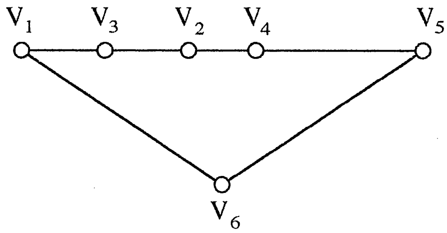

Fig. 7. The flipturn of the polygon $P = \lfloor V_1, V_2, V_3, V_4, V_5, V_6 \rceil$ with respect to the exposed pair of vertices V_2, V_4 yields a polygon congruent to P, showing that no function $\mu(P)$ can strictly increase under every flipturn. The method of proof of Theorem 4 would lead either to the exposed pair V_1, V_4, or to the exposed pair V_5, V_2; in either case $\alpha(P)$ increases, and hence $\mu(P)$ increases as well.

4. Comments

(1) The necessity of distinguishing cases in the proof of Theorem 4 is not due to a failure to find a better function μ. Indeed, as shown by the example in Fig. 7, it is possible that P is congruent with a flipturned image $g(P; A, B)$, so that no μ could strictly increase with every flipturn.

(2) The proofs of Theorems 1 and 2 found in the literature use the idea of a function μ that increases with every flip or flipturn; in all cases the area enclosed by the polygon is taken as μ. This choice does not work for Theorems 3 and 4. The function μ we used in the proof of Theorem 4 could also work for Theorem 3, but the choice we adopted makes for a more elegant proof. Naturally, the μ we used in the proof of Theorem 3 could also be used to establish Theorem 1.

(3) Theorems 3 and 4 are clearly generalizations of Theorems 1 and 2. Indeed, if the starting polygon P is simple then the resulting polygons in both cases are also simple.

(4) In [3], Erdös asked for a proof of the assertion that starting from a simple polygon P one can reach a convex polygon after a finite number of steps, where each step can be described (in the terminology of our Section 1) as performing *simultaneously* all flips possible at the given stage. Sz.-Nagy [6] observed that this construction may lead from a simple P to a selfintersecting polygon, thus halting the construction. However, if we interpret flips in the sense of the definition in Section 2, then it is possible to establish an affirmative solution to Erdös's problem, even without restriction to simple polygons.

(5) From a report of one of the referees and from a friendly communication of Professor Godfried Toussaint we learned of the existence of papers [2,7]. In an expanded version of the former, a proof of Sz.-Nagy's theorem (Theorem 1 above) is given. The latter has a variety of results that overlap our Theorems 3 and 4, and an extensive bibliography.

(6) In [5], an example due to Joss and Shannon is given which shows that even in case of Theorem 1, there is no bound on the number of steps needed to convexify all n-gons, for any fixed $n \geqslant 4$. They also conjectured that in case of Theorem 2 the universal bound $(n-1)!$ could be improved to $\frac{1}{4}n^2$. This conjecture is still open, as is the question whether the same bound applies in case of Theorem 4. The best partial

result is due to Ahn et al. [1], who show that any simple n-gon with edges in k directions can be convexified after at most $n(k-1)/2$ flipturns.

(7) Wegner [8] considered the question of inversion of flips for simple polygons. By this is meant finding a diagonal D of the simple n-gon P, which has its endpoints at vertices A and B of P, and reflecting one of the two arcs of P determined by A and B across D—*provided* the resulting polygon is simple. Wegner conjectured that for every simple polygon, any sequence of inverse flips is finite. However, this conjecture has been disproved (for each $n=4$) by Fevens et al. [4] by an elegant construction. A similar result for every $n \geqslant 4$ is to appear in an expanded version of [4].

(8) It is an open question whether any (simple? unknotted?) polygon in three-dimensional space can be transformed into a weakly convex planar polygon by a finite number of suitable flips or flipturns.

References

[1] H.-K Ahn, P. Bose, J. Czyzowicz, N. Hanusse, E. Kranakis, P. Morin, Flipping your lid, Manuscript, 2000.

[2] T. Biedl, E. Demaine, M. Demaine, S. Lazard, A. Lubiw, J. O'Rourke, M. Overmars, S. Robins, I. Streinu, G.T. Toussaint, S. Whitesides, Locked and unlocked polygonal chains in 3d, Proceedings of the 10th ACM–SIAM Symposium on Discrete Algorithms, 1999, pp. 866–867.

[3] P. Erdős, Problem 3763, Amer. Math. Monthly 42 (1935) 627.

[4] T. Fevens, A. Hernandez, A. Mesa, M. Soss, G. Toussaint, Simple polygons that cannot be deflated, Preprint, 2000.

[5] B. Grünbaum, How to convexify a polygon, Geombinatorics 5 (1995) 24–30.

[6] B. de Sz.-Nagy, Solution of Problem 3763, Amer. Math. Monthly 46 (1939) 176–177.

[7] G. Toussaint, The Erdős–Nagy theorem and its ramifications, Proceedings of the 11th Canadian Conference on Computational Geometry, 16–18 August, Vancouver, BC, Canada, 1999.

[8] B. Wegner, Partial inflation of closed polygons in the plane, Contrib. Algebra Geometry 34 (1993) 77–85.