

Flipping Edge-Labelled Triangulations

Prosenjit Bose^a, Anna Lubiw^b, Vinayak Pathak^b, Sander Verdonschot^a

^a*School of Computer Science, Carleton University, Ottawa ON, Canada*

^b*School of Computer Science, University of Waterloo, Waterloo ON, Canada*

Abstract

Flips in triangulations have received a lot of attention over the past decades. However, the problem of tracking where particular edges go during the flipping process has not been addressed. We examine this question by attaching unique labels to the triangulation edges. We introduce the concept of the *orbit* of an edge e , which is the set of all edges reachable from e via flips.

We establish the first upper and lower bounds on the diameter of the flip graph in this setting. Specifically, we prove tight $\Theta(n \log n)$ bounds for edge-labelled triangulations of n -vertex convex polygons and combinatorial triangulations, contrasting with the $\Theta(n)$ bounds in their respective unlabelled settings. The $\Omega(n \log n)$ lower bound for the convex polygon setting might be of independent interest, as it generalizes lower bounds on certain sorting models. When simultaneous flips are allowed, the upper bound for convex polygons decreases to $O(\log^2 n)$, although we no longer have a matching lower bound.

Moving beyond convex polygons, we show that edge-labelled triangulated polygons with a single reflex vertex can have a disconnected flip graph. This is in sharp contrast with the unlabelled case, where the flip graph is connected for any triangulated polygon. For spiral polygons, we provide a complete characterization of the orbits. This allows us to decide connectivity of the flip graph of a spiral polygon in linear time. We also prove an upper bound of $O(n^2)$ on the diameter of each connected component, which is optimal in the worst case. We conclude with an example of a non-spiral polygon whose flip graph has diameter $\Omega(n^3)$.

Keywords: Flip, Convex Polygon, Spiral Polygon, Quadrilateral Graph, Simultaneous Flip

Email addresses: jit@scs.carleton.ca (Prosenjit Bose), alubiw@uwaterloo.ca (Anna Lubiw), vpathak@uwaterloo.ca (Vinayak Pathak), sander@cg.scs.carleton.ca (Sander Verdonschot)

We dedicate this paper to Ferran Hurtado, whose creativity in research and seminal work on flips continues to inspire us all. A preprint of this work can be found on arXiv.org [7].

1. Introduction

The problem of reconfiguring one triangulation to another is well-studied for combinatorial triangulations, triangulations of points sets, and triangulations of polygons. In all these settings, the basic reconfiguration operation is the *flip* operation that removes one edge of the triangulation and adds another edge to obtain a new triangulation. Not every edge can always be flipped. The constraints on the edges involved in a flip depend on the setting. The geometric setting refers to a triangulation of a point set (a maximal set of non-crossing edges joining pairs of points) or polygon (a maximal set of non-crossing chords inside the polygon). In both cases, a flip removes one diagonal of a convex quadrilateral and replaces it by the other diagonal. In the combinatorial setting, a triangulation is a maximal planar graph with the clockwise order of edges around each vertex specified. Here, a flip removes one edge and replaces it by the other diagonal of the resulting quadrilateral, as long as the new edge is not already an edge of the triangulation. When we replace an edge e by an edge f with a flip, we say that e *flips to* f .

Another way to view reconfigurations via flips is through the *flip graph*, which has a vertex for each triangulation and an edge between two vertices if their corresponding triangulations differ by a single flip. The most important property of flips is that they can be used to reconfigure any triangulation into any other triangulation of the same point set, polygon, or vertex set—that is, the flip graph is connected.

Wagner [28] initiated the study of flips in 1936, by proving that the flip graph is connected in the combinatorial setting. For points sets, connectivity of the flip graph is a consequence of Lawson’s result [21] that any triangulation of a point set can be flipped to the Delaunay triangulation, which then acts as a “canonical” triangulation from which every triangulation can be reached. For triangulations of a polygon, the constrained Delaunay triangulation can be used in the same way to prove that the flip graph is connected [3]. In Section 2, we discuss these results in more detail, along with bounds on the number of flips required, i.e., bounds on the diameter of the flip graph.

Reconfiguration of triangulations via flips is of theoretical interest in the study of associahedra [26] and mixing [22]. It also has more practical benefits in mesh generation and for finding triangulations that optimize certain quality measures [3, 11]. For a broader overview of the literature on flips, we refer the reader to a survey by Bose and Hurtado [5].

Despite the extensive literature on flips in many different settings, to our knowledge, no one has specifically studied where individual edges move during the course of a sequence of flips in any setting. We say that an edge f is *reachable* from an edge e if there is a sequence of edges e, \dots, f , such that for every two consecutive edges a and b in the sequence, there exists a triangulation in which a flips to b . We define the *orbit* of an edge e to be the set of all edges reachable from e . This gives rise to a very natural question: What is the orbit of an edge? Although every triangulation can be flipped to every other, the orbits of edges can be very different. For example, some

polygons have a unique triangulation, in which case the orbits are all singletons. At the other extreme, a convex polygon has a single orbit containing all the diagonals. In the geometric setting, the orbits are exactly the connected components of a graph introduced
45 by Eppstein [12] as the *quadrilateral graph*. This graph has a vertex for every diagonal and an edge between two vertices if their corresponding diagonals e and f intersect and their endpoints form a convex quadrilateral that is empty.

An intriguing question is how orbits combine. In other words, we want to track where multiple edges go during a sequence of flips. To address this question, we study
50 flips in *edge-labelled triangulations*: triangulations of a point set, polygon, or vertex set where each edge has a unique label from $\{1, \dots, m\}$, where m is the number of edges. If we flip an edge of an edge-labelled triangulation, the new edge is assigned the label of the removed edge. In particular, this means that the set of edge labels is preserved throughout any flip sequence. In the geometric setting, we often omit the labels on the
55 edges of the convex hull or on the boundary of the polygon, as these can never be flipped.

Given two edge-labelled triangulations T and T' of the same point set, polygon, or vertex set, it is not always possible to transform T to T' via a sequence of flips. For example, when a polygon has a unique triangulation, no flips are possible, so if the labellings of T and T' differ, no flip sequence can transform one to the other. A necessary
60 condition for such a flip sequence to exist is that each edge with label i in T' must be in the orbit of the edge with label i in T . We conjecture that this condition is also sufficient.

Orbit Conjecture. *Given two edge-labelled triangulations, there is a sequence of flips that transforms one into the other if and only if for every label, the initial and final edge with that label lie in the same orbit.*

65 The Orbit Conjecture is known to hold in some cases. A result by Eppstein [12] implies that the Orbit Conjecture holds for triangulations of point sets with no empty pentagon. These point sets are highly degenerate since every set of ten points in general position contains an empty pentagon. Eppstein showed that in this case every orbit is a tree, and each triangulation contains exactly one edge from each orbit. There are a
70 few other results of a similar flavour. Cano et al. [9] proved the analogous statement for edge-labelled non-maximal plane graphs, although instead of flips, they “rotate” edges around one of their endpoints. Hernando et al. [17] proved the analogous statement for edge-labelled spanning trees of an underlying graph, where the “flip” operation consists of removing an edge and replacing it elsewhere with the same label while maintaining
75 connectivity. They showed that the orbits are exactly the 2-connected components of the underlying graph. Pathak [23] generalized this result to labelled bases of a matroid.

Edge-labelled flips in triangulations of a convex polygon were independently studied by Araujo-Pardo et al. [1], who mainly focused on the combinatorial properties of the flip graph. In the unlabelled setting, the flip graph is known to be isomorphic to the
80 1-skeleton of a polyhedron, called the associahedron. Araujo-Pardo et al. proved that this is also the case for the edge-labelled setting, and called the resulting polyhedron the colourful associahedron. As part of their proof, they show that the flip graph of

edge-labelled triangulations of a convex polygon is connected, which implies the Orbit Conjecture for this setting. A careful analysis of their argument shows that it gives
85 an upper bound on the diameter of the flip graph that is quadratic in the number of vertices of the polygon. In addition, they show that the colourful associahedron covers the regular associahedron and fully characterize its automorphism group.

In this paper we address the Orbit Conjecture for edge-labelled triangulations of polygons and edge-labelled combinatorial triangulations. We give an alternative proof
90 for the connectivity of the flip graph of edge-labelled triangulations of an n -vertex convex polygon that shows that the diameter is $O(n \log n)$, and we provide a matching $\Omega(n \log n)$ lower bound, which contrasts with the $\Theta(n)$ bound for unlabelled flips [26].

We also consider what happens when we allow multiple edges to be flipped simultaneously. We prove that the simultaneous flip distance between two edge-labelled
95 triangulations of a convex polygon is $O(\log^2 n)$, in contrast with the $\Theta(\log n)$ bound established for the unlabelled setting [14].

After convex polygons, we consider spiral polygons (polygons with at most one reflex chain). These may have multiple orbits, so the flip graph is not necessarily connected. We again prove the Orbit Conjecture, and give a tight $\Theta(n^2)$ worst case bound on the
100 diameter of each connected component of the flip graph. We also characterize the orbits, and show how to test in linear time whether there is a flip sequence that transforms one given edge-labelled triangulation into another.

Turning to more general polygons, we give an example of a polygon with two reflex chains whose flip graph is connected but has diameter $\Omega(n^3)$. This is in stark contrast to
105 the $\Theta(n^2)$ bound on the diameter of the flip graph in the unlabelled setting [3, 19], and the $\Theta(n^2)$ bound for spiral polygons.

Finally, we consider the case of n -vertex edge-labelled combinatorial triangulations. As for convex polygons, we prove the Orbit Conjecture by showing that the flip graph is connected and has diameter $\Theta(n \log n)$, contrasting with the $\Theta(n)$ bound for the
110 unlabelled case [27].

2. Background

The original result on flips was due to Wagner [28], who showed that $O(n^2)$ flips were sufficient to transform between any two given n -vertex combinatorial triangulations. In particular, he showed that any triangulation can be transformed into one with two
115 dominant vertices, often referred to as *Wagner's canonical triangulation*. The upper bound was first improved to linear by Sleator et al. [27]—this is optimal, since converting a triangulation with constant maximum degree into Wagner's canonical triangulation immediately gives a linear lower bound. Currently, the best upper bound on the diameter of the flip graph is $5n - 23$ [10]. See Bose and Verdonschot [8] for a more detailed
120 history of this problem in the combinatorial setting.

Combinatorial triangulations with vertex labels were studied by Sleator et al. [27] as an intermediate form between unlabelled combinatorial triangulations and triangulations

of a set of points in the plane. In this setting, they showed that $O(n \log n)$ flips are sufficient to transform one vertex-labelled triangulation with n vertices into any other, and that $\Omega(n \log n)$ flips are sometimes necessary. Note that by transforming both the initial and final triangulation into Wagner’s canonical form without paying attention to vertex labels, the problem becomes one of sorting vertex labels. This is essentially what they do to prove the $O(n \log n)$ upper bound. For the $\Omega(n \log n)$ lower bound, they show that there are at most 9^{n+m} distinct vertex-labelled triangulations that are reachable from a given triangulation T via m flips. Since there are already $(n - 3)!$ different ways to label the vertices of Wagner’s canonical triangulation, the $\Omega(n \log n)$ lower bound follows.

For n -vertex triangulations of point sets, Lawson [21] proved that given any triangulation, a sequence of $O(n^2)$ flips suffices to transform it to the Delaunay triangulation. For triangulations of an n -vertex polygon, Bern and Eppstein [3] showed that the flip graph is also connected and has diameter $O(n^2)$. Unlike in the combinatorial setting, Hurtado et al. [19] proved that these quadratic bounds are tight by providing an $\Omega(n^2)$ lower bound.

Flips in triangulations of a convex polygon are especially interesting because they correspond exactly to rotations in a binary tree. In this way, Sleator et al. [26] answered a question about the rotation distance between binary trees, by proving a tight bound of $2n - 10$ on the diameter of the flip graph. Recently, Pournin [25] found a purely combinatorial proof of this bound, avoiding the use of hyperbolic geometry in the original proof and establishing the bound even for small values of n .

The idea of performing flips in parallel was introduced by Hurtado et al. [18]. In the geometric setting, a set of edges may be *simultaneously flipped* if each edge is flippable and no two edges are incident to the same triangle. Galtier et al. [14] showed that $\Theta(n)$ simultaneous flips are sufficient and sometimes necessary to reconfigure one triangulation of a point set or simple polygon to another. In the case of convex polygons, they showed that $\Theta(\log n)$ simultaneous flips were sufficient and sometimes necessary. Bose et al. [4] proved that $\Theta(\log n)$ simultaneous flips are sufficient and sometimes necessary in the combinatorial setting—in this case a simultaneous flip may be performed even if some edges in the set cannot be individually flipped.

3. Convex polygons

In this section, we prove that the flip graph of edge-labelled triangulations of a convex polygon with n vertices is connected. Since all diagonals are in the same orbit, this is equivalent to proving the Orbit Conjecture in this setting. Aurajo-Pardo et al. [1] also proved this, but our proof gives a tight bound of $\Theta(n \log n)$ on the diameter of the flip graph. In addition, we show that $O(\log^2 n)$ simultaneous flips are sufficient and $\Omega(\log n)$ simultaneous flips are sometimes necessary.

Given a convex polygon with n vertices, label the vertices p_1, \dots, p_n in counter-clockwise order and give each diagonal a unique label from the set $\{1, \dots, n - 3\}$. Denote

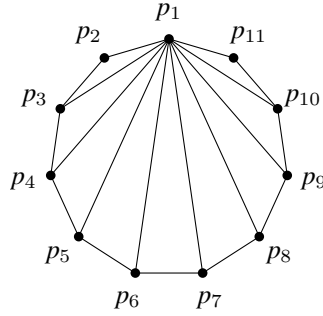


Figure 1: A canonical triangulation with 11 vertices.

by T^* the *canonical* triangulation where vertex p_1 is dominant (see Figure 1). Let ρ be the sequence of edge labels ordered in counter-clockwise order around p_1 , starting with diagonal p_1p_3 . Note that ρ is a permutation of $[1..n-3]$. We denote this canonical triangulation by (T^*, ρ) . The next lemma shows that in this canonical triangulation, ordered subsequences can easily be reversed, which plays an important role for the upper bound. Note the distinction between an *ordered* subsequence and a *contiguous* subsequence. An ordered subsequence is a sequence consisting of a subset of the elements of the original sequence in the same order as in the original. A contiguous subsequence is an ordered subsequence in which the elements are consecutive in the original sequence.

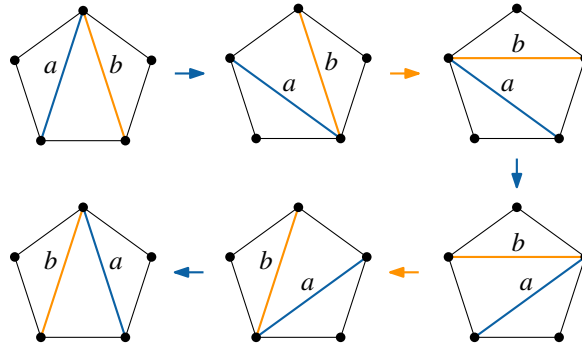


Figure 2: A sequence of five flips that reverses the order of the two diagonals of a pentagon.

Lemma 1. Given (T^*, ρ) , let S be a contiguous subsequence of diagonals adjacent to vertex p_1 , ordered in counter-clockwise order. Any ordered subsequence $R = i_1, \dots, i_r$ of S can be reversed with at most $2.5|S|$ flips.

Proof. Note that the edges of R need not be consecutive. We use induction on $|S|$. In the base case, $|S| = 0$ and we do not need any flips to reverse R , so let $|S| > 0$ and assume that the lemma holds for any contiguous subsequence S' with $|S'| < |S|$.

If there is any edge e in S that is not in R , we flip that edge. Now the remaining edges of S form a contiguous subsequence of diagonals in a smaller convex polygon. Thus, we can reverse them with $2.5(|S| - 1)$ flips by induction. Flipping e back completes the transformation, for a total of $2.5(|S| - 1) + 2 \leq 2.5|S|$ flips.

Now consider the case where $S = R$. If $|R|$ is odd, we can remove the middle diagonal from R , bringing us back to the previous case, so assume that $|R|$ is even. Then the two middle diagonals, $i_{r/2}$ and $i_{r/2+1}$, can be swapped with five flips, as shown in Figure 2. But instead of completing this swap sequence, we halt it after the first three flips. This leaves the edges out of the way of the remaining edges, which we then reverse with $2.5(|S| - 2)$ flips by induction. Afterwards, we perform the final two flips to return $i_{r/2}$ and $i_{r/2+1}$, bringing the total number of flips to $2.5(|S| - 2) + 5 = 2.5|S|$. \square

Using Lemma 1, we can sort ρ with $O(n \log n)$ flips, which gives us the following theorem.

Theorem 1. *Any edge-labelled triangulation of a convex polygon with n vertices can be transformed into any other edge-labelled triangulation using $O(n \log n)$ flips.*

Proof. We first ignore the labels and use $O(n)$ flips to transform both initial triangulations to the canonical triangulations (T^*, l'_1) and (T^*, l'_2) , respectively. Next, we show that for any permutation ρ , the canonical triangulation (T^*, ρ) can be transformed into the triangulation (T^*, ρ^*) using $O(n \log n)$ flips, where ρ^* is the sorted permutation.

We sort the permutation ρ by imitating a version of quicksort that always uses the median as pivot. This ensures that the $O(n \log n)$ bound holds even in the worst case. Let ρ_i denote the i -th element in ρ and consider all $i \in [1..n - 3]$ for which either: (1) $i \leq (n - 3)/2$ and $\rho_i > (n - 3)/2$; or (2) $i > (n - 3)/2$ and $\rho_i \leq (n - 3)/2$. Reverse the subsequence formed by these values of i . By Lemma 1, this can be done with $O(n)$ flips. This ensures that edges whose label is below $(n - 3)/2$ lie in the first half of the sequence and the larger labels lie in the second half. By recursively applying this operation to the two halves, we sort the entire sequence using $O(n \log n)$ flips. \square

As a corollary to the $O(n \log n)$ upper bound, we can get an $O(\log n)$ -factor approximation algorithm for computing the flip-distance.

Corollary 1. *We can approximate in polynomial time the flip-distance between two n -vertex edge-labelled triangulations \mathcal{T}_1 and \mathcal{T}_2 of a convex polygon within an $O(\log n)$ factor of optimal.*

Proof. Call an edge e *fixed* if it satisfies the following properties: a) it occurs in both \mathcal{T}_1 and \mathcal{T}_2 ; b) it has the same label in both \mathcal{T}_1 and \mathcal{T}_2 ; and c) the set of distinct labels that occurs to the left of e in \mathcal{T}_1 is exactly the same as the set that occurs to the left of e in \mathcal{T}_2 . Since every non-fixed edge must be flipped at least once, we need at least k flips, where k is the number of non-fixed edges. The set of fixed edges divides the polygon into several smaller convex polygons, each of which has a different edge-labelled triangulation in

\mathcal{T}_1 and \mathcal{T}_2 . Suppose the i^{th} polygon has n_i diagonals. We use the flip-sequence from Theorem 1 on the i^{th} polygon to perform the required transformation with $O(n_i \log n_i)$ flips. The total number of flips is then $\sum_i O(n_i \log n_i) \leq O(k \log k)$, thus giving an approximation factor of $O(\log n)$. \square

3.1. Lower bound

The lower bound uses a slightly modified version of the $\Omega(n \log n)$ lower bound for the vertex-labelled setting by Sleator, Tarjan, and Thurston [27]. We first give an overview of their technique, before applying it to edge-labelled triangulations of a convex polygon.

Let a *tagged half-edge graph* be an undirected graph with maximum degree Δ , whose vertices have labels called *tags*, and whose edges are split into two *half-edges*. Each half-edge is incident to one endpoint, and labelled with an *edge-end label* in $\{1, \dots, \Delta\}$, such that all edge-end labels incident on a vertex are distinct (see Figure 3a for an example). A *half-edge part* is a half-edge graph in which some half-edges do not have a twin. Note that tags are not restricted to integers: they could be tuples, or even arbitrary strings.

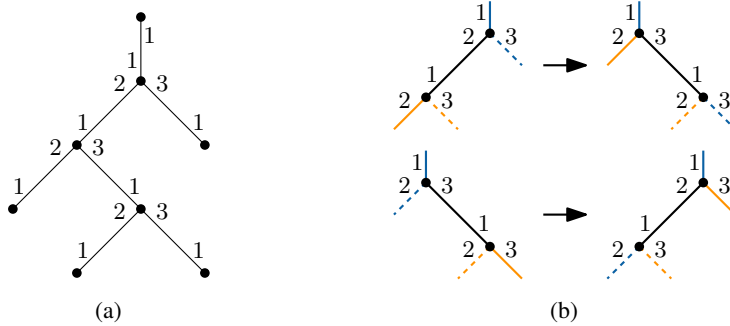


Figure 3: (a) A half-edge graph representation of a rooted binary tree. (b) A graph grammar for rotations in binary trees. Correspondence between half-edges is indicated by a combination of colour and line style.

A *graph grammar* Γ is a sequence of production rules $\Gamma_i = (L, \rightarrow, \mathcal{T}, R)$, where L and R are half-edge parts with the same number of half-edges, \rightarrow is a correspondence between the half-edges of L and R , and \mathcal{T} is a function that computes the tags of vertices in R from those in L . A possible graph grammar for rotations in (unlabelled) binary trees is depicted in Figure 3b.

Sleator, Tarjan, and Thurston prove the following theorem.

Theorem 2 (Sleator, Tarjan, and Thurston [27]). *Let G be a tagged half-edge graph with n vertices, Γ be a graph grammar, c be the number of vertices in left sides of Γ , and r be the maximum number of vertices in any right side of a production of Γ . Then $|R(G, \Gamma, m)| \leq (c + 1)^{n+r \cdot m}$, where $R(G, \Gamma, m)$ is the set of graphs obtainable from G by derivations in Γ of length at most m .*

Recall that, by definition, the maximum degree of a tagged half-edge-graph is
 245 bounded by a constant. Thus, we cannot apply this theorem directly to triangulations of
 a convex polygon. Instead, we turn to the dual graph. The *augmented dual graph* of a
 triangulation of a convex polygon is a tagged half-edge graph G with two sets of vertices:
 triangle-vertices T corresponding to the triangles of the triangulation, and edge-vertices
 E_{CH} corresponding to the boundary edges. One edge-vertex is designated as the *root*.

250 Two triangle-vertices are connected by an edge if their triangles are adjacent. All
 edge-vertices are leaves, each connected to the triangle-vertex whose triangle is incident
 to their corresponding edge (see Figure 4). As every triangle has three edges, the
 maximum degree of G is three. The edge towards the root receives edge-end label 1.
 For a triangle-vertex, the other edge-end labels are assigned in counter-clockwise order,
 255 as in Figure 3a.

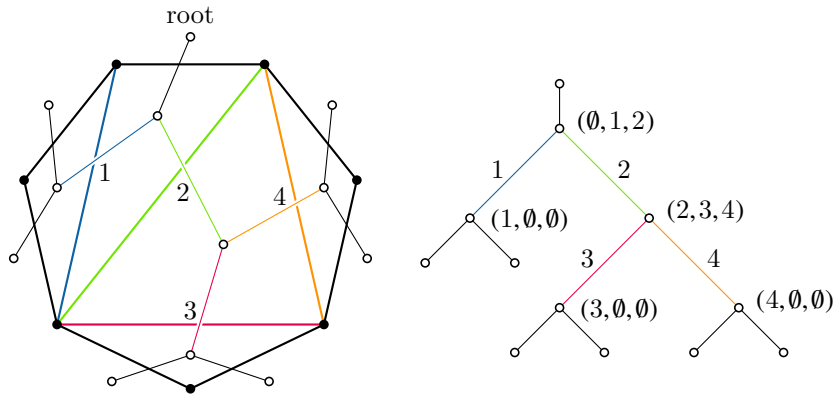


Figure 4: An edge-labelled triangulation of a convex polygon with its augmented dual graph. The edge-labels
 on the dual graph are shown to more clearly indicate the correspondence – they are actually labelled with
 edge-end labels as in Figure 3a.

This is where we deviate slightly from the original paper. Since Sleator, Tarjan,
 and Thurston were working in the vertex-labelled setting, they used the tags in the
 augmented dual graph to encode the labels of the vertices around the corresponding
 triangles. Instead, we use these tags to encode the edge-labels. Specifically, we tag
 260 each triangle-vertex with a triple containing the edge-label of each edge of its triangle,
 starting from the edge closest to the root, and proceeding in counter-clockwise order.
 Edges of the convex hull are assumed to have label \emptyset . Edge-vertices will not be involved
 in any of the production rules, so they do not need tags.

As flips in the triangulation correspond to rotations in the augmented dual graph [26],
 265 the graph grammar is identical to the graph grammar presented before. The only addition
 is the computation of new tags for the vertices on the right-hand side (see Figure 5).
 This grammar has four vertices in left sides, and a maximum of two vertices in any
 right side. Since a triangulation of an n -vertex convex polygon has $n - 2$ triangles and
 n convex hull edges, the augmented dual graph has $2n - 2$ vertices. Thus, Theorem 2

270 gives us the following.

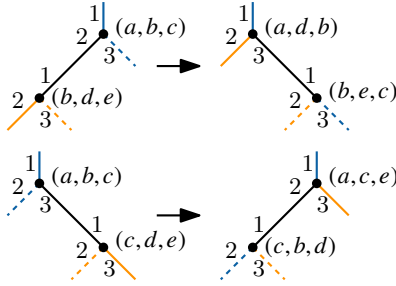


Figure 5: A graph grammar for rotations in augmented dual graphs, which correspond to flips in edge-labelled triangulations of a convex polygon.

Lemma 2. *Given an edge-labelled triangulation G of an n -vertex convex polygon, the number of distinct edge-labelled triangulations reachable from G in m flips is at most $5^{2n-2+2m}$.*

275 This bound can be further refined to 3^{n-1+2m} , using the leader-follower and zero-elimination techniques from Sleator, Tarjan, and Thurton's paper [27]. However, the cruder bound already suffices to derive the correct asymptotic lower bound.

Theorem 3. *There are pairs of edge-labelled triangulations of a convex polygon with n vertices such that transforming one into the other requires $\Omega(n \log n)$ flips.*

280 *Proof.* We first estimate the number of edge-labelled triangulations. An n -vertex convex polygon has $n-3$ diagonals, and in a fan triangulation, each sequence of labellings results in a new triangulation. Thus, there are at least $(n-3)!$ edge-labelled triangulations.

Let d be the diameter of the flip graph. Then, for every graph G , d flips suffice to reach all edge-labelled triangulations. But from Lemma 2, we know that a sequence of m flips can generate at most $5^{2n-2+2m}$ unique edge-labelled triangulations. This gives us the following bound.

$$\begin{aligned}
 5^{2n-2+2d} &\geq (n-3)! \\
 \log_5 5^{2n-2+2d} &\geq \log_5 (n-3)! \\
 2n-2+2d &\geq \log_5 (n!/n^3) \\
 2d &\geq \log_5 n! - \log_5 n^3 - 2n + 2 \\
 2d &\geq \Omega(n \log n) - O(n) \\
 d &\geq \Omega(n \log n) \quad \square
 \end{aligned}$$

In Section 4, we discuss the implications of this theorem on sorting permutations in length-weighted models.

3.2. Simultaneous flips

285 We now turn our attention to simultaneous flips. Recall that a set of edges may be simultaneously flipped if each edge is flippable and no two edges are incident to the same triangle. For edge-labelled triangulations of a convex polygon, we show that $O(\log^2 n)$ simultaneous flips suffice. Our approach essentially hinges on showing how to perform a balanced partition step of quicksort with $O(\log n)$ simultaneous flips.

290 **Theorem 4.** *Given two edge-labelled triangulations of a convex polygon with n vertices, $O(\log^2 n)$ simultaneous flips are sufficient to transform one to the other.*

Proof. We first ignore the labels and, by the result of Galtier et al. [14], use $O(\log n)$ simultaneous flips to transform both edge-labelled triangulations to the canonical triangulations (T^*, l'_1) and (T^*, l'_2) , respectively. Next, we show how to transform a general canonical triangulation (T^*, ρ) to (T^*, ρ^*) , where ρ^* is the sorted permutation. We do this by imitating quicksort. In particular, we show that the partition step of quicksort can be carried out in $O(\log n)$ simultaneous flips. Then, by recursing on both halves simultaneously, we get that the total number of simultaneous flips satisfies the recursion $T(n) = T(n/2) + O(\log n)$, which solves to $T(n) = O(\log^2 n)$.

300 We now address the partition step of quicksort. Let E be the set of all non-boundary edges, and let $E_l = \{\rho_i \mid i < n/2 \text{ and } \rho_i \geq n/2\}$ be the set of edges that are in the left half in ρ but in the right half in ρ^* . Similarly, let E_r be the set of edges that are in the right half in ρ but in the left half in ρ^* . We first flip all the edges of $E \setminus (E_l \cup E_r)$. Since one simultaneous flip can flip every other edge of the set, $O(\log n)$ simultaneous flips suffice to flip the whole set. The edges of $(E_l \cup E_r)$ now form a canonical triangulation of a smaller convex polygon. For the remainder of the proof we work only with this smaller convex polygon. In other words, we only consider canonical triangulations (T^*, ρ) where every edge on the left half wants to go to the right half and vice versa.

310 Let \mathcal{T} be a canonical triangulation where every edge on the left half is coloured red and every edge on the right half is coloured blue. Let \mathcal{T}' be the triangulation with the colours interchanged. We show how to transform \mathcal{T} to \mathcal{T}' with $O(\log n)$ simultaneous flips. We do this by converting both \mathcal{T} and \mathcal{T}' into a common coloured triangulation in which the diagonals form a path that alternates between red and blue edges, as shown in Figure 6a. We call this triangulation an *alternating zig-zag*. Note that p_1 is the high degree vertex of the canonical triangulation.

315 Imagine a line L passing through p_1 that separates the red edges from the blue edges in the canonical triangulation and crosses all the edges in the alternating zig-zag. We show how to transform the alternating zig-zag to \mathcal{T}' . We claim that with $O(1)$ simultaneous flips, we can ensure that half the blue edges are to the left of this line and half the red edges are to the right. Moreover, the edges that still cross L form an alternating zig-zag of a subpolygon at most half the size of \mathcal{T}' .

To do this, note that each four consecutive triangles of the zig-zag form a hexagon, as shown in Figure 6b. Each hexagon has three diagonals, two red and one blue, and the

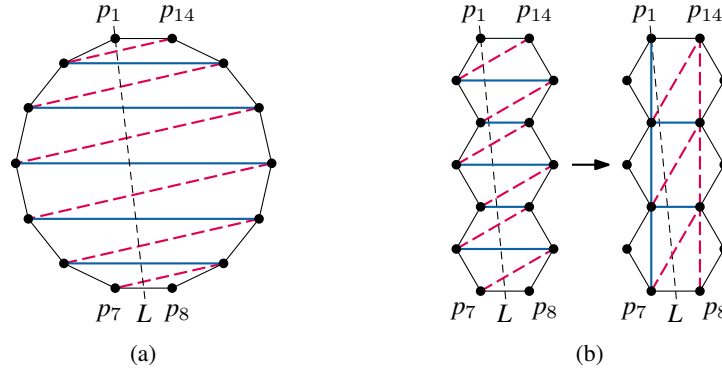


Figure 6: (a) An 11-edge zig-zag alternating between red (dotted) and blue (solid) edges. (b) An alternating zig-zag can be thought of as a collection of hexagons. In $O(1)$ simultaneous flips, we can move half of the edges to the correct side of L , with the rest forming a smaller alternating zig-zag.

hexagons are separated by the remaining blue edges. As any transformation of a hexagon
 325 takes a constant number of flips, we can reconfigure all hexagons with a constant number
 of simultaneous flips. In particular, we transform each hexagon to have one blue edge to
 the left of L and one red edge to the right of L , as shown in Figure 6b. The triangles
 that still intersect L now form an alternating zig-zag of half the size, consisting of the
 remaining red diagonal inside each hexagon and the blue edges separating the hexagons.
 330 Thus, by repeating the process $O(\log n)$ times, we ensure that every blue edge is to the
 left of L and every red edge is to the right.

Now the subpolygon to the left of L has only blue diagonals, all of which we can
 make incident to p_1 with $O(\log n)$ simultaneous flips. By doing the same with the red
 edges, we can transform the alternating zig-zag into \mathcal{T}' using $O(\log n)$ simultaneous
 335 flips. Similarly, we can transform the alternating zig-zag into \mathcal{T} by moving the blue
 diagonal of each hexagon to the right and one of the red diagonals to the left. Thus
 we can perform the partition step in $O(\log n)$ simultaneous flips, which completes the
 proof. \square

The $\Omega(\log n)$ lower bound on the worst-case number of simultaneous flips in the
 340 unlabelled case trivially provides an $\Omega(\log n)$ lower bound in the edge-labelled setting
 as well. We prove a stronger lower bound. In particular, the following theorem shows
 that even the “partition” step of quick sort requires at least $\Omega(\log n)$ simultaneous flips.

Theorem 5. *Let \mathcal{T} be a canonical triangulation on n vertices with all diagonals in the
 left half coloured red, and all edges in the right half coloured blue. Let \mathcal{T}' be the same
 345 canonical triangulation, with the colours interchanged. Transforming \mathcal{T} to \mathcal{T}' requires
 at least $\Omega(\log n)$ simultaneous flips.*

Proof. Consider the line L that passes through p_1 and separates the red edges from the
 blue ones in \mathcal{T} . For each edge, the side of L it inhabits in \mathcal{T} is different from the side it

inhabits in \mathcal{T}' . A single flip cannot replace an edge that lies completely to the left of L with an edge that lies completely to the right. Thus, in any simultaneous flip-sequence that transforms \mathcal{T} to \mathcal{T}' , for any edge e , there must be a triangulation where e intersects L .

Consider any simultaneous flip-sequence \mathcal{F} that transforms \mathcal{T} to \mathcal{T}' , i.e., \mathcal{F} is the sequence $(\mathcal{T} = \mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_{k-1}, \mathcal{T}_k = \mathcal{T}')$. For any $j \in [1..k]$, let c_j be the number of edges e for which there exists a \mathcal{T}_r with $r \leq j$ such that e intersects L in \mathcal{T}_r . From the argument above, it is clear that $c_1 = 0$ and $c_k = n$. We claim that for all $j \in [1..k-1]$, $c_{j+1} \leq 2c_j + 1$. This shows that $k \geq \Omega(\log n)$.

To see why the claim is true, consider the j^{th} simultaneous flip in the sequence. It makes $\Delta_j = c_{j+1} - c_j$ new edges cross L . Each individual flip happens in its own quadrilateral that has exactly two boundary-edges that intersect L . Since two adjacent quadrilaterals can share an edge and since the quadrilateral at the top and bottom share their boundaries with the polygon, the total number of quadrilateral edges that cross L in \mathcal{T}_j is at least $\Delta_j - 1$. But this number is also at most c_j . Therefore $\Delta_j - 1 \leq c_j$, which means $c_{j+1} \leq 2c_j + 1$. \square

4. Connections between flips and sorting models

Our upper bounds in the previous section depend on bounds for sorting permutations. In this section, we study the connection in the other direction, showing that our lower bound on the diameter of the flip graph implies an $\Omega(n \log n)$ lower bound on the cost of sorting in various length-weighted sorting models.

Sorting permutations of $[1..n]$ using a restricted set of operations has been widely studied. A main operation is the reversal of a subsequence. One of the earliest results of this kind was on “pancake” sorting [15], where a prefix of the sequence can be reversed. Bubble sort also fits into this model, as it operates by swapping two adjacent elements at a time, which is a reversal of size two. The number of size-two reversals it makes is equal to the number of inversions of the permutation and is $\Theta(n^2)$ in the worst case. More generally, any permutation can be sorted with $\Theta(n)$ reversals [20] if arbitrarily large reversals are allowed. Most of the results until now have been about the *number* of operations required. Recently, Pinter and Skiena [24] formulated a model that allows reversals of any size, but assigns them a cost proportional to the length of the reversed subsequence. This model has applications in comparative genomics, where it models a sequence of mutations in the evolution of a chromosome. Bender et al. [2] showed that any input sequence can be sorted in this model with worst-case cost $O(n \log^2 n)$, and gave an $\Omega(n \log n)$ lower bound.

We present a more general length-weighted sorting model, based on swaps. Given a permutation ρ , pick an interval $[i..j]$ with $1 \leq i < j \leq n$ and $j - i = k$ and perform a set of swaps in it. The cost of this operation is $O(k)$. Adding constraints on the kinds of swaps allowed inside the interval leads to more specific models (see Figure 7):

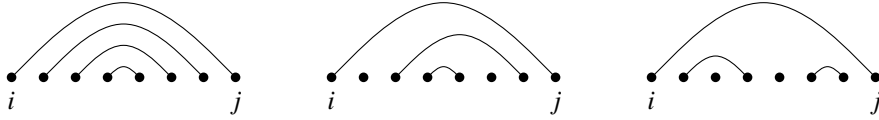


Figure 7: An illustration of the swaps permitted in the contiguous reversals model (left), the non-contiguous reversals model (middle), and the non-crossing swaps model (right).

Contiguous reversals Reverse the entire interval $[i..j]$.

Non-contiguous reversals Pick a subsequence i_1, \dots, i_m of $[i..j]$ and reverse its elements, without changing the position of any element outside the subsequence.

Non-crossing swaps Swap a set $\{(i_1, j_1), \dots, (i_m, j_m)\}$ of pairs inside $[i..j]$ such that there do not exist $x, y \in [1..m]$ with $x < y$ for which $i_x \leq i_y \leq j_x \leq j_y$, i.e. either the two intervals $[i_x, j_x]$ and $[i_y, j_y]$ nest or they are disjoint.

To the best of our knowledge, the latter two models are novel. Of the three models presented, the non-crossing swaps model is the most general, as it can simulate both other models. In fact, the non-crossing swaps model is more general than (length-weighted versions of) a large number of sorting models [13, Chapter 3]. We use our results on flips to give an $\Omega(n \log n)$ lower bound on the worst-case cost of sorting in the non-crossing swaps model and therefore in all of the sorting models that can be simulated by it. Note that our proof of Theorem 1 can be interpreted as an $O(n \log n)$ upper bound on the cost of sorting in the non-contiguous reversals model, and thereby also in the non-crossing swaps model. We first show how to simulate a non-crossing swap with flips. The proof is very similar to the proof of Lemma 1.

Lemma 3. *Let ρ and ρ' be two permutations of $[1..n]$ such that ρ' can be obtained from ρ using one set of non-crossing swaps in an interval of length k . Then the canonical triangulation (T^*, ρ) can be flipped to (T^*, ρ') with $O(k)$ flips.*

Proof. Let $S = [i..j]$ be the sequence of edges in which we perform the swaps. We proceed by induction on $|S|$. In the base case, $|S| = 0$ and we are done, so let $|S| = k$ and assume that the lemma holds for all $|S| < k$. If there is an edge e in S that is not involved in any swap, we flip it. This results in a smaller convex polygon and reduces $|S|$ by one, so we can perform the swaps there by induction and flip e back afterwards.

If every edge is involved in a swap, consider a swap $x = (i_x, j_x)$ such that no other swap is nested inside x . The edges i_x and j_x must be consecutive, since all edges are involved in some swap, and swaps do not cross. Thus, they form a convex pentagon and can be swapped with five flips, as shown in Figure 2. But instead of completing this swap sequence, we halt it after the first three flips. This leaves the edges i_x and j_x out of the way of the remaining edges, so we can perform the rest of the swaps by induction. Afterwards, we perform the final two flips to return i_x and j_x . Thus, we can perform each swap with 5 flips, giving a total cost of $O(k)$. \square

420 Combining this lemma with Theorem 3 gives the desired lower bound.

Theorem 6. *The worst-case cost of sorting a permutation of $[1..n]$ in the non-crossing swaps model is $\Omega(n \log n)$.*

5. Spiral polygons

In this section, we prove the Orbit Conjecture for spiral polygons, even when there
 425 are multiple orbits. We provide a complete characterization of the orbits, and prove a
 tight quadratic bound on the diameter of each connected component of the flip graph.
 Finally, we show how to test for the condition of the Orbit Conjecture in linear time.

A vertex of a polygon is a *convex* vertex if the interior angle at the vertex is less
 than π , otherwise it is *reflex*. A *spiral polygon* is a polygon in which all reflex vertices
 430 are consecutive along the boundary. For a spiral polygon P with m convex and k
 reflex vertices, let $C = c_1 c_2 \dots c_m$ denote the convex chain in clockwise order and
 $R = r_1 r_2 \dots r_k$ denote the reflex chain in counterclockwise order such that c_1 and r_1 are
 adjacent, and c_m and r_k are adjacent. For any vertex c on the convex chain (resp. reflex
 chain), let $V(c)$ be the set of vertices on the reflex chain (resp. convex chain) that are
 435 visible from c .

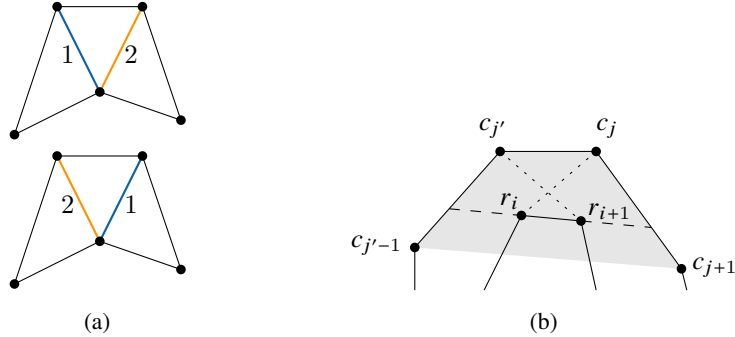


Figure 8: (a) Two edge-labelled triangulations of the same spiral polygon that cannot be transformed into each other by flips. (b) If $V(r_i) \cap V(r_{i+1}) = \{c_j, c_{j'}\}$ then the convex quadrilateral $c_{j-1}, c_j, c_{j'}, c_{j+1}$ is not empty.

Consider a triangulated 5-vertex polygon with a single reflex vertex (see Figure 8a).
 The two diagonals cannot be swapped with flips, thereby implying that the flip graph
 is disconnected. The same thing happens with any spiral polygon with only 3 convex
 vertices—there is a unique triangulation and the flip graph is either a singleton (if
 440 there is at most one reflex vertex), or disconnected. We now characterize when the flip
 graph is connected. We only consider spiral polygons with at least 4 convex vertices.
 A spiral polygon is *locally convex* if it has at least four convex vertices and for every
 four consecutive vertices a, b, c , and d on the convex chain, the quadrilateral $abcd$ is
 empty. Note that it is possible for four consecutive convex vertices to form a non-convex

445 quadrilateral, but then the quadrilateral cannot be empty. Local convexity is the key to the characterization of the orbits and connectivity of the flip graph.

Theorem 7. *For an n -vertex spiral polygon \mathcal{P} with at least 4 convex vertices the flip graph, $FG_{\mathcal{P}}$, is connected if and only if \mathcal{P} is locally convex. Furthermore, the diameter of each connected component of $FG_{\mathcal{P}}$ is $O(n^2)$.*

450 We start by showing that if \mathcal{P} is locally convex, then its flip graph is connected. Observe the following:

Lemma 4. *If \mathcal{P} is locally convex, then for any two consecutive vertices $r_i, r_{i+1} \in R$ on the reflex chain, $|V(r_i) \cap V(r_{i+1})| \geq 3$.*

Proof. Extend the segment r_i, r_{i+1} in both directions until it hits the convex chain and exits the polygon. Let $c_j, \dots, c_{j'}$ be the vertices of the convex chain between these two exit points, noting that $j > 1$ and $j' < m$. Observe that $V(r_i) \cap V(r_{i+1}) = \{c_j, \dots, c_{j'}\}$.

Suppose that $|V(r_i) \cap V(r_{i+1})| \leq 2$. We will show that \mathcal{P} is not locally convex. Consider the vertices $c_{j-1}, c_j, \dots, c_{j'}, c_{j'+1}$. They form a quadrilateral or triangle that contains r_i and r_{i+1} . If they form a quadrilateral, this violates local convexity, and if they form a triangle then we can extend it to a quadrilateral that violates local convexity. 460 □

To prove connectivity of the flip graph, we define the following canonical triangulation \mathcal{T}' , together with a partition of its diagonals into *fans*. The canonical triangulation will be defined by specifying the triangle incident to each edge of the polygon. Let $r_i r_{i+1}$ be an edge of the reflex chain with $V(r_i) \cap V(r_{i+1}) = \{c_j, \dots, c_{j'}\}$. The apex of the triangle on edge $r_i r_{i+1}$ is defined to be $c_{\lfloor (j+j')/2 \rfloor}$. The diagonals of all such triangles incident to a given convex vertex c constitute the *convex fan* at c . See Figure 9a. After this, each edge $c_j c_{j+1}$ has a unique reflex vertex that can be its apex. We add these triangles to \mathcal{T}' . The new diagonals incident to reflex vertex r constitute the *reflex fan* at r . See Figure 9b. 470 Note that every diagonal of \mathcal{T}' connects a convex and a reflex vertex, and thus there is a natural ordering of the diagonals of \mathcal{T}' consistent with the ordering of the convex and reflex chains. We label the diagonals of \mathcal{T}' with labels $\{1, \dots, |E|\}$ in this order ($|E|$ is the number of diagonals in \mathcal{T}'). The fans partition the diagonals of \mathcal{T}' and there is a natural ordering of the fans D_1, \dots, D_t such that the diagonals of D_i occur before the diagonals of D_j if $i < j$. We now provide a flip sequence to move from a given edge-labelled triangulation \mathcal{T} to the canonical triangulation \mathcal{T}' using $O(n^2)$ flips. As 475 flips are reversible, this implies that we can transform any edge-labelled triangulation into any other.

We start by ignoring labels and transforming \mathcal{T} into the unlabelled canonical triangulation with $O(n)$ flips [16]. What remains is to sort the edge labels. Our flip sequence achieves this with $O(n^2)$ flips via insertion sort. If $d_1, \dots, d_{|E|}$ is the ordered list of diagonals, the *insertion* of d_j at d_i for $i < j$ replaces the subsequence d_i, \dots, d_j by $d_j, d_i, d_{i+1}, \dots, d_{j-1}$. It suffices to show how to perform a single insertion using

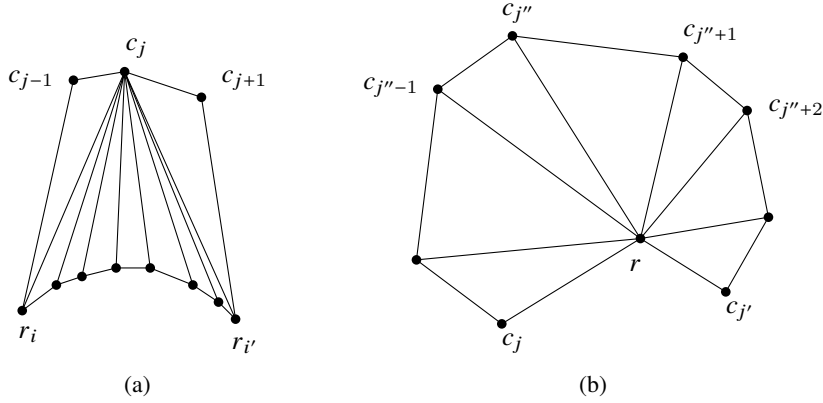


Figure 9: (a) A convex fan (c_{j-1} and c_{j+1} are not part of the fan). (b) A reflex fan.

$O(n)$ flips. Based on the structure of the canonical triangulation, it is sufficient to show
 485 how to perform an insertion within one fan and between two consecutive fans with a
 number of flips proportional to the sizes of the fans.

Lemma 5. *For each fan D_i , insertion between two edges of D_i can be performed using $O(|D_i|)$ flips. In addition, swapping the last edge of D_i with the first edge of D_{i+1} can be performed using $O(|D_i| + |D_{i+1}|)$ flips.*

490 *Proof.* We first show how to perform an insertion in a convex fan D_i (see Figure 9a). Let
 c_j be the common convex vertex and $r_i, \dots, r_{i'}$ be the sequence of reflex vertices. This
 means that c_j was chosen to be the apex vertex for each edge of the reflex chain between
 r_i and $r_{i'}$ during the construction of the canonical triangulation. Using Lemma 4,
 together with the fact that the apex is always chosen to be in the “middle” of the set of
 495 visible vertices, we get that c_{j-1} and c_{j+1} must exist and be visible to all reflex vertices
 of D_i . Also note that the canonical triangulation contains diagonals $r_i c_{j-1}$ and $r_{i'} c_{j+1}$.

We show how to insert diagonal $c_j r_k$ to $c_j r_\ell$ for any $i \leq \ell < k \leq i'$. First,
 flip all diagonals intersecting segment $r_k c_{j+1}$ and all diagonals intersecting $r_{k-1} c_{j-1}$.
 This results in the convex pentagon $r_{k-1} r_k c_{j+1} c_j c_{j-1}$. Swap the two diagonals of this
 500 pentagon with five flips as shown in Figure 2 in Section 3. Notice that this has moved
 $c_j r_k$ to $c_j r_{k-1}$. To move it to $c_j r_{k-2}$, we flip a constant number of edges to get the
 convex pentagon $r_{k-2} r_{k-1} c_{j+1} c_j c_{j-1}$. We continue in this way until $c_j r_k$ has reached
 $c_j r_\ell$. At that point, we flip all diagonals adjacent to c_{j-1} and c_{j+1} back to their original
 positions. Notice that the order of these diagonals is preserved and this results in the
 505 insertion of $c_j r_k$ to $c_j r_\ell$, using a total of $O(|D_i|)$ flips.

If D_i is a reflex fan (Figure 9b), let r be its common reflex vertex and let $c_j, \dots, c_{j'}$
 be its convex vertices. Consider two consecutive diagonals $r c_{j''}$ and $r c_{j''+1}$ of D_i . These
 form a triangle t with apex r . Consider the triangle before t . If it had apex $c_{j''}$ then the
 diagonal $r c_{j''}$ would have been placed in the convex fan D_{i-1} . Therefore it has apex r

510 and includes the diagonal (or boundary edge) $rc_{j''-1}$. Similarly, the triangle after t must have apex r as well and includes the diagonal (or boundary edge) $rc_{j''+2}$. Now consider the pentagon formed by r and the chain $c_{j''-1} \dots c_{j''+2}$. Since \mathcal{P} is locally convex, this pentagon must be empty and convex. Thus we can swap $rc_{j''}$ and $rc_{j''+1}$ with five flips as before. This shows that any two consecutive diagonals of a reflex fan can be swapped using five flips. Using a sequence of such swaps, we can perform any insertion with $O(|D_i|)$ flips.

Finally, we show how to swap the last edge d of fan D_i with the first edge d' of fan D_{i+1} . We consider four cases. The case when D_i and D_{i+1} are both convex fans is shown in Figure 10a and the case when D_i is convex and D_{i+1} is reflex is shown in Figure 10b. In both cases, let d and d' share vertex r . Our strategy is to perform a linear number of flips in the convex fans to get an empty pentagon where d and d' can be swapped. The pentagon consists of r together with 4 consecutive vertices $efgh$ on the convex chain that are visible from r . Local convexity implies that the pentagon is empty and convex.

525 The case when D_{i+1} is convex and D_i is reflex is symmetric. Finally, it is not possible to have both D_i and D_{i+1} reflex because a reflex fan was defined to contain all the diagonals between two convex fans. \square

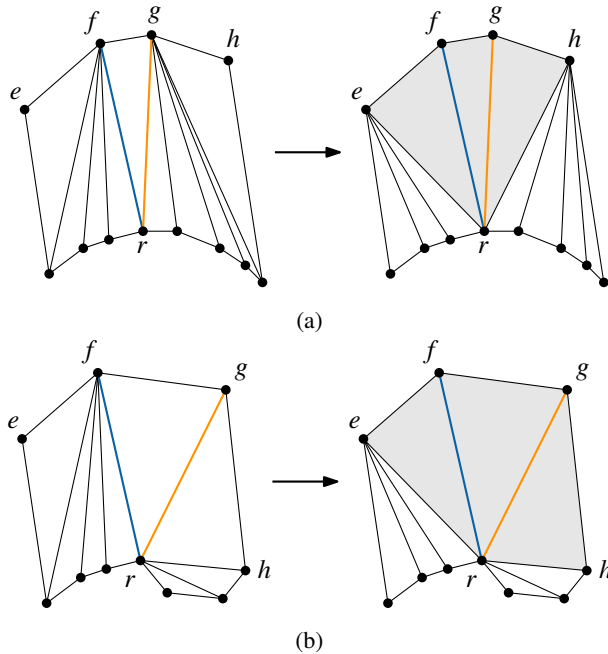


Figure 10: Swapping edges between adjacent fans D_i and D_{i+1} using the pentagon $efghr$. (a) D_i and D_{i+1} are both convex fans. (b) D_i is a convex fan but D_{i+1} is a reflex fan.

Using Lemma 5, we can reorder the labels of the edges in a locally convex spiral polygon with $O(n^2)$ flips by simulating insertion sort. What remains to be shown is that
 530 if the flip graph is connected then \mathcal{P} is locally convex. We show this in the contrapositive.

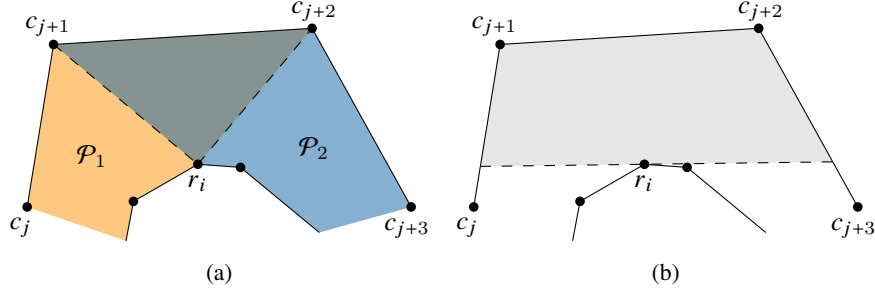


Figure 11: (a) Decomposition of a spiral polygon that is not locally convex into two sub-polygons. (b) To find r_i , we move a segment from $c_{j+1}c_{j+2}$ to c_jc_{j+3} , stopping when we hit the reflex chain.

Lemma 6. *If there exist four consecutive vertices $c_j, c_{j+1}, c_{j+2}, c_{j+3}$ on C that form a non-empty quadrilateral, then there exist two spiral polygons \mathcal{P}_1 and \mathcal{P}_2 , such that: (a) each boundary edge of \mathcal{P}_1 and \mathcal{P}_2 is either a boundary edge or a diagonal of \mathcal{P} , (b) each diagonal of \mathcal{P} is a diagonal of either \mathcal{P}_1 or \mathcal{P}_2 , and (c) no diagonal of \mathcal{P}_1 can flip
 535 to a diagonal of \mathcal{P}_2 .*

Proof. We claim that there is a vertex r_i on R that sees both c_{j+1} and c_{j+2} , and such that c_j does not see r_{i+1} and c_{j+3} does not see r_{i-1} . (Justification below.) Given such a vertex r_i , let \mathcal{P}_1 and \mathcal{P}_2 be the spiral polygons defined by the chains c_1, \dots, c_{j+2} and r_1, \dots, r_i , and c_{j+1}, \dots, c_m and r_i, \dots, r_k , respectively. Then \mathcal{P}_1 and \mathcal{P}_2 overlap on
 540 triangle $r_i c_{j+1} c_{j+2}$. See Figure 11a.

It is easy to see that \mathcal{P}_1 and \mathcal{P}_2 satisfy condition (a), as $c_{j+1}r_i$ and $c_{j+2}r_i$ are diagonals of \mathcal{P} . For condition (b) and (c), we note the following. By construction, there is no diagonal between a vertex of $\mathcal{P}_1 - \mathcal{P}_2$ and a vertex of $\mathcal{P}_2 - \mathcal{P}_1$. This implies that the diagonals of \mathcal{P} are partitioned between \mathcal{P}_1 and \mathcal{P}_2 . For condition (c), suppose a
 545 diagonal of \mathcal{P}_1 flipped to one in \mathcal{P}_2 . Then they must intersect, so they must be incident to c_{j+2} and c_{j+1} respectively. But then they cannot form an empty convex quadrilateral.

To find r_i , imagine the continuous motion of a segment whose initial position is $c_{j+1}c_{j+2}$ and whose final position is c_jc_{j+3} (see Figure 11b). One endpoint of the segment moves at uniform speed in a straight line from c_{j+1} to c_j , and the other endpoint moves at uniform speed in a straight line from c_{j+2} to c_{j+3} . The initial segment does
 550 not intersect the reflex chain; the final segment does. Therefore, by continuity, some intermediate segment, say pq , lies inside the polygon but contains a reflex vertex (or possibly two). We choose one such vertex to be r_i . The area swept by the segment until reaching pq forms an empty convex quadrilateral with c_{j+1}, c_{j+2} , and r_i on the boundary. Thus, r_i sees both c_{j+1} and c_{j+2} . Furthermore, segment pq is a chord of
 555

\mathcal{P} that goes through a reflex vertex, namely r_i . It is well known² that such a chord partitions \mathcal{P} into three subpolygons \mathcal{P}_A , \mathcal{P}_B , and \mathcal{P}_C such that no two vertices in the non-adjacent subpolygons \mathcal{P}_A and \mathcal{P}_C can see each other. As $c_j \in \mathcal{P}_A$ and $r_{i+1} \in \mathcal{P}_C$, this implies that c_j cannot see r_{i+1} . Analogously, c_{j+3} cannot see r_{i-1} , showing that r_i has all the properties we were looking for. □

This concludes the proof of Theorem 7. It also gives us a way to decide, in polynomial time, whether a given spiral polygon has a connected flip graph: for each quadrilateral formed by four consecutive vertices of the convex chain, check if it is empty. A naive implementation of this procedure will take $O(n^2)$ time. However, we can reduce the running time to $O(n)$ by walking simultaneously along the convex and reflex chains of the polygon and using the fact that visibility is monotonic. Thus we get the following theorem.

Theorem 8. *Given a spiral polygon \mathcal{P} with n vertices, we can decide in linear time whether the flip graph of the edge-labelled triangulations of \mathcal{P} is connected.*

Theorem 7 gives a complete characterization of the orbits in a spiral polygon: The orbit of a diagonal d in \mathcal{P} is precisely the set of diagonals in the maximal locally convex subpolygon of \mathcal{P} that contains d . With this characterization, we now prove the Orbit Conjecture for spiral polygons.

Theorem 9. *Given a spiral polygon \mathcal{P} with n vertices and two of its edge-labelled triangulations \mathcal{T}_1 and \mathcal{T}_2 , there exists a flip sequence that transforms \mathcal{T}_1 to \mathcal{T}_2 if and only if for each label λ , the diagonal of \mathcal{T}_1 with label λ is in the same orbit as the diagonal of \mathcal{T}_2 with label λ . Moreover, we can decide in $O(n)$ time if such a flip sequence exists, and if it does, it is of length at most $O(n^2)$.*

Proof. The ‘only if’ direction is easy. For the ‘if’ direction, we provide a flip sequence. First ignore labels and transform both \mathcal{T}_1 and \mathcal{T}_2 into the unlabelled canonical triangulation \mathcal{T} . Then rearrange the labels in \mathcal{T} . If the diagonal with label λ in \mathcal{T}_1 was in the same orbit as the diagonal with label λ in \mathcal{T}_2 , they must be diagonals of the same locally convex subpolygon \mathcal{P}_i . By applying Theorem 7, we can rearrange the labels inside each \mathcal{P}_i with at most $O(|\mathcal{P}_i|^2)$ flips. We can find the decomposition into maximal locally convex subpolygons in linear time by walking simultaneously along the convex and reflex chains of the polygon, and exploiting the fact that visibility is monotonic. Once we have the decomposition, we can verify in constant time whether two edges with the same label are in the same orbit. □

Finally, note that our quadratic bound on the flip distance is tight. Consider, for example, a locally convex spiral polygon for which every reflex vertex r has $|V(r)| \leq 4$ (see Figure 12). Let $d_1, \dots, d_{|E|}$ be the diagonals of the canonical triangulation of this

²See, for instance, Bose et al. [6], Lemma 1.

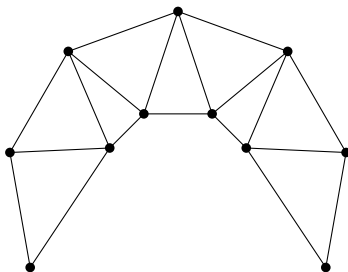


Figure 12: A spiral polygon whose flip graph has diameter $\Theta(n^2)$.

595 polygon listed left-to-right. It is clear that the distance in the quadrilateral graph between any d_i and d_j is $\Omega(|j - i|)$. Thus, we can find $\Omega(n)$ pairs of diagonals at distance $\Omega(n)$ in the quadrilateral graph. Since any flip sequence that moves a label from d_i to d_j must contain at least $\Omega(|j - i|)$ flips of that label, this gives a lower bound of $\Omega(n^2)$.

Theorem 10. *There exists a spiral polygon with n vertices that has a connected flip graph whose diameter is $\Omega(n^2)$.*

6. Lower bound on flips in edge-labelled polygons

600 In this section, we provide an example of a polygon with two reflex chains where the diameter of the flip graph is $\Theta(n^3)$. The example also demonstrates a difficulty in generalizing our results on spiral polygons to more general polygons. Our example is based on a *channel* [19] that consists of two reflex chains visible to each other and joined by two edges. We define an *augmented channel* to consist of two chains $A = a_1 \dots a_m$ and $B = b_1 \dots b_{m+2}$, as shown in Figure 13, such that: (a) every vertex on A is visible from every vertex on B , (b) the vertices on A form a reflex chain, and (c) all vertices on B , except for b_1, b_2, b_{m+1} , and b_{m+2} are reflex.
605

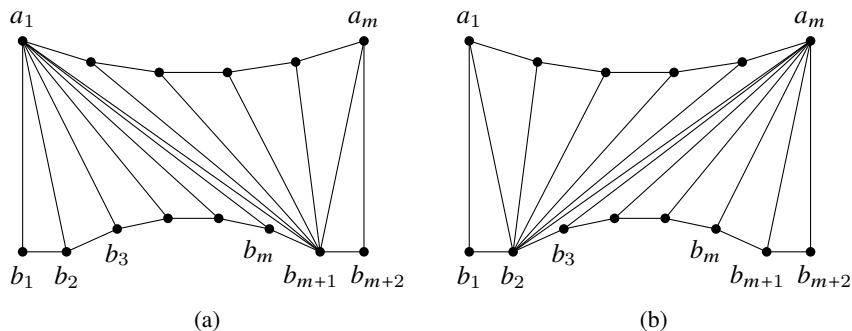


Figure 13: (a) The left-inclined triangulation of an augmented channel. (b) The right-inclined triangulation of an augmented channel.

Theorem 11. *The diameter of the flip graph of an augmented channel is $\Theta(n^3)$.*

Proof. We first show the upper bound. Any triangulation of the augmented channel induces an ordering on its diagonals as follows. If all diagonals connect a vertex on A with a vertex on B , then the ordering induced is the same as the order in which a ray passing through the polygon from left to right will hit each diagonal. If, however, one or both of the diagonals b_1b_3 and b_mb_{m+2} are present, then we take the ordering we would get by flipping those diagonals. The ordering on the diagonals defines an ordering on the labels.

Hurtado et al. [19, Theorem 3.8] showed that in the unlabelled setting, any triangulation of a channel can be transformed into any other triangulation of the channel using $O(n^2)$ flips. We observe that their transformation preserves the ordering of the labels. Thus for the $O(n^3)$ bound on the diameter of the flip graph, the challenge is in rearranging the labels. For any triangulation, let D_1 denote the set consisting of the first m labels and D_2 denote the set consisting of the last m labels. Since there are $2m - 1$ labels in total, D_1 and D_2 will have one label in common, namely, the m^{th} label. We will use two particular triangulations of the augmented channel: the *left-inclined* and *right-inclined* triangulations shown in Figure 13. Note that in the left-inclined triangulation, D_2 is exactly the set of diagonals incident on b_{m+1} and thus forms a convex fan. Similarly, in the right-inclined triangulation, D_1 forms a convex fan. This means both D_1 and D_2 can be sorted in $O(n^2)$ flips and any insertion in D_1 or D_2 can be performed in $O(n)$ flips using Lemma 5.

Thus our strategy is as follows. Given \mathcal{T}_1 and \mathcal{T}_2 , first ignore labels and transform both into the right-inclined triangulation. Then, if D_1 contains a label bigger than m , insert it (using the spiral polygon formed by b_1, b_2, b_3 , and chain A) at b_2a_m and transform the triangulation into a left-inclined triangulation while preserving the order of labels. Next, if D_2 contains a label smaller than m , insert the label at a_1b_{m+1} and transform back into the right-inclined triangulation and repeat until D_1 only contains labels $1, \dots, m$ and D_2 only contains labels $m, \dots, 2m - 1$. Finally, sort D_1 and then sort D_2 . Since transforming between the left-inclined and right-inclined triangulations takes $O(n^2)$ flips and insertion inside D_1 or D_2 takes $O(n)$ flips, we get a bound of $O(n^3)$ flips.

To show the lower bound, we use the observation that to move a label from $D_2 \setminus D_1$ to $D_1 \setminus D_2$, we must first move that label from $D_2 \setminus D_1$ to the m^{th} diagonal and then to $D_1 \setminus D_2$. The first step can be performed only if the diagonal a_1b_{m+1} is present and the second step can be carried out only if the diagonal b_2a_m is present. Going from a triangulation that has the diagonal a_1b_{m+1} to a triangulation that has the diagonal b_2a_m requires $\Omega(n^2)$ flips using the argument from Hurtado et al. [19]. Thus if we want to transform a triangulation with $D_1 = \{m, \dots, 2m - 1\}$ and $D_2 = \{1, \dots, m\}$ to a triangulation with $D_1 = \{1, \dots, m\}$ and $D_2 = \{m, \dots, 2m - 1\}$, we will need at least $\Omega(n^3)$ flips. \square

The theorem above demonstrates that deciding whether, say, a_1b_m and a_1b_{m-1} are

connected depends on the exact position of b_2 and b_{m+1} . In the example above they are connected, but if b_2 and b_{m+1} were reflex vertices, they would not be.

650 **7. Combinatorial triangulations**

In this section, we prove the Orbit Conjecture for edge-labelled combinatorial triangulations and show that the diameter of the flip graph is $\Theta(n \log n)$. Note that we consider two edge-labelled triangulations to be equivalent if they have an isomorphism that preserves the edge labels (in other words, the vertices are unlabelled).

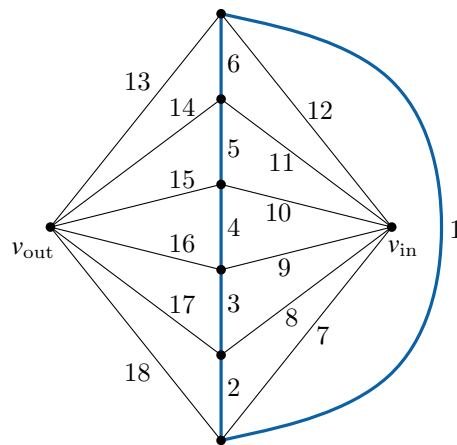


Figure 14: The labelled canonical triangulation on 8 vertices. The spine is indicated in bold.

655 **Theorem 12.** *Any edge-labelled combinatorial triangulation with n vertices can be transformed into any other by $O(n \log n)$ flips.*

Proof. We show that we can transform any edge-labelled combinatorial triangulation into a canonical one using $O(n \log n)$ flips. As flips are reversible, we can also go from the canonical triangulation to any other, which proves the theorem.

660 We use a canonical triangulation much like the one used by Sleator et al. [27] for the vertex-labelled variant. It is a double wheel: a cycle of length $n - 2$ (called the *spine*), plus a vertex v_{in} inside the cycle and a vertex v_{out} outside the cycle, each connected to every vertex on the cycle (see Figure 14). For our canonical labelling, we separate the labels into three groups. Group \mathcal{S} contains labels 1 through $n - 2$, which we place on the spine edges, starting with the edge on the outer face and continuing in clockwise order around v_{in} . The next $n - 2$ labels make up group \mathcal{C}_{in} and are placed on the edges incident to v_{in} in clockwise order, starting with the edge incident to the vertex shared by the edges with labels 1 and 2. Finally, group \mathcal{C}_{out} consists of the last $n - 2$ labels, which we place on the edges incident to v_{out} in clockwise order, starting with the edge
665
670 that shares a vertex with the edge labelled $2n - 4$.

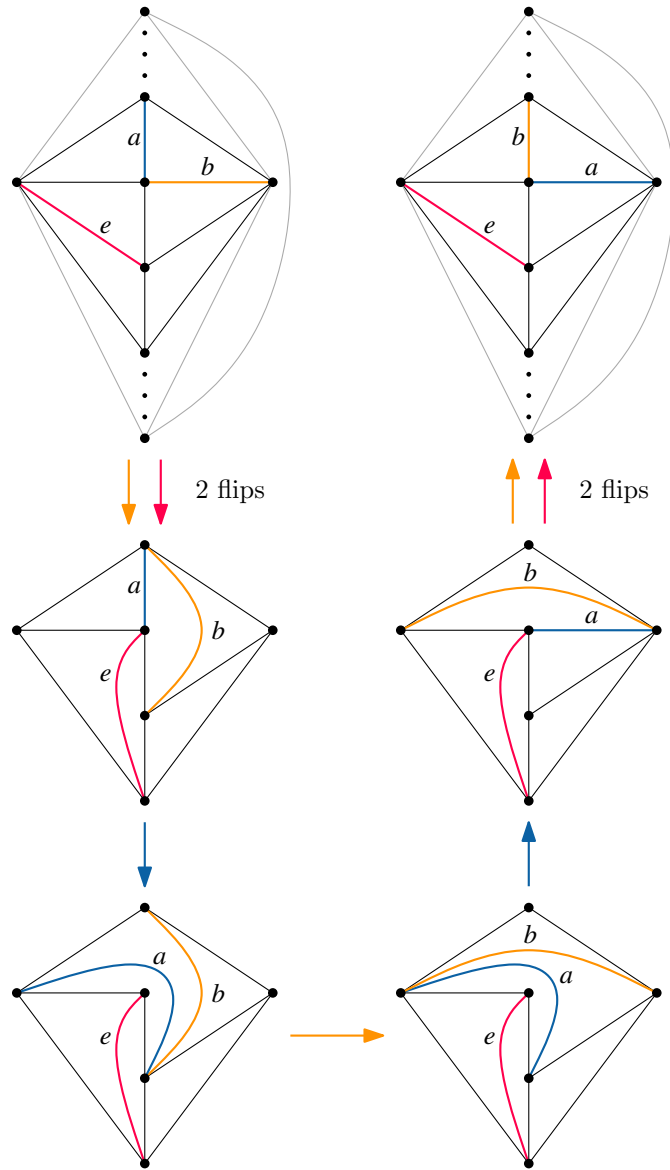


Figure 15: A sequence of seven flips that swaps two edges a and b that are consecutive around a vertex on the spine. Although edge e ends up at the same place as at the start of the sequence, it essentially acts as a catalyst here. If we did not flip it, we would not be able to flip edge a after edge b , as that would create a duplicate edge.

Our algorithm first ignores the labels and transforms the given triangulation into the unlabelled canonical triangulation. This requires $O(n)$ flips [27] and results in the correct graph, although the labels may be in arbitrary positions. To fix the labels, we first get the groups right, so all labels in group \mathcal{S} are on the spine, etc., and then rearrange labels within each group.

We use two main tools for this. The first is a *swap* that interchanges one spine edge with an incident non-spine edge in seven flips, using the flip sequence depicted in Figure 15. Our second tool is a *scramble* algorithm that reorders all labels incident to v_{in} or v_{out} using $O(n \log n)$ flips. To do this we first flip the spine edge that is part of the exterior face (labelled 1 in Figure 14) and then apply the algorithm from Theorem 1 to the outerplanar graph induced by the spine plus v_{in} (or v_{out}), observing that no flip will create a duplicate edge since the omitted edges are all incident to v_{out} (resp. v_{in}). Note that this method cannot alter the labels on the two non-spine edges that lie on the exterior face of the outerplanar graph (labelled 7 and 12 in Figure 14), but since there are only two of these, we can move them to their correct places by swapping them along the spine, using $O(n)$ flips total.

To get the labels of group \mathcal{S} on the spine, we partner every edge incident to v_{in} that has a label in \mathcal{S} with an edge on the spine that has a label in C_{in} or C_{out} . A scramble at v_{in} makes each such edge incident to its partner, and then swaps partners. By doing the same at v_{out} , all labels of \mathcal{S} are placed on the spine. Next we partner every edge incident to v_{in} that has a label in C_{out} with an edge incident to v_{out} that has a label in C_{in} . A scramble at v_{in} makes partners incident, and three swaps per pair then exchange partners.

This ensures that each edge's label is in the correct group, but the order of the labels within each group may still be wrong. Rearranging the labels in C_{in} and C_{out} is straightforward, as we can simply scramble at v_{in} and v_{out} , leaving only the labels on the spine out of order. We then use swaps to exchange the labels on the spine with those incident to v_{in} in $O(n)$ flips and scramble at v_{in} to order them correctly. Since this scramble does not affect the order of labels on the spine, we can simply exchange the edges once more to obtain the canonical triangulation. \square

7.1. Lower bound

The proof for the lower bound for combinatorial triangulations is very similar to the lower bound for triangulations of a convex polygon, described in Section 3.1. We again construct a graph grammar, which describes transformations on the dual graph that correspond to flips.

As our primary graph is a combinatorial triangulation, each vertex of the dual graph corresponds to a triangle and has degree three. As such, there is no distinction between internal nodes and leaves, and no root. This means that we need to adapt our definitions slightly. Without a root, the placement of the edge-end labels is less constrained. We only require that they occur in counter-clockwise order around each vertex. The order of labels in each tag can now follow the placement of the edge-end labels: the first label

belongs to the primary edge corresponding to the dual edge with edge-end label 1, and so on.

Finally, we need a few more production rules to deal with all possible rotations of the edge-end labels around the two triangle-vertices involved in the flip. The full collection of rules is shown in Figure 16.

715

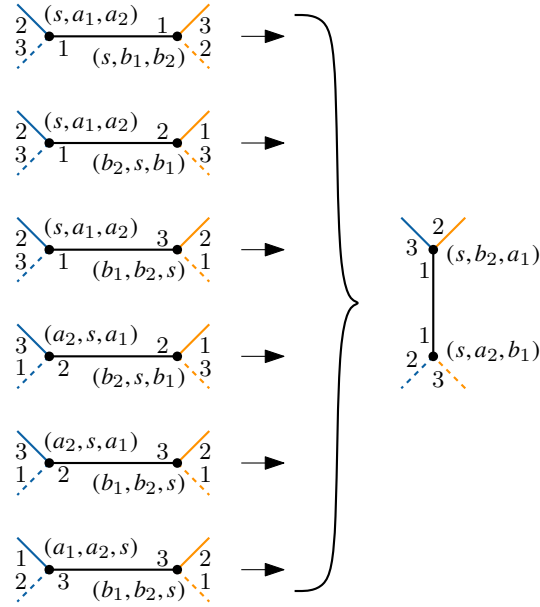


Figure 16: A graph grammar that corresponds to flips in edge-labelled combinatorial triangulations. The right-hand side of all productions is the same.

As the dual graph of an n -vertex combinatorial triangulation has $2n - 4$ vertices, Theorem 2 gives us the following bound.

Lemma 7. *Given an n -vertex edge-labelled combinatorial triangulation G , the number of distinct edge-labelled triangulations reachable from G in m flips is at most $13^{2n-4+2m}$.*

720

Again, Sleator, Tarjan, and Thurston [27] show that this bound can be significantly reduced (to $3^{2n-4}8^m$), but the simple bound suffices for our purposes.

Theorem 13. *There are pairs of edge-labelled combinatorial triangulations with $n \geq 5$ vertices such that transforming one into the other requires $\Omega(n \log n)$ flips.*

Proof. If we fix the labelling of the spine edges in the canonical triangulation from the proof of Theorem 12, any relabelling of the remaining edges is unique. Thus, there are at least $(2n - 6)!$ distinct edge-labelled combinatorial triangulations. Combined with Lemma 7, this implies that $13^{2n-4+2d} \geq (2n - 6)!$, where d is the diameter of the flip

graph. We derive the following.

$$\begin{aligned}
13^{2n-4+2d} &\geq (2n-6)! \\
13^{2n-4+2d} &\geq n! && (\text{for } n \geq 5) \\
\log_{13} 13^{2n-4+2d} &\geq \log_{13} n! \\
2n-4+2d &\geq \log_{13} n! \\
2d &\geq \log_{13} n! - 2n + 4 \\
2d &\geq \Omega(n \log n) - O(n) \\
d &\geq \Omega(n \log n) && \square
\end{aligned}$$

7.2. Simultaneous flips

725 Recall that, in a triangulation of a convex polygon, a simultaneous flip is a set of flips that are executed in parallel, such that no two flipped edges share a triangle. In a combinatorial triangulation, we have the additional requirement that the resulting graph may not contain duplicate edges.

730 Simultaneous flips in combinatorial triangulations were first studied by Bose et al. [4]. They showed a tight $\Theta(\log n)$ bound on the diameter of the flip graph. As part of their proof, they showed that every combinatorial triangulation can be made 4-connected with a single simultaneous flip. Recently, Cardinal et al. [10] proved that it is possible to find such a simultaneous flip that consists of fewer than $2n/3$ individual flips. They used this result to obtain arc drawings of planar graphs in which only $2n/3$ edges are represented by multiple arcs.

735 In this section, we show that, just as in the non-simultaneous setting, we obtain the same bounds for edge-labelled convex polygons and edge-labelled combinatorial triangulations. That is, we can transform any edge-labelled combinatorial triangulation into any other with $O(\log^2 n)$ simultaneous flips, and $\Omega(\log n)$ simultaneous flips are sometimes necessary. The lower bound holds already in the unlabelled setting, if one vertex has linear degree in the first triangulation, while every vertex has constant degree in the second. We now prove the upper bound.

Theorem 14. *Any edge-labelled combinatorial triangulation with n vertices can be transformed into any other by $O(\log^2 n)$ simultaneous flips.*

745 *Proof.* We closely follow the strategy of the proof of Theorem 12. We first transform the given triangulations into the canonical one with $O(\log n)$ simultaneous flips, using the result of Bose et al. [4]. This reduces the problem to sorting the edge labels on the canonical triangulation. In the non-simultaneous setting, we did this by reordering the labels on the edges incident to v_{in} or v_{out} (called *scrambling*), and swapping a subset of spine edges with incident non-spine edges. Thus, the theorem follows if we can show how to perform these operations with $O(\log^2 n)$ simultaneous flips.

750 Since the sequence of flips from Figure 15, which swaps a single spine edge with an incident non-spine edge, only involves a constant number of triangles, it is tempting

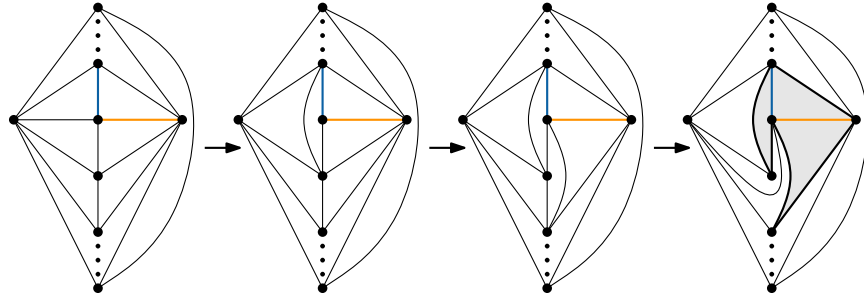


Figure 17: A sequence of three flips that creates a pentagon (shaded) in which the two highlighted edges can be swapped. All new edges and all diagonals of the pentagon are incident to one of the four spine vertices shown.

to think we can simply perform many of these swaps simultaneously. Unfortunately, this is not the case, since the sequence creates the edge $(v_{\text{out}}, v_{\text{in}})$. This means that trying to perform this sequence simultaneously in different locations would create a duplicate edge. Therefore we use a slightly longer sequence that creates a pentagon containing the edges to be swapped (illustrated in Figure 17), performs the swap inside this pentagon, and restores the canonical triangulation, using a total of eleven flips. The crucial property of this sequence is that it only creates edges incident to four spine vertices near the edge to be swapped. Thus, we can perform any number of swaps simultaneously without creating duplicate edges, as long as each swap is at distance four or more from the others. This means that, given a set of spine edges to swap, we can divide them into four rounds such that the edges to be swapped in each round are at distance four or more, and perform the swaps in each round simultaneously. Thus, we can swap any subset of spine edges with $O(1)$ simultaneous flips.

To scramble the edges incident on v_{out} , we first flip to create $(v_{\text{out}}, v_{\text{in}})$ and then apply the algorithm from Theorem 4 to the outerplanar graph induced by the edges incident to v_{out} . This uses $O(\log^2 n)$ simultaneous flips to rearrange all labels, except for those on the two outermost edges that are part of the boundary. In the non-simultaneous setting, we fixed this by swapping these labels along the spine, but this would take too many flips here. Instead, if the labels that need to be on the outermost edges are in the interior, we use Theorem 4 to place these labels on the interior edges closest to the outermost edges. Then, we can exchange them with the labels on the outermost edges with only three swaps. This ensures that the outermost edges have the correct labels, so a second application of Theorem 4 can place the remaining labels in the right order. If the label for one of the outermost edges is not in the interior and not already in place, it must be on the other outermost edge. In this case, we can first exchange it with the label on a nearby interior edge with a constant number of swaps. The entire sequence requires $O(\log^2 n)$ simultaneous flips.

Since these operations use $O(1)$ and $O(\log^2 n)$ simultaneous flips, and we can sort the labels with a constant number of applications, the theorem follows. \square

8. Conclusions

We have initiated the exploration of flips in edge-labelled triangulations and formulated the Orbit Conjecture: that for any two edge-labelled triangulations of a polygon, point set, or vertex set (in the combinatorial setting), there is a sequence of flips that transforms one into the other if and only if for every label, the initial and final edge with that label lie in the same orbit. Furthermore, we conjecture that the diameter of any connected component of the flip graph is bounded by a polynomial in the number of vertices.

We have established the conjecture for combinatorial triangulations and triangulations of convex polygons—all edges are in one orbit so the flip graph is connected, and its diameter is $\Theta(n \log n)$. This means that the worst case number of flips is a $\log n$ factor more than for the unlabelled case. With simultaneous flips, edge labels raise the worst case bound by at most a $\log n$ factor, but we could not prove that this is tight. We also established the conjecture for spiral polygons, characterizing the orbits and showing that each connected component of the flip graph has diameter $\Theta(n^2)$. For general polygons there are examples of a connected flip graph of diameter $\Theta(n^3)$.

References

- [1] G. Araujo-Pardo, I. Hubard, D. Oliveros, and E. Schulte. Colorful associahedra and cyclohedra. *J. Combin. Theory Ser. A*, 129:122–141, 2015.
- [2] M. A. Bender, D. Ge, S. He, H. Hu, R. Y. Pinter, S. Skiena, and F. Swidan. Improved bounds on sorting with length-weighted reversals. In *Proc. 15th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '04*, pages 919–928, 2004.
- [3] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. In *Computing in Euclidean geometry*, volume 1 of *Lecture Notes Ser. Comput.*, pages 23–90. 1992.
- [4] P. Bose, J. Czyzowicz, Z. Gao, P. Morin, and D. R. Wood. Simultaneous diagonal flips in plane triangulations. *J. Graph Theory*, 54(4):307–330, 2007.
- [5] P. Bose and F. Hurtado. Flips in planar graphs. *Comput. Geom.*, 42(1):60–80, 2009.
- [6] P. Bose, A. Lubiw, and J. I. Munro. Efficient visibility queries in simple polygons. *Comput. Geom.*, 23(3):313–335, 2002.
- [7] P. Bose, A. Lubiw, V. Pathak, and S. Verdonschot. Flipping edge-labelled triangulations. *ArXiv e-prints*, 2013. arXiv:1310.1166.

- [8] P. Bose and S. Verdonschot. A history of flips in combinatorial triangulations. In *Computational Geometry - XIV Spanish Meeting on Computational Geometry, EGC 2011, Dedicated to Ferran Hurtado on the Occasion of His 60th Birthday, Alcalá de Henares, Spain, June 27-30, 2011, Revised Selected Papers*, pages 29–44, 2011.
- [9] J. Cano, J.-M. Díaz-Báñez, C. Huemer, and J. Urrutia. The edge rotation graph. *Graphs Combin.*, 29(5):1207–1219, 2013.
- [10] J. Cardinal, M. Hoffmann, V. Kusters, C. D. Tóth, and M. Wettstein. Arc diagrams, flip distances, and Hamiltonian triangulations. In *32nd International Symposium on Theoretical Aspects of Computer Science*, pages 197–210, 2015.
- [11] H. Edelsbrunner. *Geometry and topology for mesh generation*. Cambridge University Press, Cambridge, 2001.
- [12] D. Eppstein. Happy endings for flip graphs. *J. Comput. Geom.*, 1(1):3–28, 2010.
- [13] G. Fertin, A. Labarre, I. Rusu, É. Tannier, and S. Vialette. *Combinatorics of genome rearrangements*. Computational Molecular Biology. MIT Press, Cambridge, MA, 2009.
- [14] J. Galtier, F. Hurtado, M. Noy, S. Pérennes, and J. Urrutia. Simultaneous edge flipping in triangulations. *Internat. J. Comput. Geom. Appl.*, 13(2):113–133, 2003.
- [15] W. H. Gates and C. Papadimitriou. Bounds for sorting by prefix reversal. *Discrete Mathematics*, 27:47–57, 1979.
- [16] S. Hanke. Ebene triangulationen. Master’s thesis, Insitut für Informatik, University of Freiburg, 1994.
- [17] C. Hernando, F. Hurtado, M. Mora, and E. Rivera-Campo. Grafos de árboles etiquetados y grafos de árboles geométricos etiquetados. In *Proc. X Encuentros de Geometra Computacional*, pages 13–19, 2003.
- [18] F. Hurtado, M. Noy, and J. Urrutia. Parallel edge flipping. In *Proc. 10th Canadian Conf. Comput. Geom.*, pages 26–27, 1998.
- [19] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete Comput. Geom.*, 22(3):333–346, 1999.
- [20] J. Kececioğlu and R. Ravi. Exact and approximation algorithms for sorting by reversals, with application to genome rearrangement. *Algorithmica*, 13:180–210, 1995.
- [21] C. Lawson. Software for C^1 surface interpolation. In *Mathematical Software III*, pages 161–194. Academic Press, New York, 1977.

- 850 [22] M. Molloy, B. Reed, and W. Steiger. On the mixing rate of the triangulation walk. In *Proc. Randomization Methods in Algorithm Design*, volume 43 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 179–190. Amer. Math. Soc., 1999.
- [23] V. Pathak. *Reconfiguring Triangulations*. PhD thesis, Cheriton School of Computer Science, University of Waterloo, Canada, 2014.
- 855 [24] R. Pinter and S. Skiena. Sorting with length-weighted reversals. In *Proc. 13th International Conference on Genome Informatics, GIW '02*, pages 173–182, 2002.
- [25] L. Pournin. The diameter of associahedra. *Adv. Math.*, 259:13–42, 2014.
- [26] D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1(3):647–681, 1988.
- 860 [27] D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Short encodings of evolving structures. *SIAM J. Discrete Math.*, 5(3):428–450, 1992.
- [28] K. Wagner. Bemerkungen zum Vierfarbenproblem. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 46:26–32, 1936.