Flipturning Polygons*

Oswin Aichholzer†
Technische Universität Graz
oaich@igi.tu-graz.ac.at

Carmen Cortés
Universidad de Sevilla
ccortes@cica.es

Erik D. Demaine
University of Waterloo
eddemaine@uwaterloo.ca

Vida Dujmović
McGill University
vida@cs.mcgill.ca

Jeff Erickson‡
University of Illinois
jeffe@cs.uiuc.edu

Henk Meijer
Queen's University
henk@cs.queensu.ca

Mark Overmars
Universiteit Utrecht
markov@cs.uu.nl

Belén Palop§
Universidad Rey Juan Carlos
b.palop@escet.urjc.es

Suneeta Ramaswami
Rutgers University
rsuneeta@crab.rutgers.edu

Godfried T. Toussaint¶
McGill University
godfried@cs.mcgill.ca

Submitted to *Discrete & Computational Geometry*: August 15, 2000
Revised and resubmitted: March 20, 2002

Abstract

A *flipturn* that transforms a nonconvex simple polygon into another simple polygon by rotating a concavity 180 degrees around the midpoint of its bounding convex hull edge. Joss and Shannon proved in 1973 that a sequence of flipturns eventually transforms any simple polygon into a convex polygon. This paper describes several new results about such flipturn sequences. We show that any orthogonal polygon is convexified after at most $n-5$ arbitrary flipturns, or at most $\lfloor 5(n-4)/6 \rfloor$ well-chosen flipturns, improving the previously best upper bound of $(n-1)!/2$. We also show that any simple polygon can be convexified by at most n^2-4n+1 flipturns, generalizing earlier results of Ahn *et al.* These bounds depend critically on how degenerate cases are handled; we carefully explore several possibilities. We prove that computing the longest flipturn sequence for a simple polygon is NP-hard. Finally, we show that although flipturn sequences for the same polygon can have significantly different lengths, the shape and position of the final convex polygon is the same for all sequences and can be computed in $O(n \log n)$ time.

*Portions of this work were done at the Workshop on Polygonal Entanglement Theory held February 4–11, 2000 at McGill University's Bellairs Research Center, Holetown, Barbados. For the most recent version of this paper, see http://www.uiuc.edu/~jeffe/pubs/flipturn.html.

†Supported by the Austrian Programme for Advanced Research and Technology (APART).

‡Contact author. Partially supported by a Sloan Fellowship.

§Partially supported by Proyecto DGES-MEC-PB98-0933.

¶Partially supported by NSERC Grant No. A9293.

1 Introduction

A central problem in polymer physics and molecular biology is the reconfiguration of large molecules (modeled as polygons) such as circular DNA [14]. Most of the research in this area involves computer-intensive Monte-Carlo simulations. To simplify these simulations they are usually restricted to the integer lattices $\mathbb{Z}^2$ and $\mathbb{Z}^3$, although some work has also been done on the FCC lattice [23]. Like the related algorithmic robotics research on linkages, the problems of interest to physicists and biologists involve closed simple polygons [10], open simple polygonal chains [20] and simple polygonal trees [13], *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 self-avoiding walk from scratch is difficult, especially if it must return to its starting point as in the case of polygons. The waiting time is too long due to attrition; 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 random chains or polygons is to modify one such object into another using a simple operation called a *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, pivots are seen as instantaneous combinatorial changes, not continuous processes. In general the pivots used are selected from a large variety of transformations such as reflections, rotations, or 'cut and paste' operations on certain subchains. We refer the reader to a multitude of such problems and results in [18]. For example, Madras and Sokal [19] have shown that for all $d \geq 2$, every simple lattice polygonal chain with n edges in $\mathbb{Z}^d$ 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 right angles.

In order to prove the ergodicity of their self-avoiding walks, polymer physicists are interested in convexifying polygons (and straightening open polygonal chains). If a polygon can be transformed to some canonical convex configuration, then any simple polygon can be reconfigured to any other via this intermediate position. This theoretical aspect of polymer physics research resembles the algorithmic robotics work on convexification of polygonal linkages. We refer the reader to survey papers of O'Rourke [21] and Toussaint [27] for further references in the latter area.

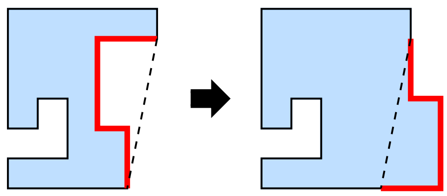

Figure 1.1. A flipturn. The edges of the pocket are bold (red), and its lid is dashed.

In this paper, we are concerned with one type of pivot of central concern in polymer physics research. This pivot is usually called an *inversion* in the physics literature, but since it seems to have been first proposed in an unpublished 1973 paper of Joss and Shannon [15], we will follow their terminology and call it a *flipturn*. Flipturns are defined as follows. Any nonconvex polygon has at least one concavity, or *pocket*. Formally, a pocket of a nonconvex polygon P is a maximal connected sequence of polygon edges disjoint from the convex hull of P except at its endpoints. The line segment joining the endpoints of a pocket is called the *lid* of the pocket. A flipturn rotates

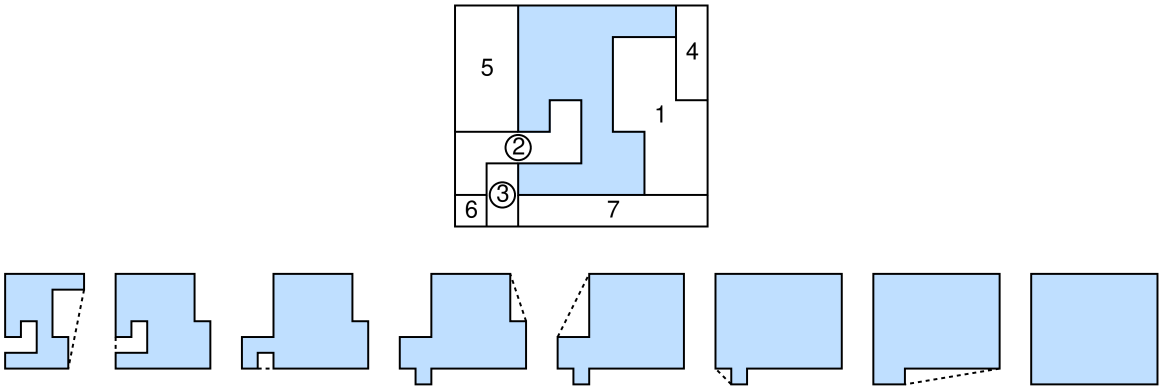

Figure 1.2. A convexifying flipturn sequence.

a pocket 180 degrees about the midpoint of its lid, or equivalently, reverses the order of the edges of a pocket without changing their lengths or slopes. Figure 1.1 shows the effect of a single flipturn on a nonconvex orthogonal polygon, and Figure 1.2 shows a sequence of flipturns transforming this polygon into a rectangle. We will illustrate such sequences by overlaying the resulting polygons and labeling the area added by each flipturn by its position in the sequence. (The circled numbers will be explained in Section 2.)

1.1 Previous and Related Results

Joss and Shannon proved that any simple polygon with n sides can be convexified by a sequence of at most $(n-1)!$ flipturns, by observing that each flipturn produces a new cyclic permutation of the edges. Since each flipturn increases the polygon's area, each of the $(n-1)!$ cyclic permutations can occur at most once. We can immediately improve this bound to $(n-1)!/2$ by observing that at most half of the $(n-1)!$ cyclic permutations describe a simple polygon with the proper orientation. Although this is the best bound known, it is extremely loose; Joss and Shannon conjectured that $n^2/4$ flipturns are always sufficient. Grünbaum and Zaks [16] showed that even crossing polygons could be convexified with a finite number of flipturns. Biedl [5] discovered a family of polygons that are convexified only after $(n-2)^2/4$ badly chosen flipturns, nearly matching Joss and Shannon's conjectured upper bound. Ahn *et al.* [1] recently proved that any simple polygon can be convexified by a sequence of at most $n(n-3)/2$ so-called *modified* flipturns (which we define in Section 2). Better results are known for orthogonal and lattice polygons in the plane. Dubins *et al.* [10] showed that any simple lattice polygon in the plane can be convexified with $n-4$ well-chosen flipturns [18]. Until recently this was the best upper bound known. Ahn *et al.* [1] show that any polygon with s distinct edge slopes can be convexified by $\lceil n(s-1)/2-s \rceil$ modified flipturns; in particular, $n/2-2$ modified flipturns suffice to convexify any orthogonal polygon.

There are significant differences between flipturns and another common pivoting rule, the *Erdős-Nagy flip* [12, 15, 26, 28], in which a pocket is *reflected* across its lid. As with flipturns, any convex polygon can be convexified using a finite number of flips. Unlike flipturns, however, the number of flips required is not bounded by any function of n. Joss and Shannon constructed a family of quadrilaterals that require an unbounded number of flips to convexify [15].

Another important difference is that flipturns preserve the *slopes* of polygon edges, while flips preserve their *order* around the polygon. If we always direct polygon edges so that they form a counterclockwise cycle, then flipturns do not change the direction of any edge. Since flipturns also do not change edge lengths, the final convex shape is the same for any convexifying flipturn sequence. We can easily compute this shape in $O(n \log n)$ time by sorting the edges of the original polygon by orientation, breaking ties arbitrarily. On the other hand, starting from the same simple polygon, different sequences of flips can lead to different convex polygons—see Figure 1.3(a).

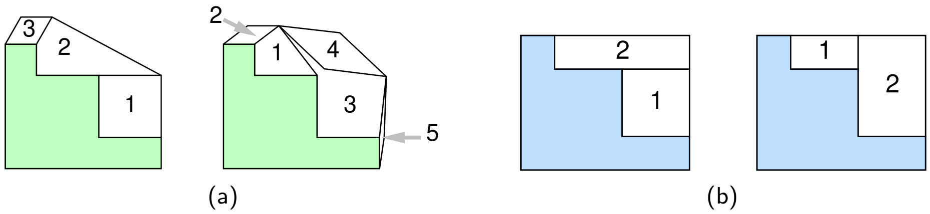

Figure 1.3. (a) Different Erdős-Nagy flip sequences can lead to different convex shapes. (b) Different flipturn sequences always lead to the same convex shape.

For further results on both flips and flipturns for general polygons, simpler algorithms, and a more complete history of the problem, see [28].

1.2 New Results

Our results depend critically on the behavior of flipturns in degenerate cases. In Section 2, we offer three alternate definitions: *standard*, *extended*, and *modified* flipturns. As our naming suggests, we believe that standard flipturns are closest to the original definition of Joss and Shannon. Modified flipturns were introduced by Ahn *et al.* [1].

In Section 3, we show that under all three definitions, both the shortest and longest flipturn sequences required to convexify any orthogonal n-gon have length $O(n)$, and that our upper bounds are tight in the worst case up to small constant factors. Our new bounds are summarized in the first two rows of Tables 1.1 and 1.2; the last row of each table gives the corresponding results of Ahn *et al.* for modified flipturns. We also show that the shortest and longest flipturn sequences for the same orthogonal polygon can differ in length by at least $(n-4)/4$.

Using techniques developed in Section 3, we prove in Section 4 that any simple n-gon is convexified after at most n^2-4n+2 standard or extended flipturns, generalizing the modified flipturn results of Ahn *et al.* [1]. This matches both Biedl's $(n-2)^2/4$ lower bound [5] and Joss and Shannon's conjectured $n^2/4$ upper bound up to a small constant factor. It remains open whether the *shortest* flipturn sequence for every polygon has linear length, or whether some polygon always requires a quadratic number of flipturns.

Section 5 considers the complexity of computing optimal flipturn sequences. We show that computing the longest flipturn sequence for a given simple polygon and finding the shortest convexifying sequence of *generalized* flipturns are both NP-hard.

In Section 6, we prove that for any simple polygon, every sequence of flipturns eventually leads to the same convex polygon, which we can compute in $O(n \log n)$ time. As we already mentioned, the fact that the *shape* of the final convex polygon is independent of the flipturn sequence is rather obvious, but the independence of the final polygon's *position* requires considerably more effort.

In an expanded version of this paper [2], we also describe a data structure to maintain a simple n-gon and its convex hull, so that any flipturn can be performed in $O(\log^4 n)$ amortized time. Our data structure is an extension of the dynamic planar convex hull structure of Overmars and van

Flipturn type	Shortest flipturn sequence	Longest flipturn sequence
standard	$\lfloor 3(n-4)/4 \rfloor \leq ?? \leq \lfloor 5(n-4)/6 \rfloor$	$\lfloor 5(n-4)/6 \rfloor \leq ?? \leq n-5$
extended	$\lfloor 3(n-4)/4 \rfloor$	$\lfloor 3(n-4)/4 \rfloor \leq ?? \leq n-5$
modified [1]	$(n-4)/2$	$(n-4)/2$

Table 1.1. Bounds for shortest and longest flipturn sequences for orthogonal polygons. See Section 3.

Flipturn type	s-oriented polygons	arbitrary polygons
standard	$ns - \lfloor (n+5s)/2 \rfloor - 1$	$n^2 - 4n + 1$
extended	$ns - \lfloor (n+5s)/2 \rfloor - 1$	$n^2 - 4n + 1$
modified [1]	$\lceil n(s-1)/2 \rceil - s$	$n(n-3)/2$

Table 1.2. Upper bounds for longest flipturn sequences of more general polygons. For arbitrary polygons, the lower bound is $(n-2)^2/4$ [5]. See Section 4.

Leeuwen [22]. Together with the results of Sections 3 and 4, this implies that we can compute a convexifying sequence of flipturns for any polygon in $O(n^2 \log^4 n)$ time, or for any orthogonal polygon in $O(n \log^4 n)$ time.

2 The Importance of Being Degenerate

The behavior of flipturn sequences depends critically on how flipturns are defined in degenerate cases. In the general case, a lid is an edge of the polygon's convex hull. However, in degenerate cases where three or more vertices are collinear[1], a lid can be a proper subset of a convex hull edge according to Joss and Shannon's original definition [15]. Although there are several different types of degeneracies, only one type will actually affect our results. We call a pocket or flipturn *degenerate* whenever the two edges just outside the pocket lie on the same line. For almost all simple polygons—that is, for all but a measure-zero subset—every flipturn sequence consists entirely of non-degenerate flipturns. In our illustrations of flipturn sequences such as Figure 1.2, degenerate flipturns are indicated by circled numbers.

Since flipturning about a proper subset of a convex hull edge may seem unnatural, we offer the following alternative definition. An *extended* pocket of a polygon is a chain of at least two edges joining an adjacent pair of convex hull vertices. An extended flipturn rotates an extended pocket 180 degrees about the midpoint of its lid, which is a complete convex hull edge. An extended pocket or flipturn is *degenerate* if and only if the two edges just *inside* the pocket lie on the same line.

Another alternative is proposed by Ahn *et al.* [1], who define *modified* pockets as follows. Consider a standard pocket with vertices $v_i, v_{i+1}, \ldots, v_j$ (where index arithmetic is modular). If the nearby vertex v_{j+1} lies on the line through v_i and v_j, then the chain of edges from v_i to v_{j+1} is a modified pocket; otherwise, the standard pocket from v_i to v_j is a modified pocket. If the standard pocket is degenerate, the modified pocket contains one of the two colinear boundary edges.

Figure 2.1 illustrates a standard flipturn, an extended flipturn, and one of two possible modified flipturn of the 'same' degenerate pocket of a polygon. Note that a single extended flipturn can simultaneously invert several standard or modified pockets.

[1] A vertex of a polygon is a boundary point with internal angle not equal to π. Although computer representations of polygons can store points in the interior of an edge, we will not consider such points to be vertices.

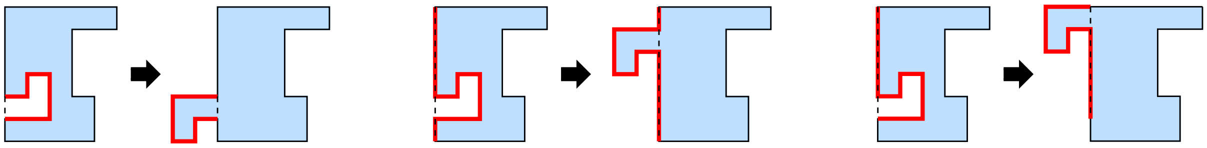

(a) Standard (b) Extended (c) Modified

Figure 2.1. Three types of degenerate flipturns. Compare with Figure 1.1.

In the next section, we will focus entirely on *orthogonal* polygons, each of whose edges is either horizontal or vertical. A pocket of an orthogonal polygon is degenerate if and only if its lid is horizontal or vertical. To emphasize this point, we will refer to degenerate and non-degenerate pockets of orthogonal polygons as *orthogonal* and *diagonal*, respectively. As Figure 2.1 illustrates, the exact behavior of an orthogonal flipturn depends on which of the three definitions we use; diagonal flipturns are the same under any definition.

3 Orthogonal Polygons

In this section, we derive bounds on the maximum length of either the shortest or longest convexifying flipturn sequences for orthogonal polygons. The bounds for the shortest sequence tell us how quickly we can convexify a polygon if we choose flipturns intelligently; the longest sequence bounds tell us how many flipturns we can perform even if we choose flipturns blindly. Our results are summarized in the first two rows of Table 1.1.

By any of our three definitions, a diagonal flipturn reduces the number of vertices of an orthogonal polygon by two; specifically, the endpoints of the flipturned pocket lie in the interior of edges of the new polygon. This observation immediately implies the following upper and lower bounds.

Theorem 3.1. *Any orthogonal n-gon is convexified by any sequence of $(n-4)/2$ diagonal flipturns.*

Theorem 3.2. *At least $(n-4)/2$ flipturns are required to convexify any orthogonal n-gon.*

For almost all orthogonal polygons, every flipturn sequence contains only diagonal flipturns. In this case, exactly $(n-4)/2$ flipturns are necessary and sufficient to convexify the polygon, and these flipturns can be chosen arbitrarily. Thus, any discussion of flipturn sequences on orthogonal polygons only becomes interesting if orthogonal flipturns are possible.

The presence of degeneracies has little effect on the behavior of modified flipturns; *every* modified flipturn on an orthogonal polygon removes two vertices. We immediately obtain the following result, most of which is a special case of a theorem of Ahn *et al.* [1].

Theorem 3.3. *Exactly $(n-4)/2$ modified flipturns are necessary and sufficient to convexify any orthogonal n-gon, and these flipturns can be chosen arbitrarily.*

Since this theorem completely characterizes the lengths of modified flipturn sequences for orthogonal polygons, the rest of this section will focus entirely on standard and extended flipturns.

3.1 Shortest Flipturn Sequences

Here we develop upper and lower bounds on the length of the shortest sequence of flipturns required to convexify an orthogonal polygon.

Theorem 3.4. *For all n, there is an orthogonal n-gon that requires $\lfloor 3(n-4)/4 \rfloor$ standard or extended flipturns to convexify.*

Proof: When n is a multiple of 4, the polygon consists of a horizontally symmetric rectangular 'comb' with $n/4$ 'teeth'; if n is not a multiple of 4, we add a small rectangular notch in a bottom corner of the polygon. See Figure 3.1. (We consider a rectangle to be a comb with one tooth.) Both the teeth and the gaps between them decrease in height as they approach the middle of the polygon. Since the polygon is symmetric about its vertical bisecting line, standard and extended flipturns have exactly the same effect. The only way to eliminate the comb is through a sequence of orthogonal flipturns across the top edge of the polygon's bounding box; each such flipturn eliminates exactly one tooth. It easily follows that *every* flipturn sequence for this polygon has length $\lfloor 3(n-4)/4 \rfloor$. $\square$

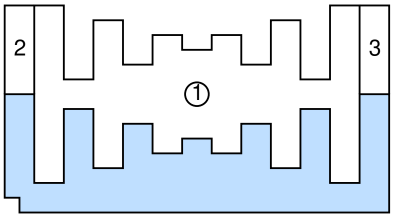

Figure 3.1. An orthogonal n-gon requiring $\lfloor 3(n-4)/4 \rfloor$ flipturns to convexify.

For any polygon P, let $\square(P)$ denote its axis-aligned bounding rectangle.

Lemma 3.5. *Let P be an orthogonal polygon.*

- **(a)** *If some vertex of $\square(P)$ is not a vertex of P, then P has a diagonal pocket.*
- **(b)** *If two adjacent vertices of $\square(P)$ are not vertices of P, then we can perform at least two consecutive diagonal flipturns on P.*

Proof: (a) Suppose some corner of $\square(P)$ is not a vertex of P. Some edge of $\mathrm{conv}(P)$ lies on a line separating the missing corner from the interior of P. This edge contains a diagonal lid.

(b) Without loss of generality, suppose P does not contain the top left and top right vertices of $\square(P)$. Part (a) implies that P has at least two diagonal pockets. Let Q be the result of flipturning one of these pockets. Since the width of the flipturned pocket is less than the width of P, and thus less than the width of Q, at least one of the upper corners of $\square(Q)$ is not a vertex of Q. (As Figure 3.2 shows, flipturning one pocket can hide the opposite corner.) Thus, by part (a), Q still has at least one diagonal pocket. $\square$

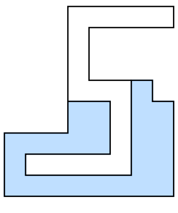

Figure 3.2. Flipturning one diagonal pocket can hide another one.

Theorem 3.6. *Any orthogonal n-gon can be convexified by a sequence of at most $\lfloor 3(n-4)/4 \rfloor$ extended flipturns.*

Proof: We achieve the stated upper bound by performing an orthogonal extended flipturn only when no diagonal pockets are available. By Lemma 3.5, we are forced to perform an orthogonal flipturn on a polygon P only if all four corners of $\Box(P)$ are also vertices of P.

Let P be a nonconvex orthogonal n-gon with no diagonal pockets. Without loss of generality, suppose P has an extended orthogonal pocket whose lid lr is the top edge of $\Box(P)$. This pocket lies strictly between the vertical lines through l and r. Let P_1 be the polygon obtained by flipturning this extended pocket. The highest vertices of P_1 are vertices of the newly flipturned pocket, and thus lie between the vertical lines through l and r. Thus, neither of the top vertices of $\Box(P_1)$ is a vertex of P_1, so by Lemma 3.5, we can perform at least two consecutive diagonal flipturns on P_1. See Figure 3.3.

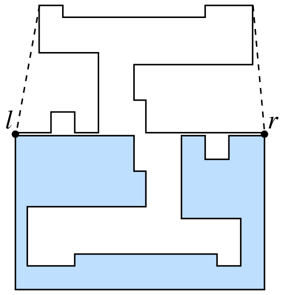

Figure 3.3. Any orthogonal extended flipturn creates at least two diagonal pockets.

In other words, any orthogonal extended flipturn can be followed by at least two diagonal flipturns. Thus, if we perform orthogonal flipturns only when no diagonal flipturn is available, any three consecutive flipturns eliminate at least four vertices. □

Theorem 3.4 implies that this result is the best possible for extended flipturns. For standard flipturns, we obtain the following slightly weaker upper bound.

Theorem 3.7. *Any orthogonal n-gon can be convexified by a sequence of at most $\lfloor 5(n-4)/6 \rfloor$ standard flipturns.*

Proof: As in the previous theorem, we achieve the upper bound by performing orthogonal flipturns only when no diagonal flipturn is available. However, we also choose orthogonal flipturns carefully if more than one is available. Say that an orthogonal flipturn is *good* if it can be followed by at least two diagonal flipturns and *bad* otherwise. We will perform a bad orthogonal flipturn only if no good orthogonal flipturn or diagonal flipturn is available.

Let P be an orthogonal polygon. Without loss of generality, consider a forced orthogonal flipturn whose lid bc lies on the top edge of $\Box(P)$, and let P_1 be the polygon resulting from this flipturn. See Figure 3.4(a). The lid endpoints b and c must lie in two different pockets of P_1, since the flipturned pocket touches the top of $\Box(P_1)$. The horizontal width of the pocket must be less than the horizontal width of P, so P_1 cannot have both the upper left and upper right corners of $\Box(P_1)$ as vertices. Thus, by Lemma 3.5, any forced orthogonal flipturn can be followed first by a diagonal flipturn and then by at least one more (possibly orthogonal) flipturn. In particular, any bad flipturn can be followed by exactly one diagonal flipturn.

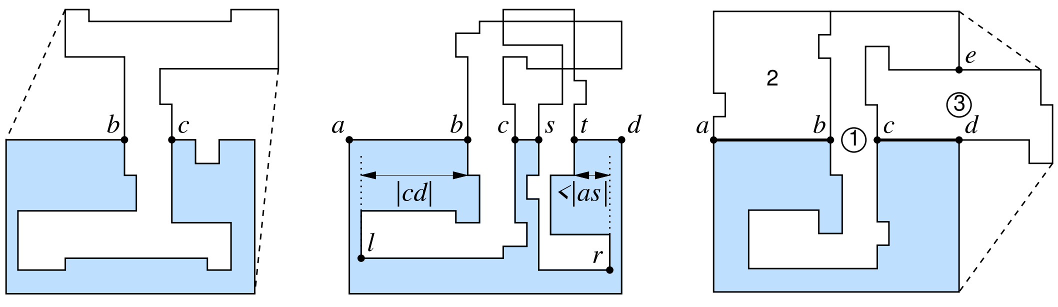

Figure 3.4. (a) A forced orthogonal flipturn creates at least two pockets, at least one of which is diagonal. (b) A polygon with only bad pockets cannot have both dexter and sinister pockets on the same edge. (c) A forced bad orthogonal flipturn (flipturn ①) creates a good orthogonal pocket (flipturn ③).

Let P be a polygon with no diagonal pockets or good orthogonal pockets. Consider a bad orthogonal flipturn whose lid bc is a subset of the top edge ad of $\square(P)$, and let P_1 be the resulting polygon. Exactly one of the top corners of $\square(P_1)$ is a vertex of P_1. If this is the top right corner, call pocket bc *dexter*; otherwise, call it *sinister*. Without loss of generality, suppose the pocket bc is dexter. Let P_2 be the polygon resulting from the only available diagonal flipturn, whose lid is the upper left edge of $\operatorname{conv}(P_1)$. Since P_2 must have no diagonal pockets, this flipturn moves vertex b to the upper left corner of $\square(P_2)$. See Figure 3.4(c).

If some pocket had a lid in ab, that pocket would be inverted by the diagonal flipturn on P_1 and P_2 would have a diagonal pocket, contradicting our assumption that pocket bc is bad. Similarly, if there is a bad pocket with lid in cd, it cannot be dexter. Suppose there is a sinister pocket with lid $st \subset cd$. Let l be a leftmost point in pocket bc, and let r be a rightmost point in pocket st. See Figure 3.4(b). The horizontal distance from l to b must be equal to $|cd|$, and the horizontal distance from t to r must equal to $|as|$, since both pockets are bad. But this is impossible, since $|cd| + |as| > |ad|$. We conclude that bc must be the *only* lid on the top edge of $\square(P)$.

Now consider the orthogonal pocket of P_2 created when pocket bc is flipturned. Its lid de lies on the right edge of $\square(P_2)$. We claim that this pocket must be good. Let P_3 be the resulting polygon when this pocket is flipturned. Since cd is the bottommost edge of pocket de, nothing in P_3 lies above and to the right of vertex e, so the upper right vertex of $\square(P_3)$ is not a vertex of P_3. Since the height of pocket de is less than the height of the original polygon P, the bottom right vertex of $\square(P_3)$ is also not a vertex of P_3. Therefore, by Lemma 3.5, P_3 can undergo at least two consecutive flipturns.

We have just shown that any forced bad flipturn is immediately followed by a diagonal flipturn, a good orthogonal flipturn, and then two diagonal flipturns. Thus, any five consecutive flipturns include at least three diagonal flipturns, which remove at least six vertices from the polygon. $\square$

3.2 Longest Flipturn Sequences

We now prove upper and lower bounds on the maximum number of flipturns that an orthogonal polygon can undergo before becoming convex.

We derive our upper bounds by counting certain special edges of the polygon. We call a polygon edge a *bracket* if either both its vertices are convex or both its vertices are concave. Any orthogonal n-gon has at least four brackets (its highest, leftmost, lowest, and rightmost edges) and unless $n = 4$, at most $n - 2$ brackets.

Theorem 3.8. *For all $n > 4$, the longest standard or extended flipturn sequence for any orthogonal n-gon has length at most $n - 5$.*

Proof: We claim that flipturns do not increase the number of brackets, and that any orthogonal flipturn decreases the number of brackets by two. Let P be an orthogonal polygon and let Q be the result of one flipturn. Any bracket of P that lies completely outside the flipturned pocket is still a bracket in Q; any bracket completely inside the flipturned pocket is inverted, but remains a bracket. Thus, to prove our claim, it suffices to consider just four edges, namely, the two edges adjacent to each endpoint of the lid. If any of these four edges is a bracket, we call it a *lid bracket*. We distinguish between *inner* and *outer* lid brackets, which lie inside and outside the pocket, respectively.

After symmetry considerations, there are only three cases to check for orthogonal pockets and ten cases for diagonal pockets. Most of these cases are completely specified by the number of inner and outer lid brackets. The only exception is a pocket with one lid bracket of each type, which can either share a lid endpoint or not. The cases are enumerated for standard flipturns in Figure 3.5; the cases for extended flipturns are almost identical.

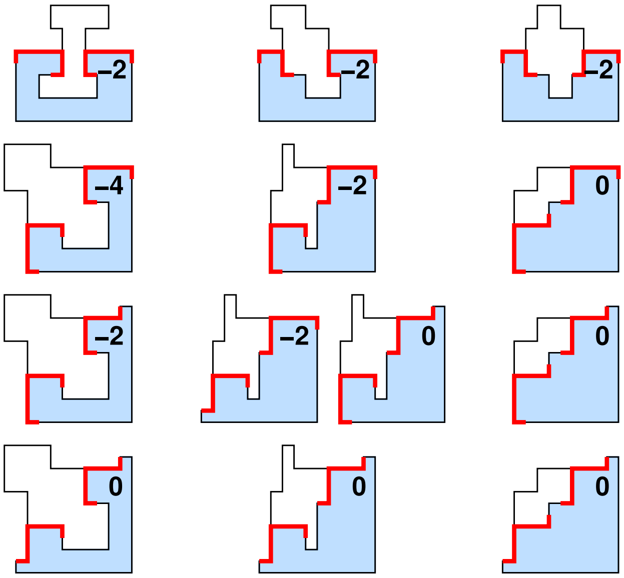

Figure 3.5. Thirteen classes of standard flipturns and the number of brackets they remove. Only the bold (red) edges are important. The top row shows orthogonal flipturns; the other rows show diagonal flipturns with two, one, and no outer lid brackets. The columns show flipturns with two, one, and no inner lid brackets. Compare with Figure 4.1.

Since each orthogonal flipturn removes two brackets and no diagonal flipturn adds brackets, there can be at most $(n - 6)/2$ orthogonal flipturns. Since each diagonal flipturn removes two vertices and each orthogonal flipturn leaves the number of vertices unchanged, any convexifying flipturn sequence contains exactly $(n - 4)/2$ diagonal flipturns. Thus, there can be at most $(n - 6)/2 + (n - 4)/2 = n - 5$ flipturns altogether. $\square$

How tight is this upper bound? As for the shortest flipturn sequence, the answer depends on whether we consider standard or extended flipturns. Unfortunately, we do not obtain an exact answer in either case.

Theorem 3.9. *For all n, there is an orthogonal n-gon that can undergo $\lfloor 3(n-4)/4 \rfloor$ extended flipturns.*

Proof: This follows directly from Theorem 3.4. □

Theorem 3.10. *For all n, there is an orthogonal n-gon that can undergo $\lfloor 5(n-4)/6 \rfloor$ standard flipturns.*

Proof: We construct an orthogonal n-gon P_n essentially by following the proof of Theorem 3.7. P_4 is a rectangle. P_6 is an **L**-shaped hexagon, which is convexified by one flipturn. P_8 is a rectangle with a rectangular orthogonal pocket in one side, which requires three flipturns to convexify. For all $n \geq 10$, P_n consists of a rectangle with a single **L**-shaped pocket, where the tail of the **L** is an inverted and reflected copy of P_{n-6}. See Figure 3.6. In the language of the proof of Theorem 3.7, P_n's only pocket is *bad*—flipturning it creates one diagonal pocket and one orthogonal pocket. If we flipturn diagonal pockets whenever possible, the first five flipturns eliminate six vertices and leave a distorted P_{n-6}. The theorem follows by induction. □

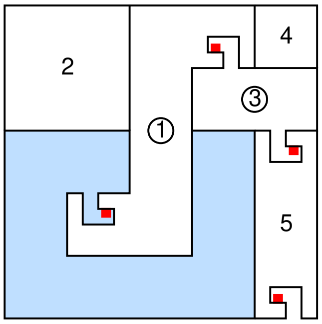

Figure 3.6. An orthogonal n-gon that can undergo $\lceil 5(n-4)/6 \rceil$ standard flipturns. Two levels of recursion are shown. The small squares contain a recursive copy of the polygon.

3.3 Order Matters

We close this section by observing that the shortest and longest flipturn sequences for the same orthogonal polygon can differ significantly in length.

Theorem 3.11. *For infinitely many n, there is an orthogonal n-gon whose shortest and longest standard or extended flipturn sequences differ in length by at least $(n-4)/4$.*

Proof: Figure 3.7 illustrates the recursive construction of such a polygon, for all n of the form $16k+4$. The shortest flipturn sequence for the polygon includes only diagonal flipturns and therefore has length $(n-4)/2$. Another sequence, which we believe to be the longest, requires twelve flipturns to remove every 16 vertices. Figure 3.7 illustrates this long sequence of standard flipturns. The corresponding extended flipturn sequence is essentially equivalent. □

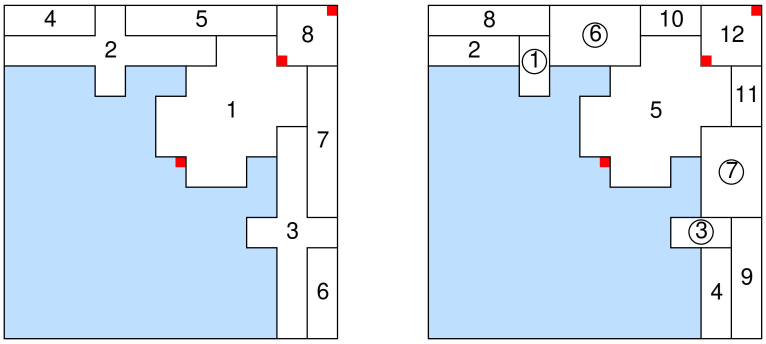

Figure 3.7. An orthogonal polygon that can be convexified with either $(n-4)/2$ or $3(n-4)/4$ flipturns. The small squares contain a recursive copy of the polygon.

4 More General Polygons

In this section, we derive upper bounds for the longest flipturn sequences of arbitrary polygons, generalizing both our earlier results for orthogonal polygons and the modified flipturn results of Ahn *et al.* [1].

Consider an arbitrary polygon P whose boundary is oriented counterclockwise. Let $\vec{e}$ denote the direction of any (oriented) edge e in P, let S be the set of all such edge directions and their reversals. We clearly have $s \le |S| \le 2s$, where s is the number of distinct edge slopes. Ahn *et al.* define the *discrete angle* at a vertex $v = e \cap e'$ to be one more than the number of elements of S strictly inside the angle between $\vec{e}$ and $\vec{e}'$. The *total discrete angle* $D(P)$ is the sum of the discrete angles at the vertices of P. For example, any orthogonal n-gon has total discrete angle n.

Ahn *et al.* prove the following lemma [1]. (Only the first half of this lemma is stated explicitly, but their proof implies the second half as well.)

Lemma 4.1 (Ahn et al. [1]). *Every non-degenerate flipturn decreases* $D(P)$ *by at least two, and every degenerate flipturn leaves* $D(P)$ *unchanged.*

Ahn *et al.* also prove that $D(P) \le n(s-1)$ in general and $D(P) = 2s$ if P is convex. Thus, Lemma 4.1 immediately implies that $\lceil (ns - n - 2s)/2 \rceil \le n(n-3)/2$ nondegenerate flipturns suffice to convexify any polygon. However, since no bound was previously known for the number of degenerate flipturns, this bound does not apply to degenerate polygons. To avoid this problem, Ahn *et al.* introduce modified flipturns, for which degeneracies do not exist. To account for degenerate flipturns under the standard definition, we study the change in the number of *brackets*, which we denote by $B(P)$. Recall from Section 3.2 that a bracket is an edge with either two convex or two concave vertices.

Lemma 4.2. *Every non-degenerate standard or extended flipturn increases* $B(P)$ *by at most two, and every degenerate standard or extended flipturn decreases* $B(P)$ *by at least two.*

Proof: Let P be a simple polygon and let P' be the result of one flipturn. As we argued in the proof of Theorem 3.8, it suffices to focus on the lid brackets, *i.e.*, the brackets touching the endpoints of the lid. Let b and b' denote the number of lid brackets in P and P', respectively, so that $B(P') = B(P) - b + b'$.

For nondegenerate flipturns, we need to show that $b' - b \leq 2$. This is trivial if $b \geq 2$, because we must have $b' \leq 4$. There are three remaining cases to consider: no lid brackets, one outer lid bracket, and one inner lid bracket. For each of these, there are nine subcases, depending on whether each lid endpoint becomes a concave vertex, disappears, or becomes a convex vertex after the flipturn. These cases are enumerated in Figure 4.1.

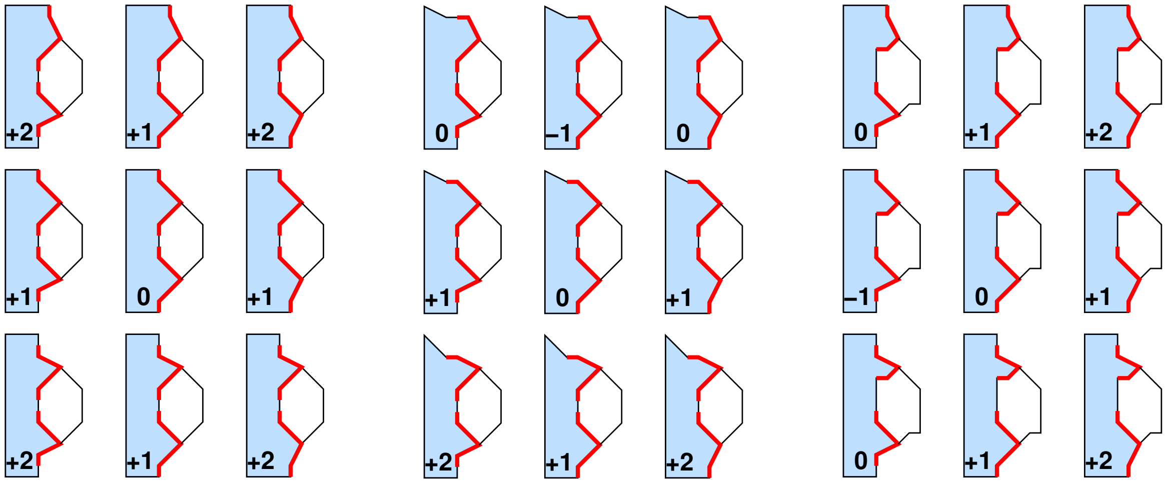

Figure 4.1. Twenty-seven classes of nondegenerate flipturns and the number of brackets they add or remove. Only the bold (red) edges are important. Symmetric cases are omitted. Compare with Figure 3.5.

For degenerate flipturns, we need to show that $b - b' \geq 2$. Degenerate standard pockets always have two outer lid brackets, and both lid endpoints always become concave vertices. Thus, there are only three cases to consider, depending on the number of inner lid brackets, precisely as in Theorem 3.8. (See the top row of Figure 3.5.) Similar arguments apply to degenerate extended flipturns. □

Theorem 4.3. *Every s-oriented polygon is convexified after any sequence of $ns - \lfloor (n+5s)/2 \rfloor - 1$ standard or extended flipturns.*

Proof: We define the *potential* $\Phi(P)$ of a polygon P to be its discrete angle plus half the number of brackets: $\Phi(P) = D(P) + B(P)/2$. For the initial polygon P, we have $D(P) \leq n(s-1)$ and $B(P) \leq n-2$, so the initial potential $\Phi(P)$ is at most $ns - n/2 - 1$. For the final convex polygon P^*, we have $D(P^*) = 2s$ and $B(P^*) \geq s$, so the final potential $\Phi(P^*)$ is at least $5s/2$. By Lemmas 4.1 and 4.2, every flipturn reduces the potential by at least one. Thus, after any sequence of $\lceil \Phi(P^*) - \Phi(P) \rceil = \lceil ns - n/2 - 5s/2 - 1 \rceil$ flipturns, the polygon must be convex. □

If we set $s = n$, we obtain an upper bound $n^2 - 3n - 1$ for arbitrary simple polygons. However, if $s = n$, there can be no degenerate flipturns, so the discrete angle results from Ahn *et al.* apply directly, giving us the upper bound $n(n-3)/2$. Hence, the actual worst case arises when $s = n - 1$.

Corollary 4.4. *Every simple polygon is convexified after any sequence of $n^2 - 4n + 2$ standard or extended flipturns.*

We can improve our results in some cases using a different definition of discrete angle. Let T denote the set of edge directions (*without* their reversals), let $t = |T|$, and let $h \leq t - 1$ be the maximum number of edge directions that fit in an open half-circle. Now define the discrete angle at a vertex $e \cap e'$ to be one more than the number of elements of T in the open interval between $\vec{e}$ and $\vec{e}'$. The discrete angle at any vertex is at most $h - 1$, so $D(P) \leq n(h-1) \leq n(t-2)$ for any polygon P; if P is convex, then $D(P) = t$. Lemma 4.1 still holds under this new definition. Thus, we obtain the following upper bounds.

Theorem 4.5. *Every simple polygon is convexified after any sequence of* $\lceil (nh - n - t)/2 \rceil \leq \lceil t(n-1)/2 \rceil - n$ *modified flipturns or* $nh - \lfloor (n+3t)/2 \rfloor - 1 \leq nt - \lfloor 3(n+t)/2 \rfloor - 1$ *standard or extended flipturns.*

This theorem improves all earlier results whenever h is significantly smaller than t. For arbitrary simple polygons, we have $h \leq t-1 \leq n-1$. Setting $h = t-1 = n-1$ gives us the same $n(n-3)/2$ upper bound for modified flipturns. For standard or extended flipturns, however, we obtain a very slight improvement since the worst case is actually $h = t - 1 = n - 2$.

Corollary 4.6. *Every simple polygon is convexified after any sequence of* $n^2 - 4n + 1$ *standard or extended flipturns.*

5 Extreme Orders are Hard to Find

In this section, we show that two problems related to computing optimal flipturn sequences are NP-hard: computing the shortest convexifying sequence of *generalized* flipturns (defined below), and computing the longest sequence of standard or extended flipturns. The complexity of computing the shortest convexifying sequence of standard flipturns remains an open problem.

A *generalized flipturn* rotates a contiguous chain of edges 180 degrees around the midpoint of its endpoints without introducing self-intersections. The endpoints of the rotated chain need not lie on a convex hull edge. Generalized flipturns include standard, extended, and modified flipturns as special cases.

Theorem 5.1. *Computing the shortest sequence of generalized flipturns that convexifies a simple polygon is NP-hard.*

Proof: We prove the theorem by relating flipturn sequences to the problem of *sorting by reversals.* Here, the input is a permutation $(\pi_1, \pi_2, \ldots, \pi_n)$, and a reversal is any permutation of the form $(1, 2, \ldots, i-1, j, j-1, \ldots, i+1, i, j+1, \ldots, n)$ for some $i < j$. Any permutation can be sorted—that is, transformed into the identity permutation $(1, 2, \ldots, n)$—by applying at most $n - 1$ reversals, and this is tight in the worst case [4]. Caprara [6] proved that computing the minimum number of reversals required to sort a permutation is NP-hard.

Given a permutation $(x_1, x_2, \ldots, x_n)$, we construct a simple polygon P as follows. We start with a right isosceles triangle whose horizontal hypotenuse has been subdivided into n equal fragments. We then perturb the vertices of these fragments vertically, so that for each i, the slope of the ith fragment is $(2x_i - n - 1)/2n^2$. Convexifying this polygon is equivalent to sorting the fragments by slope. No generalized flipturn can separate the two long diagonal edges without introducing a self-intersection, so every legal generalized flipturns is a reversal of the permutation of fragments and vice versa. It follows immediately that the shortest generalized flipturn sequence that convexifies P corresponds precisely to the shortest sequence of reversals that sorts the input permutation. $\square$

Computing the shortest generalized flipturn sequence is trivial for orthogonal polygons; any sequence of $(n-4)/2$ modified flipturns is a solution. Computing longest flipturn sequences, on the other hand, is NP-hard even for orthogonal polygons.

Theorem 5.2. *Computing the longest sequence of standard or extended flipturns for an orthogonal polygon is NP-hard.*

Proof: A flipturn sequence for an orthogonal polygon has length greater than $(n-4)/2$ if and only if it contains an orthogonal flipturn. Thus, to prove the theorem, we only need to show the NP-hardness of the decision problem ORTHOGONAL FLIPTURN: Given an orthogonal polygon, does *any* flipturn sequence contain an orthogonal flipturn? We prove that this problem is NP-complete by a reduction from SUBSET SUM: Given a set of positive integers $A = \{a_1, a_2, \ldots, a_n\}$ and another integer T, does any subset of A sum to T?

The reduction algorithm is given in Figure 5.1, and an example of its output is shown in Figure 5.2. The algorithm constructs a polygon in linear time by walking along its edges in clockwise order, starting and ending at the top of the first step. (The algorithm assumes without loss of generality that n is even.) The basic structure of the polygon is a staircase, with one square step for each of the a_i, plus one long step of height T splitting the other steps in half. Just below each of the upper steps is an an inward horizontal spike; just above each of the lower steps is an outward horizontal spike; and just behind the long step is a vertical *test spike* of length exactly T. The horizontal spikes all have length greater than T, and they increase in length as they get closer to the top and bottom of the polygon.

```
SUBSETSUM(A, T) ↦ ORTHOGONALFLIPTURN:
  ⟨⟨Upper steps and inward spikes⟩⟩
  for i ← 1 to n/2
      SOUTH(a_{2i−1}); EAST(a_{2i−1}); SOUTH(1);
      WEST(T + 2n − 4i + 4); SOUTH(1); EAST(T + 2n + 4i − 4)

  ⟨⟨Test spike⟩⟩
  SOUTH(T + 2); EAST(1); NORTH(T); EAST(1); SOUTH(T + 1); WEST(2);

  ⟨⟨Lower steps and outward spikes⟩⟩
  for i ← 1 to n/2
      SOUTH(1); EAST(a_{n−2i+2}); SOUTH(a_{n−2i+2})
      EAST(T + 4i + 2); SOUTH(1); WEST(T + 4i + 2);

  ⟨⟨Close off the polygon⟩⟩
  Σ ← ∑_{i=1}^{n} a_i
  WEST(T + Σ + 2n + 2); NORTH(T + Σ + 2n + 3); EAST(T + 2n + 2)
```

Figure 5.1. The algorithm to reduce SUBSETSUM to ORTHOGONAL FLIPTURN.

At any point during the flipturning process, the polygon has one main pocket containing the test spike and several secondary pockets containing one or more steps. Initially, there is only one secondary pocket, containing only the step with height a_1. For each i, the step with height a_i is exposed after the main pocket is flipturned $i-1$ times. No matter which flipturns we perform before flipturning the test spike, the vertical distance Δ between the top endpoint of the main pocket's lid and the top edge of the polygon's bounding box is always the sum of elements of A. Specifically, if we flipturn every step whose height is an element of some subset $B \subseteq A$ as soon as it

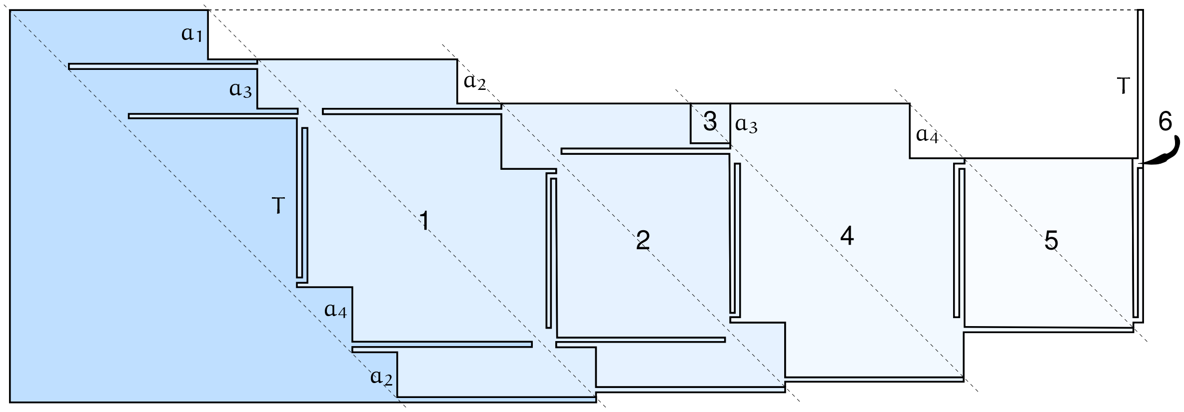

Figure 5.2. A sample reduction from SUBSET SUM to ORTHOGONAL FLIPTURN. The steps store the set $\{a_1, a_2, a_3, a_4\}$ and the vertical spike stores the target sum $T = a_1 + a_2 + a_4$. If we flipturn the step of height a_3 as soon as possible (flipturn 3) and leave the other steps alone, then flipturning the test spike (flipturn 6) creates an orthogonal pocket.

becomes available, then just before the test spike is flipped, Δ is the sum of the elements of $A \setminus B$; see Figure 5.2. Thus, since the test spike has length T, flipturning it can create an orthogonal pocket if and only if some subset of A sums to T. □

Note that the polygon produced by our reduction never has more than one orthogonal pocket; the longest flipturn sequence has either $(n-4)/2$ or $(n-2)/2$ flipturns. Thus, even approximating the maximum number of orthogonal flipturns is NP-hard.

6 Order Doesn't Matter

Joss and Shannon showed that any simple polygon can be transformed into a convex polygon by a sufficiently long sequence of flipturns. As we observed in the introduction, every flipturn sequence results in the same convex shape. We can easily compute this shape in $O(n \log n)$ time by sorting the edges of P by their orientation. (If the polygon has parallel edges, there may be several possible sorted edge sequences, but they all describe the same convex shape.) For s-oriented polygons, this requires only $O(n \log s)$ time.

In this section, we show that the *position* of the final convex polygon is also independent of the flipturn sequence. To prove this result, we show how to predict the y-coordinate of the top edge of the final convex polygon's bounding box. The position of the left edge follows from a symmetric argument, and these two edges determine the polygon's final position. Our result actually holds for arbitrary sequences of generalized flipturns, and therefore for all three specific flipturn types. Recall from the previous section that a generalized flipturn rotates a chain of edges 180 degrees around the midpoint of its endpoints without introducing self-intersections. Following our earlier usage, we call the rotated chain a *pocket* and the segment between its endpoints the *lid*.

Consider a horizontal trapezoidal decomposition of the exterior of a polygon P, obtained by casting rays left and right from every vertex. We classify the trapezoids in this decomposition into several groups. If a region is unbounded, we call it an *outer region*; otherwise, we call it an *inner region*. We further classify outer regions into the infinite *strips* above or below P (including the top and bottom halfplanes), and the semi-infinite *side regions* to the left or right of P. We also classify inner regions as *up-regions* and *down-regions* as follows. Consider the shortest path through the exterior of P from a point in the interior of an inner region ρ to a point at infinity. If the first segment of this path goes up from the starting point, ρ is an up-region; otherwise, ρ is a down-region. We emphasize that this classification is independent of the starting point within ρ.

Our key insight is that the total height of the up-regions is precisely the distance between the top of the current polygon's bounding box and the top of the final convex polygon's bounding box. Specifically, let U denote the sum of the heights of the up-regions, and let $\hat{y}$ denote the y-coordinate of the top vertex of P. We prove our main result inductively, by showing that the quantity $U + \hat{y}$ is an invariant preserved by any generalized flipturn.

Our proof uses the following refinement of the trapezoidal decomposition. Let ab be a lid of some pocket in P, and let c be the midpoint of ab. We subdivide the plane into horizontal strips using the horizontal line ℓ_0 through c, the horizontal lines L passing through every vertex of P, and the reflection L' of L across ℓ_0. Number the strips $1, 2, 3, \ldots$ counting upwards from ℓ_0 and $-1, -2, -3, \ldots$ counting downwards from ℓ_0. With this numbering, any strip i is the reflection of strip $-i$ across ℓ_0. In particular, strips i and $-i$ have the same width, which we denote w_i. There are at most $2n + 2$ strips altogether. These strips subdivide the exterior of the polygon into trapezoids, which we classify as up-regions, down-regions, strips, and side regions exactly as above. Notice that refining the trapezoids has no effect on the total height of the up-regions.

For each $i > 0$, let u_i denote the number of up-regions in strips i and $-i$, and let x_i be the indicator variable equal to 1 if strip i intersects P and 0 otherwise. See Figure 6.1(a). Let P' be the result of flipturning the pocket ab. This flipturn moves any point on the boundary of the pocket from some strip i to the corresponding strip $-i$. The strips subdivide the exterior of P' into regions exactly as the exterior of P, and we define the corresponding variables u'_i and x'_i mutatis mutandis. See Figure 6.1(b).

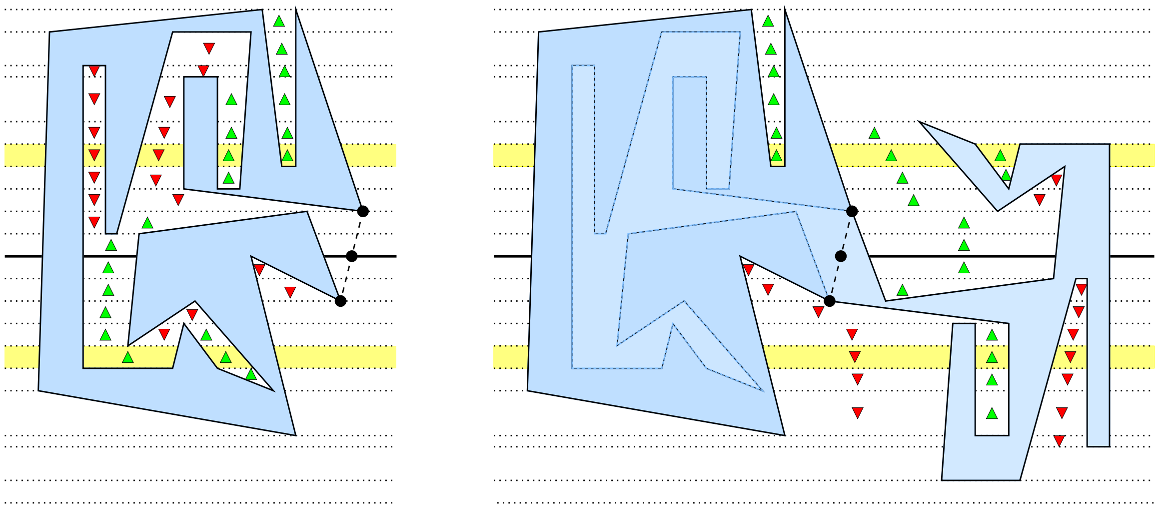

(a) (b)

Figure 6.1. Strips defined by a polygon and one of its pockets. Strips 4 and -4 are highlighted. Triangles indicate up-regions and down-regions. (a) The original polygon P, with $u_4 = 3$ and $x_4 = 1$. (b) The flipturned polygon P', with $u'_4 = 3$ and $x'_4 = 1$.

Our core lemma is the following.

Lemma 6.1. $u_i + x_i = u'_i + x'_i$ *for all* i.

Proof: Fix an index $i > 0$; for the sake of readability we omit the subscript i from all our notation. We prove the theorem by induction on the number of inner regions in the flipturned pocket. If the pocket contains no inner regions, it must be y-monotone. Flipturning such a pocket changes neither the number of up-regions nor the height of the polygon, so $u = u'$ and $x = x'$.

The inner regions of P have a natural forest structure, defined by connecting each region to the next region encountered on a shortest path to infinity. The roots of this forest are inner regions directly adjacent to outer regions, and its leaves are inner regions adjacent to only one other region. We define a simpler polygon $\tilde{P}$ by filling in some leaf region ρ inside the pocket ab; more formally, $\tilde{P} = P \cup \rho$. Let $\tilde{P}'$ be the result of flipturning the now-simpler pocket ab of $\tilde{P}$, and let ρ' be the image of ρ under this flipturn (so that $\tilde{P}' = P' \setminus \rho'$). Finally, define $\tilde{u}$, $\tilde{u}'$, $\tilde{x}$, and $\tilde{x}'$ analogously to u, u', x, and x' for these new polygons. The inductive hypothesis implies that $\tilde{u} + \tilde{x} = \tilde{u}' + \tilde{x}'$.

It suffices to consider the case where ρ lies either in strip i or in strip $-i$, since otherwise we immediately have $\tilde{u} = u$, $\tilde{x} = x$, $\tilde{u}' = u'$, and $\tilde{x}' = x'$.

Suppose ρ is an up-region. Since $\tilde{P}$ has one fewer up-region in strip i than P, we have $\tilde{u} = u - 1$. Some region σ of $\tilde{P}'$ is split into two regions by ρ'. If we imagine a continuous transformation from $\tilde{P}'$ to P', the trapezoid ρ' grows upward from the bottom edge of σ.

To express u' in terms of $\tilde{u}'$, we consider four cases, illustrated in Figure 6.2. If σ is an up-region, then ρ' splits it into two up-regions, so $u' = \tilde{u}' + 1$. If σ is a down-region, then ρ' splits it into an up-region and a down-region, so $u' = \tilde{u}' + 1$. If σ is a side region, then ρ' splits it into an up-region and a side region, so $u' = \tilde{u}' + 1$. Finally, if σ is a strip, ρ' splits σ into two side regions, so $u' = \tilde{u}'$.

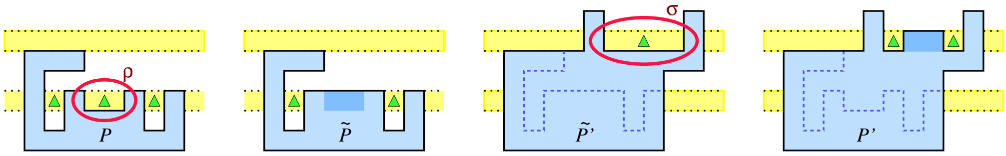

(a) σ is an up-region

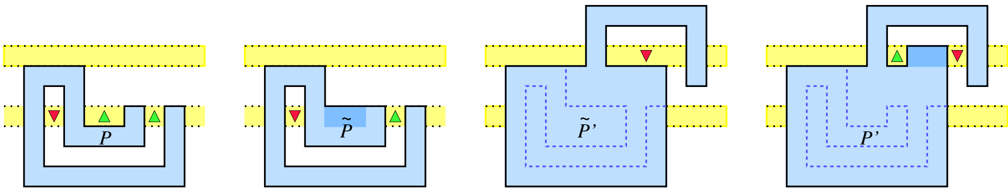

(b) σ is a down-region

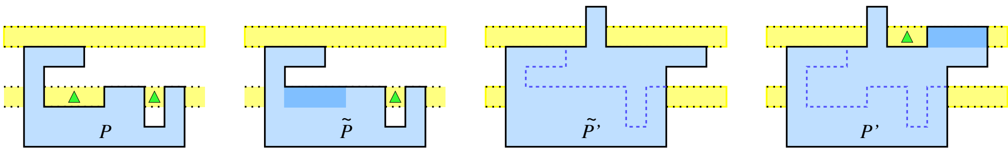

(c) σ is a side-region

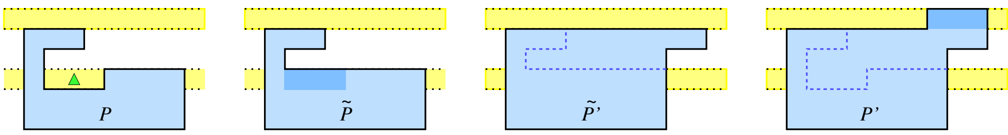

(d) σ is a strip

Figure 6.2. Cases for the proof of Lemma 6.1 where ρ is an up-region in strip $-i$.

The values of the indicator variables $x, x', \tilde{x}, \tilde{x}'$ depend on whether the leaf region ρ lies in strip i or in strip $-i$. If ρ is in strip i, then we immediately have $x = \tilde{x} = x' = \tilde{x}' = 1$. If ρ is in strip $-i$ and σ is not a strip, then $\tilde{x}' = x' = 1$ and $x = \tilde{x}$ (but these might be either 0 or 1). Finally, if σ is a strip, then ρ' must lie above $\tilde{P}'$ in strip i (because ρ is an up-region) and therefore $x' = 1$ and $x = \tilde{x} = \tilde{x}' = 0$.

We conclude that if ρ is an up-region, then $u' + x' = \tilde{u}' + \tilde{x}' + 1 = \tilde{u} + \tilde{x} + 1 = u + x$, as required.

A similar (and slightly simpler) case analysis applies when ρ is a down-region. In each case, we have $\tilde{u} = u$, $\tilde{u}' = u'$, $\tilde{x} = x$, and $\tilde{x}' = x'$. We omit further details. □

Theorem 6.2. *The final convexified position of a polygon is independent of the convexifying generalized flipturn sequence. Moreover, this position can be determined in $O(n)$ time.*

Proof: Let w_i denote the vertical width of strip i (and strip $-i$). Lemma 6.1 implies that

$$\sum_{i>0} (u_i + x_i) w_i = \sum_{i>0} (u'_i + x'_i) w_i. \tag{1}$$

Let $\hat{y}$ and $\hat{y}'$ denote the y-coordinates of the top of P and P', respectively, and let y_0 be the y-coordinate of the lid midpoint c. We easily observe that

$$\sum_{i>0} x_i w_i = \hat{y} - y_0 \qquad \text{and} \qquad \sum_{i>0} x'_i w_i = \hat{y}' - y_0. \tag{2}$$

Finally, define $U = \sum_{i>0} u_i w_i$ and $U' = \sum_{i>0} u'_i w_i$. Combining equations (1) and (2), we obtain the identity $U + \hat{y} = U' + \hat{y}'$. In other words, the total height of all the up-regions plus the maximum y-coordinate of the polygon is an invariant preserved by any generalized flipturn.

Let P^* be the convex polygon produced by some sequence of generalized flipturns starting from P, and define U^* and $\hat{y}^*$ analogously to U and $\hat{y}$. Obviously, P^* has no up-regions, so $U^* = 0$. Thus, by induction on the number of flipturns, we have the identity $\hat{y}^* = U + \hat{y}$. Since $U + \hat{y}$ is independent of the convexifying flipturn sequence, so is the vertical position of P^*.

We can compute U in linear time by computing a horizontal trapezoidal decomposition of P, using Chazelle's algorithm [8] or the more recent randomized variant by Amato, Goodrich, and Ramos [3], and then performing a depth-first search of its dual graph.

The argument for the horizontal position of P^* is symmetric. □

7 Conjectures and Open Problems

We have proven several upper and lower bounds on the lengths of shortest and longest flipturn sequences for several types of polygons. Most of our upper and lower bounds match within small constant factors, but there is still considerable room for improvement.

Perhaps the easiest open problem is to improve our $\lfloor 5(n-4)/6 \rfloor$ upper bound on the shortest flipturn sequence for orthogonal polygons. Our upper bound proof (Theorem 3.7) uses an algorithm that prefers diagonal flipturns to orthogonal flipturns and good orthogonal flipturns to bad orthogonal flipturns. Our example polygon P_n from Theorem 3.10 can be used to show that this algorithm is not always optimal. If we modify our algorithm to ignore rectangular 'notches' in the corners of the polygon, we can convexify P_n with less than $2n/3$ flipturns. (See the expanded version of this paper for further details [2].) The ignored notches are precisely the diagonal flipturns that do not remove brackets; see Theorem 3.8. A modified algorithm that tries to reduce the number of brackets as quickly as possible, as well as the number of vertices, might lead to a tighter upper bound.

Asymptotically, our upper bounds for simple polygons agree with Joss and Shannon's original conjecture [15]—any polygon can indeed be convexified by $O(n^2)$ flipturns—but there is still a significant gap between our upper bounds and the $(n-2)^2/4$ lower bound construction of Biedl [5]. We, like Joss and Shannon, conjecture that the correct answer is closer to $n^2/4$.

A more interesting open question concerns the length of *shortest* flipturn sequences for general polygons. The best lower bounds are those derived for orthogonal polygons in Section 3, but not subquadratic upper bounds are known. Can arbitrary polygons be convexified with only $O(n)$ flipturns, or does some polygon require a super-linear number of flipturns to convexify?

Finally, how hard is it to find the shortest sequence of standard or extended flipturns that convexifies a given simple polygon? We conjecture that this problem is NP-hard, even for orthogonal polygons. It would be surprising if the additional constraint of flipturning only across convex hull edges makes the optimization problem significantly easier. One possible solution would be to prove that the following problem is co-NP-hard: Given an orthogonal polygon, does every flipturn sequence contain at least one orthogonal flipturn?

Acknowledgments

All but one of the authors thank Godfried Toussaint for organizing the Barbados workshop where most of this work was done. Thanks also to David Bremner, Ferran Hurtado, Vera Sacristán, and Mike Soss for sharing coffee, rum, and ideas.

References

[1] H.-K. Ahn, P. Bose, J. Czyzowicz, N. Hanusse, E. Kranakis, and P. Morin. Flipping your lid. *Geombinatorics* X(2):57–63, 2000. ⟨http://www.scs.carleton.ca/˜morin/publications/linkage/flipturn-tr.ps⟩.

[2] O. Aichholzer, C. Cortés, E. D. Demaine, V. Dujmović, J. Erickson, H. Meijer, M. Overmars, B. Palop, S. Ramaswami, and G. T. Toussaint. Flipturning polygons. Technical Report UU-CS-2000-31, Department of Computer Science, Utrecht University, 2000. arXiv:cs.CG/0008010.

[3] N. M. Amato, M. T. Goodrich, and E. A. Ramos. Linear-time polygon triangulation made easy via randomization. *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, 201–212, 2000.

[4] V. Bafna and P. A. Pevzner. Genome rearrangements and sorting by reversals. *SIAM J. Comput.* 25(2): 272–289, 1996.

[5] T. Biedl. Polygons needing many flipturns. Technical Report CS-2000-04, Department of Computer Science, University of Waterloo, January 2000. ⟨ftp://cs-archive.uwaterloo.ca/cs-archive/CS-2000-04/⟩.

[6] A. Caprara. Sorting permutations by reversals and Eulerian cycle decompositions. *SIAM J. Discrete Math.* 12(1):91–110, 1999.

[7] T. M. Chan. Dynamic planar convex hull operations in near-logarithmic amortized time. *Proc. 40th Annu. IEEE Sympos. Found. Comput. Sci.*, 92–99, 1999.

[8] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete Comput. Geom.* 6(5):485–524, 1991.

[9] B. Chazelle and D. P. Dobkin. Intersection of convex objects in two and three dimensions. *J. ACM* 34(1):1–27, 1987.

[10] L. E. Dubins, A. Orlitsky, J. A. Reeds, and L. A. Shepp. Self-avoiding random loops. *IEEE Trans. Inform. Theory* 34:1509–1516, 1988.

[11] H. Edelsbrunner. Computing the extreme distances between two convex polygons. *J. Algorithms* 6:213–224, 1985.

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

[13] P. W. Finn, D. Halperin, L. E. Kavraki, J.-C. Latombe, R. Motwani, C. Shelton, and S. Venkatasubramanian. Geometric manipulation of flexible ligands. *Applied Computational Geometry*, pp. 67–78. Lecture Notes Comput. Sci. 1148, Springer-Verlag, 1996.

[14] M. D. Frank-Kamenetskii. *Unravelling DNA.* Addison-Wesley, 1997.

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

[16] B. Grünbaum and J. Zaks. Convexification of polygons by flips and by flipturns. *Discrete Mathematics* 241(1–3): 333–342, 2001.

[17] L. J. Guibas and R. Sedgewick. A dichromatic framework for balanced trees. *Proc. 19th Annu. IEEE Sympos. Found. Comput. Sci.*, 8–21, 1978.

[18] N. Madras and G. Slade. *The Self-Avoiding Walk.* Birkhäuser, Boston, 1993.

[19] N. Madras and A. D. Sokal. The pivot algorithm: A highly efficient Monte Carlo method for the self-avoiding walk. *Journal of Statistical Physics* 50:109–186, 1988.

[20] F. M. McMillan. *The Chain Straighteners.* The MacMillan Press, 1979.

[21] J. O'Rourke. Folding and unfolding in computational geometry. *Discrete and Computational Geometry (Proc. JCDCG '98)*, pp. 142–147. Lecture Notes Comput. Sci. 1763, Springer-Verlag, 2000.

[22] M. H. Overmars and J. van Leeuwen. Maintenance of configurations in the plane. *J. Comput. Syst. Sci.* 23:166–204, 1981.

[23] E. J. J. 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.

[24] R. Seidel and C. R. Aragon. Randomized search trees. *Algorithmica* 16:464–497, 1996.

[25] D. D. Sleator and R. E. Tarjan. Self-adjusting binary search trees. *J. ACM* 32(3):652–686, 1985.

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

[27] G. T. Toussaint. Computational polygonal entanglement theory. *Proc. VIII Encuentros de Geometría Computacional*. July 7-9, 1999. ⟨http://www-cgrl.cs.mcgill.ca/~godfried/publications/castellon.ps.gz⟩.

[28] G. T. Toussaint. The Erdős-Nagy theorem and its ramifications. *11th Canadian Conf. Comput. Geom.*, 9–12, 1999. ⟨http://www.cs.ubc.ca/conferences/CCCG/elec_proc/fp19.ps.gz⟩.