A Proof of the Orbit Conjecture for Flipping Edge-Labelled Triangulations

Anna Lubiw¹ · Zuzana Masárová² · Uli Wagner²

Received: 23 September 2017 / Revised: 31 August 2018 / Accepted: 5 September 2018 /

Published online: 26 September 2018

© The Author(s) 2018

Abstract

Given a triangulation of a point set in the plane, a *flip* deletes an edge e whose removal leaves a convex quadrilateral, and replaces e by the opposite diagonal of the quadrilateral. It is well known that any triangulation of a point set can be reconfigured to any other triangulation by some sequence of flips. We explore this question in the setting where each edge of a triangulation has a label, and a flip transfers the label of the removed edge to the new edge. It is not true that every labelled triangulation of a point set can be reconfigured to every other labelled triangulation via a sequence of flips, but we characterize when this is possible. There is an obvious necessary condition: for each label l , if edge e has label l in the first triangulation and edge f has label l in the second triangulation, then there must be some sequence of flips that moves label l from e to f , ignoring all other labels. Bose, Lubiw, Pathak and Verdonchot formulated the *Orbit Conjecture*, which states that this necessary condition is also sufficient, i.e. that *all* labels can be simultaneously mapped to their destination if and only if *each* label individually can be mapped to its destination. We prove this conjecture. Furthermore, we give a polynomial-time algorithm (with $O(n^8)$ being a crude bound on the run-time) to find a sequence of flips to reconfigure one labelled triangulation to another, if such a sequence exists, and we prove an upper bound of $O(n^7)$ on the length of the flip sequence. Our proof uses the topological result that the sets of pairwise non-crossing edges on a planar point set form a simplicial complex that is homeomorphic to a high-dimensional ball (this follows from a result of Orden and Santos; we give a different proof based on a shelling argument). The dual cell complex of this simplicial ball, called the *flip complex*, has the usual flip graph as its 1-skeleton. We use properties of the 2-skeleton of the flip complex to prove the Orbit Conjecture.

Editor in Charge: Kenneth Clarkson

A preliminary version of this paper appeared at the 33rd International Symposium on Computational Geometry (SoCG 2017).

Extended author information available on the last page of the article

Keywords Triangulations · Reconfiguration · Flip · Constrained triangulations · Delaunay triangulation · Shellability · Piecewise linear balls

1 Introduction

The flip operation is fundamental to the study of triangulations of point sets in the plane. A flip removes one edge and replaces it by the opposite diagonal of the resulting quadrilateral, so long as that quadrilateral is convex. Lawson [21] proved the foundational result that any triangulation can be transformed into any other triangulation of the same point set via a sequence of flips. His second proof of this result [22] used the approach that is more widely known—showing that any triangulation can be flipped to the Delaunay triangulation, which then acts as a “hub” through which we can flip any triangulation to any other.

The result that every triangulation can be flipped to every other is captured succinctly by saying that the *flip graph* is connected, where the *flip graph* has a vertex for each triangulation of the given point set, and an edge when two triangulations differ by one flip. The special case of a point set in convex position has been very thoroughly studied. In this case triangulations correspond to binary trees, and a flip corresponds to a rotation. The flip graph in this case is the 1-skeleton of a polyhedron called the *associahedron*.

The use of flips to reconfigure triangulations is relevant to the study of associahedra [31] and mixing [25]. Flips are also important in practice for mesh generation and for finding triangulations that optimize certain quality measures [3,14]. The survey by Bose and Hurtado [6] discusses these and many other aspects of flips.

Despite the extensive work on flips, it is only recently that the question of where edges go under flip operations has been investigated. This can be formalized by attaching a label to each edge in a triangulation. Throughout, we fix a set P of n points in general position, and we identify triangulations with their edge sets (i.e., a triangulation of P is a maximal set T of pairwise non-crossing edges spanned by P). A *labelled triangulation* $\mathcal{T}$ of P is a pair (T, ℓ) where T is a triangulation of P and ℓ is a *labelling function* that maps the edges of T one-to-one onto the labels $1, 2, \dots, t_P$. Here t_P is the number of edges in any triangulation of P . When we perform a flip operation on $\mathcal{T}$, the label of the removed edge is transferred to the new edge.

We can now capture “where an edge goes” under flip operations. We say that edges e and f lie in the same *orbit* if we can attach label l to e in some triangulation and apply some sequence of flips to arrive at a triangulation in which edge f has label l . The orbits are exactly the connected components of a graph that Eppstein [15] called the *quadrilateral graph*—this graph has a vertex for every one of the possible $\binom{n}{2}$ edges formed by point set P , with e and f being adjacent if they cross and their four endpoints form a convex quadrilateral that is empty of other points. In particular, this implies that there is a polynomial-time algorithm to find the orbits. The orbits can be very different depending on P . For a point set in convex position, all the edges not in the convex hull are in a single orbit [7], but at the other extreme, a point set with no empty convex pentagon has the property that in any triangulation, the edges are all in distinct orbits [15].

Fig. 1 Five flips swap the edge labels (a and b) of two diagonals of a convex pentagon. In the flip graph these five flips form a 5-cycle

Orbits tell us where each individual edge label can go, but not how they combine. The main question we address in this paper is: when is there a sequence of flips to reconfigure one labelled triangulation of point set P to another labelled triangulation of P ? A necessary condition is that, for each label l , the edges with label l in the two triangulations must lie in the same orbit. Bose et al. [7] conjectured that this condition is also sufficient. As our main result we prove this “Orbit Conjecture,” and strengthen it by providing a polynomial-time algorithm and a bound on the length of the flip sequence.

Theorem 1.1 (Orbit Theorem) *Given two edge-labelled triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ of a point set, there is a flip sequence that transforms one into the other if and only if for every label l , the edges of $\mathcal{T}_1$ and $\mathcal{T}_2$ having label l belong to the same orbit. Furthermore, there is a polynomial-time algorithm that tests whether the condition is satisfied, and if it is, computes a flip sequence of length $O(n^7)$ to transform $\mathcal{T}_1$ to $\mathcal{T}_2$.*

The orbit theorem is stated for triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ that may have different edge sets, but—since we know how to use flips to change the edge set—the crux of the matter is the special case where the two triangulations have the same edge set T but different label functions ℓ_1 and ℓ_2 . In other words, we are given a permutation of the edge labels of a triangulation, and we seek a flip sequence to realize the permutation. Furthermore, since every permutation is a composition of transpositions, we concentrate first on finding a flip sequence to transpose (or “swap”) two labels. This idea of reducing the problem to the case of swaps appears in [7].

One insight to be gained from previous work is that empty convex pentagons in the point set seem to be crucial for swapping edge labels. Certainly, an empty convex pentagon provides a label swap—Fig. 1 shows how the edge labels of two diagonals of an empty convex pentagon can be swapped by a sequence of five flips. In the other direction, the special cases of the orbit theorem that were proved by Bose et al. [7] for convex and spiral polygons involved moving pairs of labels into empty convex pentagons and swapping them there. Furthermore, Eppstein [15] showed that in a triangulation of a point set with no empty convex pentagons, no permutations of edge labels are possible via flips.

The foundation of our proof is to make this intuition about empty convex pentagons rigorous. In particular, we show that the only elementary operation that is needed for label permutation is to transpose two labels by moving them into an empty convex pentagon and swapping them there. More formally, given a labelled triangulation $\mathcal{T} = (T, \ell)$, an *elementary swap* of edges e and f in T is a transposition of the labels of e and f that is accomplished as follows: perform a sequence, σ , of flips on $\mathcal{T}$ to get to a triangulation $\mathcal{T}'$ in which the labels $\ell(e)$ and $\ell(f)$ are attached to the two diagonals of an empty convex pentagon; then perform the 5-flip sequence, π , that

transposes these two labels; then perform the sequence σ^{-1} . We say that the sequence $\sigma\pi\sigma^{-1}$ *realizes* the elementary swap. Observe that the effect of $\sigma\pi\sigma^{-1}$ on $\mathcal{T}$ is to transpose the labels of e and f while leaving all other labels unchanged. Also notice that any other flip sequence that exchanges the labels of e and f must, by Theorem 1.2 below, be a sequence of elementary swaps. We will prove that an elementary swap can always be realized by a flip sequence of length $O(n^6)$, and furthermore, that such a sequence can be found in polynomial time.

One of our main results is the following, from which the Orbit Theorem can readily be derived:

Theorem 1.2 *In a labelled triangulation $\mathcal{T}$, two edges are in the same orbit if and only if there is an elementary swap between them.*

In order to prove Theorem 1.2, we use the following key result:

Theorem 1.3 (Elementary Swap Theorem) *Given a labelled triangulation $\mathcal{T}$, any permutation of the labels that can be realized by a sequence of flips can be realized by a sequence of elementary swaps.*

This theorem is proved using topological properties of the *flip complex*, whose 1-skeleton is the flip graph. A result of Orden and Santos [27] can be used to show that the flip complex has the topology of a high-dimensional ball.¹ We give an alternate proof of this. We use the 2-skeleton of the flip complex, and show that its 2-cells correspond to cycles in the flip graph of two types: quadrilaterals, which do not permute labels; and pentagons, which correspond precisely to the 5-cycles of flips shown in Fig. 1. Then we prove the Elementary Swap Theorem by translating it into a result about decomposing closed walks in the flip graph into simpler *elementary walks*.

Although there is a rich literature on associahedra and on cell complexes associated with triangulations of point sets, we are not aware of any previous combinatorial results on triangulations that require topological proofs, as our proof of the Orbit Theorem seems to.

We now briefly describe the rest of our method after the Elementary Swap Theorem is established. In order to prove Theorem 1.2, we need one more ingredient about the structure of elementary swaps: we will show that any sequence of elementary swaps that moves the label of edge e to edge f can be “completed” to get the label of f back to e , and that, in fact, the resulting sequence provides an elementary swap of e and f .

The high-level idea of our proof of Theorem 1.2 is then as follows: From our hypothesis that two edges e and f lie in the same orbit, we show that there is a sequence of flips that permutes the labels of triangulation $\mathcal{T}$, taking the label of e to f . The Elementary Swap Theorem then gives us a sequence of elementary swaps to do the same (this is the significant step of the proof). Finally, from the structure of elementary swaps we can then find an elementary swap of e and f .

Our paper is organized as follows. In Sect. 3 we prove the Elementary Swap Theorem using topological methods. In Sect. 4 we prove the properties of elementary swaps that were mentioned above. In top-down fashion, we begin in Sect. 2 by expanding on the high-level ideas, and proving the Orbit Theorem assuming the results in the later sections.

¹ Technically speaking, the flip complex is homotopy equivalent to a ball.

1.1 Background

The diameter of the flip graph of a point set gives the worst-case number of flips required to reconfigure one triangulation to another. For unlabelled triangulations, the diameter of the flip graph is known to be $\Theta(n^2)$, with the upper bound proved by Lawson [21] and the lower bound proved by Hurtado et al. [17]. For the special case of points in convex position, there is an exact bound of $2n - 10$ [29,31]. The problem of finding the distance in the flip graph between two given triangulations of a point set is NP-hard [23], and even APX-hard [28]. It has recently been shown to be fixed-parameter tractable [20]. The problem remains NP-hard for triangulations of a polygon [1], but the complexity status is open for the case of points in convex position. For further results on flips, see the survey by Bose and Hurtado [6].

The *labelled flip graph* of a point set has a vertex for every labelled triangulation of the point set and an edge between two labelled triangulations that differ by a flip. Bose et al. [7] formulated the Orbit Conjecture and proved it for the special case of triangulations of any convex polygon, showing that the labelled flip graph has a single connected component (ignoring convex hull edges, which cannot flip), and giving a tight bound of $\Theta(n \log n)$ on its diameter. Araujo-Pardo et al. [2] independently proved the Orbit Conjecture for convex polygons, and introduced “colorful associahedra” which generalize associahedra to the setting of labelled (or coloured) triangulations. Bose et al. also proved the Orbit Conjecture for spiral polygons. In this case the labelled flip graph may be disconnected but each connected component has diameter $O(n^2)$, which is a tight bound.

The best known lower bound on the diameter of a connected component of the labelled flip graph for a point set is $\Omega(n^3)$ [7]. There is a large gap between this lower bound and our upper bound of $O(n^7)$.

The Orbit Theorem holds for combinatorial triangulations [7], and for pseudotriangulations [8]. In both these cases there is a single orbit, so the labelled flip graph is connected. There are also some related results using variants of the flip operation, for example, Cano et al. [10] reconfigured edge-labelled non-maximal plane graphs by “rotating” edges around one of their endpoints; again there is a single orbit. A related result where there are multiple orbits is an analogue of the Orbit Theorem for labelled (or “ordered”) bases of a matroid—one labelled basis can be turned into another labelled basis via basis exchange steps if and only if elements with the same label lie in the same connected component of the matroid [24].

For more general problems of reconfiguring one structure to another via elementary steps, see [18,33].

1.2 Preliminaries and Definitions

Most definitions were given above, but we fill in a few missing details. Throughout, we assume a set of n points in general position in the plane. A point set determines $\binom{n}{2}$ *edges* which are the line segments between pairs of points. Two edges *cross* if they intersect in a point that is interior to at least one of the two edges. An *empty convex k -gon* is a subset of k points that forms a convex polygon with no point of P in its

interior. A *diagonal* of a convex polygon is an edge joining two points that are not consecutive on the polygon boundary.

Several times in our proofs we will use the result that if two unlabelled triangulations of the same point set have a subset, S , of *constrained* edges in common, then there is a sequence of flips that transforms one triangulation into the other, without ever flipping any edge of S , i.e. the edges in S remain fixed throughout the flip sequence. This was first proved by Dyn et al. [13], and can alternatively be proved using constrained Delaunay triangulations [3].

2 Proof of the Orbit Theorem

In this section we prove the Orbit Theorem assuming the Elementary Swap Theorem (Theorem 1.3, proved in Sect. 3), and assuming the following two results on elementary swaps. The first result shows that every elementary swap can be realized by a relatively short flip sequence that can be found efficiently, and the second result gives us a way to combine elementary swaps so that, after moving e 's label to f , we can get f 's label back to e . These lemmas will be proved in Sect. 4.

Lemma 2.1 *If there is an elementary swap between two edges in a triangulation $\mathcal{T}$ then there is a flip sequence of length $O(n^6)$ to realize the elementary swap, and, furthermore, this sequence can be found in polynomial time.*

Lemma 2.2 *Let $\mathcal{T}$ be a labelled triangulation containing two edges e and f . If there is a sequence of elementary swaps on $\mathcal{T}$ that takes the label of edge e to edge f , then there is an elementary swap of e and f in $\mathcal{T}$.*

As we show in Sect. 4, a simple group-theoretic argument suffices to prove a weaker version of Lemma 2.2, namely, that under the stated assumptions, there is a sequence of elementary swaps exchanging the labels of e and f in $\mathcal{T}$. Proving the stronger version, which we need for our bounds on the length of flip sequences, requires using the properties of elementary swaps.

We prove the Orbit Theorem in stages, first Theorem 1.2 (the case of swapping two labels in a triangulation), then the more general case of permuting edge labels in a triangulation, and finally the full result.

Proof of Theorem 1.2 The “if” direction is clear, so we address the “only if” direction. Suppose that $\mathcal{T} = (T, \ell)$ is the given edge-labelled triangulation and that e and f are edges of T that are in the same orbit. Then there is a sequence of flips that changes $\mathcal{T}$ to an edge-labelled triangulation $\mathcal{T}' = (T', \ell')$ where T' contains f and $\ell'(f) = \ell(e)$. We now apply the result that any constrained triangulation of a point set can be flipped to any other. Fix edge f and flip T' to T . Applying the same flip sequence to the labelled triangulation $\mathcal{T}'$ yields an edge-labelling of triangulation T in which edge f has the label $\ell(e)$. Thus we have a sequence of flips that permutes the labels of $\mathcal{T}$ and moves the label of e to f .

By the Elementary Swap Theorem (Theorem 1.3) there is a sequence of elementary swaps whose effect is to move the label of edge e to edge f . By Lemma 2.2 there is an elementary swap of e and f in $\mathcal{T}$. $\square$

Theorem 2.3 (Edge Label Permutation Theorem) *Let T be a triangulation of a point set with two edge-labellings ℓ_1 and ℓ_2 such that for each label l , the edge with label l in ℓ_1 and the edge with label l in ℓ_2 are in the same orbit. Then there is a sequence of $O(n)$ elementary swaps to transform the first labelling to the second. Such a sequence can be realized via a sequence of $O(n^7)$ flips, which can be found in polynomial time.*

Proof The idea is to effect the permutation as a sequence of swaps. If every edge has the same label in ℓ_1 and ℓ_2 we are done. So consider a label l that is attached to a different edge in ℓ_1 and in ℓ_2 . Suppose $\ell_1(e) = l$ and $\ell_2(f) = l$, with $e \neq f$. By hypothesis, e and f are in the same orbit. By Theorem 1.2 there is an elementary swap of e and f in (T, ℓ_1) which results in a new labelling ℓ'_1 that matches ℓ_2 in one more edge (namely the edge f) and still has the property that for every label l , the edge with label l in ℓ'_1 and the edge with label l in ℓ_2 are in the same orbit. Thus we can continue this process until all edge labels match those of ℓ_2 . In total we use $O(n)$ elementary swaps. These can be realized via a sequence of $O(n^7)$ flips by Lemma 2.1. Furthermore, the sequence can be found in polynomial time. $\square$

We can now prove the Orbit Theorem.

Proof of Theorem 1.1 The necessity of the condition is clear, and we can test it in polynomial time by finding all the orbits, so we address sufficiency. The idea is to reconfigure $\mathcal{T}_1$ to have the same underlying unlabelled triangulation as $\mathcal{T}_2$ and then apply the previous theorem. The details are as follows. Let $\mathcal{T}_1 = (T_1, \ell_1)$ and $\mathcal{T}_2 = (T_2, \ell_2)$. There is a sequence σ of $O(n^2)$ flips to reconfigure the unlabelled triangulation T_1 to T_2 , and σ can be found in polynomial time. Applying σ to the labelled triangulation $\mathcal{T}_1$ yields a labelled triangulation $\mathcal{T}_3 = (T_2, \ell_3)$. Note that for every label l , the edges of $\mathcal{T}_1$ and $\mathcal{T}_3$ having label l belong to the same orbit. This is because flips preserve orbits (by definition of orbits). Thus by Theorem 2.3 there is a flip sequence τ that reconfigures $\mathcal{T}_3$ to $\mathcal{T}_2$, and this flip sequence can be found in polynomial time and has length $O(n^7)$. The concatenation of the two flip sequences, $\sigma\tau$, reconfigures $\mathcal{T}_1$ to $\mathcal{T}_2$, has length $O(n^7)$, and can be found in polynomial time. $\square$

3 Proof of the Elementary Swap Theorem

As mentioned in the introduction, we prove the Elementary Swap Theorem using topological properties of the *flip complex*, whose 1-skeleton (i.e. vertices and edges) is the flip graph. In fact, we will only need the 2-cells of the flip complex, not any higher-dimensional structure. We will show that 2-cells of the flip complex correspond to 4- and 5-cycles in the flip graph.

The basic idea is as follows. We will translate the Elementary Swap Theorem to a statement about walks in the flip graph. The hypothesis of the Elementary Swap Theorem is that we have a sequence of flips that permutes the edge labels of a triangulation T . In the flip graph, this sequence corresponds to a closed walk w that starts and ends at triangulation T . Our main topological result is that the flip complex has a trivial fundamental group, which will imply that such a closed walk w can be decomposed into simpler *elementary walks*. Each elementary walk starts at T , traces a path in the

Fig. 2 (a) Triangulations that differ in the diagonals of two internally disjoint quadrilaterals form an *elementary 4-cycle* in the flip graph. The cycle does not permute the labels (shown as red and blue). (b) Triangulations that differ in the diagonals of a convex pentagon form an *elementary 5-cycle* in the flip graph. This cycle permutes labels as shown in Fig. 1

flip graph, then traverses the edges of a 2-cell, then retraces the path back to T . The edge-label permutation induced by an elementary walk depends on the 2-cell. If the 2-cell is a 4-cycle, the permutation is the identity; and if the 2-cell is a 5-cycle, then the permutation is a transposition, and the elementary walk corresponds to an elementary swap. Altogether, this implies that the permutation induced by the closed walk w can be expressed as a composition of elementary swaps, which proves the Elementary Swap Theorem.

Before stating our main topological theorem, we first define the special cycles that will be shown to correspond to 2-cells of the flip complex. In the same way that an edge of the flip complex corresponds to two triangulations that differ on one edge, every 2-cell of the flip complex corresponds to a set of triangulations that differ on two edges. Define an *elementary 4-cycle* to be a cycle of the flip graph obtained in the following way. Take a triangulation T and two edges $e, f \in T$ whose removal leaves two internally disjoint convex quadrilaterals in T . Each quadrilateral can be triangulated in two ways, which results in four triangulations that contain $F := T \setminus \{e, f\}$. These four triangulations form a 4-cycle in the flip graph, as shown in Fig. 2(a). Observe that a traversal of the cycle corresponds to a sequence of flips that returns edge-labels to their original positions.

Define an *elementary 5-cycle* to be a cycle of the flip graph obtained in the following way. Take a triangulation T and two edges $e, f \in T$ whose removal leaves a convex pentagon in T . There are five triangulations that contain $F := T \setminus \{e, f\}$, and they form a 5-cycle in the flip graph, as shown in Fig. 2(b). Observe that the sequence of flips around such a cycle permutes labels of e and f as shown in Fig. 1.

As a side remark, note that it can be shown that, in fact, any cycle in the flip graph of length less than 6 is an elementary 4- or 5-cycle. However, we will not need this in what follows.

Our main topological theorem is the following.

Theorem 3.1 *Let P be a set of n points in general position in the plane. There is a high-dimensional cell complex $\mathbb{X} = \mathbb{X}(P)$, which we call the flip complex, such that:*

1. *The 1-skeleton of $\mathbb{X}$ is the flip graph of P ;*
2. *There is a one-to-one correspondence between the 2-cells of $\mathbb{X}$ and the elementary 4-cycles and elementary 5-cycles of the flip graph of P ;*
3. *$\mathbb{X}$ has the topology of (i.e., is homotopy equivalent to) a high-dimensional ball; therefore its fundamental group, $\pi_1(\mathbb{X})$, is trivial.*

In what follows, we will use a number of notions from combinatorial topology; some of these we will recall along the way, but others we will only describe informally or leave undefined and instead refer the reader to standard textbooks for further background (in particular, we refer the reader to [5, Appendix 4.7] and [16] for background on *regular cell complexes*, *shellability*, and *piecewise linear balls and spheres*, to [32] for background on the fundamental group of cell complexes, and to [16, 26] for background on *dual complexes*; we will provide more detailed references for specific results below).

Theorem 3.1 follows from a result of Orden and Santos [27], see Remark 3.9 at the end of Sect. 3.3 below for more details; we are grateful to F. Santos for bringing this reference to our attention.

Before becoming aware of the work of Orden and Santos, we found a different proof of Theorem 3.1 that starts out by considering the simplicial complex $\mathbb{T} = \mathbb{T}(P)$ whose faces are the sets of pairwise non-crossing edges (line segments) spanned by P . This complex $\mathbb{T}$ is shown to be a *shellable simplicial ball* (by an argument based on constrained Delaunay triangulations), and $\mathbb{X}$ is then constructed as the *dual complex* of $\mathbb{T}$. We hope that this alternative proof of Theorem 3.1 is of some independent interest and present it in Sects. 3.2 and 3.3 below. Before that, in Sect. 3.1, we show how to derive the Elementary Swap Theorem from Theorem 3.1.

3.1 From Topology to the Elementary Swap Theorem

In this section we use Theorem 3.1 to prove the Elementary Swap Theorem. We begin by defining elementary walks. A *walk* in the flip graph is a sequence $T_0, T_1, \dots, T_k$ of triangulations (possibly with repetitions) such that T_{i-1} and T_i differ by a flip. We will refer to T_0 and T_k as the start and the end of the walk, respectively. A walk is *closed* if it starts and ends at the same triangulation. If w_1 and w_2 are walks such that the end of w_1 equals the start of w_2 then we can define their *composition* $w_1 w_2$ in the obvious way. Furthermore, if $w = (T = T_0, T_1, \dots, T_k)$ is a walk, we will use the notation $w^{-1} = (T_k, T_{k-1}, \dots, T_0)$ for the *inverse walk*.

Fix a triangulation T_0 . An *elementary quadrilateral walk* is a closed walk of the form $wz w^{-1}$, where z is an elementary 4-cycle in the flip graph, and w is a walk from T_0 to some triangulation on z . An *elementary pentagonal walk* is defined analogously, with z an elementary 5-cycle.

It is straightforward to check the effect of these elementary walks on labellings:

Lemma 3.2 *Let (T_0, ℓ) be a labelled triangulation. An elementary quadrilateral walk does not permute the labels. An elementary pentagonal walk swaps the labels of two edges (e and f in Fig. 2(b)) and leaves all other labels fixed; this corresponds exactly to the notion of an elementary swap introduced earlier.*

Another operation that does not affect the permutation of labels induced by a closed walk is the following. A *spur* ww^{-1} starting and ending at T is an arbitrary walk w starting at T , immediately followed by the *inverse walk*. If w_1 and w_2 are walks in the flip graph such that w_1 ends at a triangulation T and w_2 starts there, and if s is a spur at T , then we say that the walk w_1sw_2 differs from w_1w_2 by a *spur insertion*. The inverse operation is called a *spur deletion*.

Lemma 3.3 *If two closed walks w and w' in the flip graph differ only by a finite number of spur insertions and deletions then they yield the same permutation of edge labels.*

Proof A flip immediately followed by its inverse flip has no effect on labels. The lemma follows by induction on the length of a spur and the number of spur insertions and deletions. $\square$

By Lemmas 3.2 and 3.3, the Elementary Swap Theorem directly reduces to the following, which we prove using Theorem 3.1:

Proposition 3.4 *Let w be a closed walk in the flip graph starting and ending at T_0 . Then, up to a finite number of spur insertions and deletions, w can be written as the composition of finitely many elementary walks.*

Proof We use the well-known fact that the fundamental group of a cell complex can be defined *combinatorially* in terms of closed walks in the 1-skeleton and this definition is equivalent to the usual topological definition in terms of continuous loops, see [30, Chap. 7] or [32, Chap. 4]. In particular, in a cell complex with trivial fundamental group any two closed walks in the 1-skeleton starting at the same vertex are related by a finite number of spur insertions, deletions and so-called 2-cell relations.

We describe the combinatorial definition of the fundamental group of the flip complex $\mathbb{X}$ in detail. By Theorem 3.1, the 1-skeleton of $\mathbb{X}$ is the flip graph of P . Fix a *base triangulation* T_0 , and, for every triangulation T , fix a walk p_T from T_0 to T . Given two triangulations T_1, T_2 that differ by a flip, we form the closed walk w_{T_1, T_2} in the flip graph, called a *generating walk*, that goes from T_0 to T_1 along p_{T_1} , then flips to T_2 , and then returns to T_0 along $p_{T_2}^{-1}$. It is easy to see that, up to a finite number of spur insertions and deletions, every closed walk starting and ending at T_0 can be written as a composition of generating walks.

We say that walks w and w' are *2-cell related* if we can express them as $w = w_1w_2$ and $w' = w_1zw_2$, where z is a closed walk traversing the boundary of a 2-cell (an elementary cycle) exactly once in either orientation. Notice that w_1w_2 and $w_1zz^{-1}w_2$ differ only by the spur zz^{-1} , hence, up to spur insertion and deletion, being 2-cell related is symmetric.

Also, notice the *precomposition property*: if w and w' are 2-cell related as above and if $w = w_1w_2$ is precomposed with the closed walk $w_1zw_1^{-1}$ then the result $w'' = (w_1zw_1^{-1})w = w_1z(w_1^{-1}w_1)w_2$ differs from w' only by the spur $w_1^{-1}w_1$. By

Theorem 3.1, a boundary of a 2-cell is an elementary 4- or 5-cycle and so the walk $w_1 z w_1^{-1}$ above is an elementary walk.

Two walks in the flip graph are called *equivalent* if they differ by a finite number of spur insertion and/or deletions and by applying a finite number of 2-cell relations. It is not hard to check that this defines an equivalence relation, and the fundamental group $\pi_1(\mathbb{X})$ is given as the set of equivalence classes of closed walks starting and ending at T_0 .

By Theorem 3.1, the fundamental group of the flip complex $\mathbb{X}$ is trivial. This translates into the fact that every closed walk starting and ending at T_0 is equivalent to the trivial walk. By the precomposition property, this means that, up to a finite number of spur insertions and deletions, every closed walk is a composition of finitely many elementary walks. $\square$

3.2 The Simplicial Complex of Plane Graphs

In this section and the following one, we give a proof of Theorem 3.1. This section is about the simplicial complex $\mathbb{T} = \mathbb{T}(P)$ whose faces are the sets of pairwise non-crossing edges (line segments) spanned by P .

Let P be a set of n points in general position in the plane, i.e., no three points lie on a line and no four points lie on a common circle. Let E be the set of edges (closed line segments) spanned by P . Two edges $e, f \in E$ are said to be *non-crossing* if they are disjoint or if they intersect in a single point of P that is an endpoint of both edges. We say that a subset $F \subseteq E$ is *non-crossing* if every pair of distinct edges $e, f \in F$ is non-crossing. If G is non-crossing and $F \subseteq G$ then F is non-crossing as well. Thus, the non-crossing sets of edges form an abstract simplicial complex

$$\mathbb{T} = \mathbb{T}(P) := \{F : F \subseteq E, F \text{ non-crossing}\},$$

which we call the *complex of plane graphs on P* . We collect some basic properties of $\mathbb{T}$:

1. The *facets* (inclusion-maximal faces) of $\mathbb{T}$ are exactly the triangulations of P (every non-crossing set of edges $F \subseteq E$ can be extended to a triangulation). Thus, the simplicial complex $\mathbb{T}$ is of dimension $m - 1$, where m is the number of edges in any triangulation of P , and it is *pure*, i.e., every face of $\mathbb{T}$ is contained in a face of dimension $m - 1$.
2. Every face F of $\mathbb{T}$ of dimension $m - 2$ is contained in either one or two triangulations. In the latter case, F corresponds to a flip between these two triangulations.

We will show that the topology of $\mathbb{T}$ is particularly simple, namely that $\mathbb{T}$ is homeomorphic to an $(m - 1)$ -dimensional ball. Furthermore, there is a combinatorial certificate (*shellability*) for this homeomorphism. This implies that the homeomorphism is particularly nice and that $\mathbb{T}$ is a *piecewise-linear ball*, which means that there is a subdivision $\mathbb{T}'$ of $\mathbb{T}$ such that $\mathbb{T}'$ is simplicially isomorphic to a subdivision $\mathbb{B}'$ of the d -dimensional simplex Δ^d . The only property of piecewise-linearity that we will need is that it ensures that the construction of the *dual cell complex* $\mathbb{T}^*$ is well-behaved (see Proposition 3.8 below and Footnote 3 preceeding it).

We refer to [9,16,19] and [5, Appendix 4.7] for more details and further references on shellability and piecewise-linear balls, spheres, and manifolds.

We recall that a pure d -dimensional simplicial complex $\mathbb{K}$ is *shellable* if there exists a total ordering of its facets $F_1, F_2, \dots, F_N$ (called a *shelling order*) such that, for every $2 \leq j \leq N$, the intersection of F_j with the simplicial complex generated by the preceding facets is pure of dimension $d - 1$, i.e., for every $i < j$ and $F := F_i \cap F_j$, there exists some $k < j$ such that $G := F_j \cap F_k$ is of dimension $d - 1$ and $F \subseteq G$.

We will need the following result (which appears implicitly in [4], and explicitly in [11]; see [5, Prop. 4.7.22] for a short proof).

Proposition 3.5 *Suppose $\mathbb{K}$ is a finite d -dimensional simplicial complex that is a pseudomanifold, i.e., $\mathbb{K}$ is pure and every $(d - 1)$ -dimensional face of $\mathbb{K}$ is contained in at most two d -faces. If $\mathbb{K}$ is shellable then $\mathbb{K}$ is either a piecewise-linear ball or a piecewise-linear sphere. The former case occurs iff there is at least one $(d - 1)$ -dimensional face that is contained in only one d -face of $\mathbb{K}$, in which case the pseudomanifold is said to have nonempty boundary.²*

Theorem 3.6 *$\mathbb{T}$ is a shellable $(m - 1)$ -dimensional pseudomanifold with nonempty boundary, and hence a piecewise-linear ball.*

Proof We observed earlier that $\mathbb{T}$ is a pure $(m - 1)$ -dimensional simplicial complex, and that every $(m - 2)$ -dimensional face of $\mathbb{T}$ is contained in at most two $(m - 1)$ -dimensional faces, hence $\mathbb{T}$ is a pseudomanifold. Moreover, if T is a triangulation of P and if $e \in T$ is a non-flippable edge (e.g., if e is a convex hull edge) then $F := T \setminus \{e\}$ is an $(m - 2)$ -dimensional face of $\mathbb{T}$ that is contained in a unique $(m - 1)$ -face, namely T . Hence, $\mathbb{T}$ has nonempty boundary.

Thus, by Proposition 3.5, it suffices to show that $\mathbb{T}$ is shellable, i.e., to exhibit a shelling order for the facets of $\mathbb{T}$.

With every triangulation T of P , we associate the sorted vector of angles $\alpha(T) = (\alpha_1(T), \alpha_2(T), \dots, \alpha_{3t}(T))$, where $\alpha_1(T) \leq \alpha_2(T) \leq \dots \leq \alpha_{3t}(T)$ are the angles occurring in the triangulation T . We order the triangulations of P by sorting the corresponding angle vectors $\alpha(T)$ lexicographically from largest to smallest. Since we assume P to be in general position, this defines a total ordering

$$T_1, T_2, \dots, T_N, \quad \alpha(T_1) >_{\text{LEX}} \alpha(T_2) >_{\text{LEX}} \dots >_{\text{LEX}} \alpha(T_N), \quad (1)$$

where N is the number of triangulations of P .

It is well known (see, for example, [12, Chap. 3.4]) that in this ordering, T_1 is the Delaunay triangulation of P . Moreover, if we consider only triangulations containing a particular plane subgraph corresponding to a face F of $\mathbb{T}$ and the corresponding subsequence of the angle vectors, the first of these vectors corresponds to the Delaunay triangulation constrained to F .

We claim that the triangulation ordering (1) defines a shelling. We need to prove that for every $i < j \leq N$ and $F := T_i \cap T_j$, there exists some $k < j$ such that $G := T_k \cap T_j$ is of dimension $m - 2$ and $F \subseteq G$.

² We remark that the property of being a shellable pseudomanifold (which is a combinatorial and algorithmically verifiable condition) is strictly stronger than being a piecewise-linear ball or sphere, which in turn is strictly stronger than being a simplicial complex homeomorphic to a ball or sphere.

To see this, consider the subsequence $T_{k_1}, T_{k_2}, \dots$ of the sequence (1) consisting only of those triangulations that contain the edge set F . Then T_{k_1} is the constrained Delaunay triangulation with respect to the edge set F , and T_i and T_j both appear in that subsequence; in particular, $T_j \neq T_{k_1}$ since T_i precedes it. Since every triangulation containing F can be transformed to the constrained Delaunay triangulation T_{k_1} , (see, e.g., the description of the Lawson flip algorithm in [12]) there must exist an edge $e \in T_j \setminus T_{k_1}$ such that flipping e (a Lawson flip) increases the angle vector; thus, the triangulation resulting from flipping e is some T_k with $k < j$ and satisfies $F \subseteq T_k \cap T_j$ as desired. $\square$

Finally, we need a characterization of interior versus boundary faces of $\mathbb{T}$. Let $\mathbb{B}$ be a piecewise-linear ball of dimension d . By definition, the *boundary* $\partial\mathbb{B}$ of $\mathbb{B}$ is the subcomplex of $\mathbb{B}$ consisting of all faces F for which there exists a $(d-1)$ -dimensional face G of $\mathbb{B}$, with $F \subseteq G$, such that G is contained in a unique d -dimensional face of $\mathbb{B}$, see, for example, [9, 19]. (In the case $\mathbb{B} = \mathbb{T}$, the latter condition means that $G = T \setminus \{e\}$ for some triangulation T and some edge $e \in T$ that is not flippable.) A face F of $\mathbb{B}$ that does not lie in $\partial\mathbb{B}$ is called an *interior face*.

For the proof of Theorem 3.1 we need properties of interior faces of $\mathbb{T}$ of dimensions $m-1$, $m-2$ and $m-3$. The following proposition characterizes interior faces more generally.

Proposition 3.7 *Let $\mathbb{T}$ be the simplicial complex of plane graphs on the point set P . A non-crossing set of edges F on P is an interior face of $\mathbb{T}$ if and only if the following conditions hold:*

- (i) *F contains all convex hull edges of P ,*
- (ii) *Every bounded region in the complement of the plane graph (P, F) is convex.*

Proof Note that a polygon is non-convex iff it has a reflex vertex. More generally, a bounded region in the complement of the plane graph (P, F) is non-convex iff there is an interior point p of P and a half-plane H through p with no edge of F from p to a point interior to H —in this case we say that p “has no edge in a half-plane”. The statement of the proposition is then equivalent to the following: F is a boundary face if and only if F misses a convex hull edge or there is an interior point p of P with no edge in a half-plane. We prove this statement.

For the forward direction, suppose that F is a boundary face. Then there is a triangulation T , $F \subseteq T$, and an edge $e \in T - F$ such that e is not flippable in T . If e is a convex hull edge, then F does not contain all convex hull edges. Otherwise e is a diagonal of a non-convex quadrilateral in T . Set p to be the reflex vertex of the non-convex quadrilateral and H to contain the other end of e but not the two other vertices of the quadrilateral. Then p has no edge in the half-plane H .

For the other direction, first note that if F misses a convex hull edge then F is a boundary face. For the other case, suppose there is a non-convex hull point p of P that has no edge in the half-plane H . Augment F to a maximal set F' of non-crossing edges without using any edge from p into H . This will not yet be a triangulation (because in a triangulation p is surrounded by triangles and they have angles bounded by π). Now augment F' further to a triangulation T . Then $T - F'$ contains some edge

e incident to p , and e is not flippable otherwise we could have further augmented F' . Thus F is a boundary face. $\square$

3.3 The Dual Flip Complex $\mathbb{X}$

To define the flip complex $\mathbb{X}$, we need the notion of *dual cells* and the dual cell decomposition of a piecewise-linear ball; for the precise definition, we refer to [16, Sect. I.6] or [26, §64 and §70].³ Here, we simply collect the properties that we will need:

Proposition 3.8 *Let $\mathbb{B}$ be a d -dimensional piecewise-linear ball.*

1. *For each interior k -dimensional face F of $\mathbb{B}$, one can define a dual cell F^* (a certain subcomplex of the barycentric subdivision of $\mathbb{B}$ that is a piecewise-linear ball of dimension $d - k$ [16, Lem. I.19]).*
2. *The construction reverses inclusion, i.e., for interior faces F, G of $\mathbb{B}$, $F \subseteq G$ iff $F^* \supseteq G^*$.*
3. *The dual cells of the interior faces of $\mathbb{B}$ form a regular cell complex, denoted $\mathbb{B}^*$ and called the dual cell complex. $\mathbb{B}^*$ need not be a manifold or pure d -dimensional, but it is homotopy equivalent to $\mathbb{B}$ [26, Lem. 70.1].⁴*

We define the *flip complex* $\mathbb{X} := \mathbb{T}^*$ as the dual complex of the simplicial complex $\mathbb{T}$.

Proof of Theorem 3.1 By Proposition 3.8, $\mathbb{X} = \mathbb{T}^*$ is a regular cell complex that is homotopy equivalent to the ball $\mathbb{T}$; consequently, the fundamental group $\pi_1(\mathbb{X})$ vanishes.

It remains to show the characterization of the vertices, edges, and 2-cells of $\mathbb{X}$.

The vertices of $\mathbb{X}$ correspond (are dual) to the faces of $\mathbb{T}$ of the highest dimension $(m - 1) = \dim \mathbb{T}$, i.e., to the triangulations of P (these are automatically interior faces of $\mathbb{T}$).

The edges of $\mathbb{X}$ correspond to interior $(m - 2)$ -dimensional faces F of $\mathbb{T}$, i.e., faces F that are contained in two triangulations T and T' that differ by a flip. Thus, the 1-skeleton of $\mathbb{X}$ is exactly the flip graph of P .

Every 2-cell of $\mathbb{X}$ is the dual cell F^* of an interior face F of $\mathbb{T}$ of dimension $m - 3 = \dim F$. Consider an arbitrary triangulation T containing F , i.e., F is obtained from T by deleting two edges e, f . By Proposition 3.7, e and f are both flippable in T since they lie in a convex polygon in T .

If e and f are not incident to a common triangle in T , (or, equivalently, removing both e and f from T creates two internally disjoint convex quadrilaterals) then there

³ In [26], the terminology *dual blocks* is used instead of dual cells, since the construction is described in a more general setting (for arbitrary triangulated manifolds or homology manifolds) in which the dual blocks F^* need not be cells (homeomorphic to balls). However, in the setting of piecewise-linear manifolds, in particular of piecewise-linear balls, as in our case, this technical issue does not arise.

⁴ More specifically, the dual complex of a piecewise-linear manifold with boundary is a deformation retraction of the manifold. For manifolds without boundary, the dual complex is piecewise-linearly homeomorphic to the original manifold.

exist four triangulations containing F and these form an elementary 4-cycle in the flip graph. It follows from the definition of $\mathbb{X} = \mathbb{T}^*$ that the 4-cycle is the boundary of the dual cell F^* .

Otherwise, e and f are incident to a common triangle in T . By Proposition 3.7 the union of the three triangles of T containing either e or f forms a convex polygon, necessarily a pentagon. There are five triangulations containing F and these form an elementary 5-cycle in the flip graph. It follows from the definition of $\mathbb{X} = \mathbb{T}^*$ that the 5-cycle is the boundary of the dual cell F^* .

Hence, every 2-cell of $\mathbb{X}$ corresponds to an elementary 4- or 5-cycle of the flip graph.

Conversely, every elementary 4- or 5-cycle of the flip graph gives rise to a 2-cell F^* of $\mathbb{X}$: more precisely, F^* corresponds to the intersection of the triangulations in the elementary cycle. $\square$

Remark 3.9 As remarked above, the flip complex $\mathbb{X}$ and Theorem 3.1 are closely related to a result of Orden and Santos [27]. Specifically, Orden and Santos showed that for every point set P , there exist a simple polytope $\mathbb{Y} = \mathbb{Y}(P)$ and a distinguished face F_0 of $\mathbb{Y}$ with the following properties: The vertices of $\mathbb{Y}$ that do *not* lie in the distinguished face F_0 are in one-to-one correspondence with the triangulations of P .⁵ More generally, the faces of $\mathbb{Y}$ that are disjoint from the distinguished face F_0 are in one-to-one correspondence with the non-crossing sets of edges of P that contain all convex hull edges of P . Furthermore, the correspondence reverses inclusion.

It follows from this that the cell complex $\mathbb{K} = \mathbb{K}(P)$ of all faces of $\mathbb{Y}$ disjoint from F_0 has the flip graph of P as its 1-skeleton, and the fundamental group of $\mathbb{K}$ is trivial (since $\mathbb{K}$ is the complement of the star of a face in the boundary of a convex polytope); analogously to the proof of Theorem 3.1, it can be shown that the 2-faces of $\mathbb{K}$ correspond to the elementary 4-cycles and 5-cycles in the flip graph. Thus, the complex $\mathbb{K}$ could be used instead of the flip complex $\mathbb{X}$ to prove the Elementary Swap Theorem.

A different way of viewing the complex $\mathbb{K}$ is as follows: Let C be the set of convex hull edges of P . Then C is a face of the complex $\mathbb{T}$ of plane graphs on P , and since $\mathbb{T}$ is a shellable ball, the link $\mathbb{L}$ of C in $\mathbb{T}$ is a shellable (and hence piecewise-linear) ball or sphere. The complex $\mathbb{K}$ of Orden–Santos is the dual cell complex of $\mathbb{L}$.

4 Proofs of Properties of Elementary Swaps

In this section we prove Lemmas 2.1 and 2.2.

To prove Lemma 2.1, the idea is to look at paths in the *double quadrilateral graph* G_D that we will define below. Informally speaking, G_D captures where pairs of non-crossing edges can go via flips, similar to the way the quadrilateral graph captures where a single edge can go via flips. We will show that there is an elementary swap

⁵ The vertices of F_0 , by contrast, correspond to *pseudotriangulations* with vertex set P that are not triangulations. Here, a pseudotriangulation of P is a decomposition of the convex hull of P into *pseudotriangles*, i.e., possibly non-convex polygons with exactly three non-reflex vertices.

between two labels in a triangulation if and only if there exists a path of certain type in the double quadrilateral graph.

Proof of Lemma 2.1 Construct a graph G_D called the *double quadrilateral graph*. Vertices of the graph G_D are pairs of non-crossing edges on the point set P , and we define two vertices (e_1, f_1) and (e_2, f_2) of G_D to be adjacent if either $e_1 = e_2$ and f_1 and f_2 are adjacent in the quadrilateral graph, or if $f_1 = f_2$ and e_1 and e_2 are adjacent in the quadrilateral graph. (Recall that two edges a and b are adjacent in the quadrilateral graph if a and b cross and their four endpoints form an empty quadrilateral.)

In the graph G_D we identify some vertices as “swap vertices”. These are the vertices (g, h) such that g and h are diagonals of some empty convex pentagon in the point set. Note that the swap vertices can be identified in polynomial time.

We claim that there is an elementary swap of e and f in the labelled triangulation $\mathcal{T} = (T, \ell)$ if and only if there is a path in G_D from vertex (e, f) to a swap vertex. For the forward direction, suppose there is such an elementary swap. It begins with a sequence σ of flips from $\mathcal{T}$ to a labelled triangulation $\mathcal{T}'$ in which labels $\ell(e)$ and $\ell(f)$ are attached to two diagonals g and h of some empty convex pentagon. The subsequence of σ consisting of those flips that apply to an edge whose current label is $\ell(e)$ or $\ell(f)$ corresponds to a path in G_D from (e, f) to the swap vertex (g, h) .

For the other direction, let π be a path in G_D from (e, f) to a swap vertex. It suffices to show that the path π provides a sequence of flips, σ , that takes $\mathcal{T}$ to some labelled triangulation $\mathcal{T}'$ in which labels $\ell(e)$ and $\ell(f)$ are attached to two diagonals of an empty convex pentagon, because the rest of the elementary swap is then determined. Consider the first edge of π and suppose without loss of generality that it goes from (e, f) to (e, f') (the case when e changes is similar). Then e and f' are non-crossing. Because f and f' are adjacent in the quadrilateral graph, they cross and form an empty convex quadrilateral Q . Note that e does not intersect the interior of Q , since Q is empty and e does not cross f or f' . We apply the result that any constrained triangulation can be flipped to any other with $O(n^2)$ flips. Fix edges e and f in T and flip T to a labelled triangulation that contains the edges of Q . In this triangulation, we can flip f to f' , transferring $\ell(f)$ to f' . We continue in this way to realize each edge of π via $O(n^2)$ flips, arriving finally at a labelled triangulation in which labels $\ell(e)$ and $\ell(f)$ are attached to edges that are the diagonals of some empty convex pentagon in the point set. Fixing the two diagonals, we can flip to a triangulation that contains the edges of the convex pentagon, and at this point we are done.

Because the graph G_D has $O(n^4)$ vertices, the diameter of any of its connected components is $O(n^4)$. Thus, if there is an elementary swap that exchanges the labels of edges e and f , then there is one corresponding to a path in G_D of length $O(n^4)$. We can explicitly construct G_D and find such a path in polynomial time. As argued above, every edge of G_D can be realized by $O(n^2)$ flips. This proves that, for any elementary swap, we can construct a sequence of $O(n^6)$ flips to realize it, and the construction takes polynomial time. $\square$

As mentioned in Sect. 2, there is a group-theoretic argument proving a weaker version of Lemma 2.2. The argument depends on the following claim: If a permutation group is generated by transpositions and contains a permutation that maps element e to f then the group contains the transposition of e and f . To prove this claim, notice that

if the group contains transpositions (ab) and (bc) , then it also contains transposition $(ac) = (ab)(bc)(ab)$; and apply induction.

To apply this claim in our situation, observe that by the Elementary Swap Theorem, all label permutations achievable by flips in a triangulation $\mathcal{T}$ are compositions of elementary swaps, hence, these label permutations indeed form a group G generated by transpositions. Moreover, by the assumption of Lemma 2.2, G contains a permutation taking the label of edge e to edge f . Hence, by the above claim, the group G also contains a permutation, which is a composition of elementary swaps, whose effect is to transpose labels of edges e and f .

In order to prove the full result of Lemma 2.2, i.e., that the label transposition of e and f can be done with a single elementary swap, we combine the techniques used in the proof of the group theory claim above with the structure of elementary swaps.

Proof of Lemma 2.2 An elementary swap in triangulation $\mathcal{T}$ acts on two edges of $\mathcal{T}$. We define a graph G_S called the *elementary swap graph* of $\mathcal{T}$. G_S has a vertex for every edge of $\mathcal{T}$, and we define vertices e and f to be adjacent in G_S if there is an elementary swap of e and f in $\mathcal{T}$.

By hypothesis, there is a sequence of elementary swaps that takes the label of edge e to edge f . Observe that no sequence of elementary swaps will take the label of edge e outside the connected component of G_S that contains e . Therefore e and f must lie in the same connected component of G_S . We will now show that each connected component of G_S is a clique. This implies that there is an elementary swap of e and f , and completes our proof.

Consider a simple path $(e_0, e_1), (e_1, e_2), \dots, (e_{k-1}, e_k)$ in G_S . Let σ_i , $i = 1, \dots, k$, be a flip sequence that realizes the elementary swap (e_{i-1}, e_i) , and let $\sigma = \sigma_1 \sigma_2 \dots \sigma_{k-1}$. Observe that σ takes the label of e_0 to e_{k-1} , and does not change the label of e_k (by the assumption that the path is simple). By definition of an elementary swap, the flip sequence σ_k has the form $\rho \pi \rho^{-1}$ where ρ is a sequence of flips that moves the labels of e_{k-1} and e_k into an empty convex pentagon, and π is the sequence of five flips that exchanges the labels of e_{k-1} and e_k .

Consider the flip sequence $\sigma \sigma_k \sigma^{-1} = \sigma \rho \pi \rho^{-1} \sigma^{-1} = \sigma \rho \pi (\sigma \rho)^{-1}$. The first part of this flip sequence, $\sigma \rho$, moves the labels of e_0 and e_k into an empty convex pentagon; the middle part, π , exchanges them; and the final part, $(\sigma \rho)^{-1}$ reverses the first part. Therefore this flip sequence realizes an elementary swap of e_0 and e_k . $\square$

5 Conclusions and Open Problems

We have characterized when two labelled triangulations of a set of n points belong to the same connected component of the labelled flip graph, and proved that the diameter of each connected component is bounded by $O(n^7)$. We conclude with some open problems:

1. Reduce the gap between the upper bound, $O(n^7)$, and the best known lower bound of $O(n^3)$ [7] on the diameter of a component of the labelled flip graph.
2. We have studied the case where each edge in a triangulation has a unique label, and given a bound of $O(n^7)$ on the diameter of a component of the labelled flip graph.

The case where edges are unlabelled can be viewed as the case where every edge has the same label—in this case the bound becomes $O(n^2)$. A unifying scenario is when the edges have labels and labels may appear on more than one edge. Is there a bound on the diameter of connected components of the flip graph that depends on the number of labels, or on the maximum number of edges with the same label?

3. We did not analyze the run-time of our algorithms in the main text. A crude bound is $O(n^8)$, with the bottleneck being the explicit construction in the proof of Lemma 2.1 of the double quadrilateral graph which has $O(n^4)$ vertices and thus $O(n^8)$ edges. This bound can surely be improved.
4. What is the complexity of the following flip distance problem for labelled triangulations: Given two labelled triangulations and a number k , is there a flip sequence of length at most k to transform the first triangulation to the second one? This problem is NP-complete in the unlabelled setting, but knowing the mapping of edges might make the problem easier.

Acknowledgements Open access funding provided by Institute of Science and Technology (IST Austria). We are grateful to F. Santos for bringing the result of Orden and Santos [27] to our attention. This research was initiated at the 2016 Bellairs Workshop on Geometry and Graphs. We thank anonymous reviewers and participants of the 2017 Symposium on Computational Geometry for helpful suggestions.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Aichholzer, O., Mulzer, W., Pilz, A.: Flip distance between triangulations of a simple polygon is NP-complete. *Discrete Comput. Geom.* **54**(2), 368–389 (2015)
2. Araujo-Pardo, G., Hubard, I., Oliveros, D., Schulte, E.: Colorful associahedra and cyclohedra. *J. Comb. Theory Ser. A* **129**, 122–141 (2015)
3. Bern, M., Eppstein, D.: Mesh generation and optimal triangulation. In: Du, D.Z., Hwang, F. (eds.) *Computing in Euclidean Geometry. Lecture Notes Series on Computing*, vol. 1, pp. 23–90. World Scientific, River Edge (1992)
4. Bing, R.H.: Some aspects of the topology of 3-manifolds related to the Poincaré conjecture. In: Saaty, T.L. (ed.) *Lectures on Modern Mathematics*, vol. II, pp. 93–128. Wiley, New York (1964)
5. Björner, A., Las Vergnas, M., Sturmfels, B., White, N., Ziegler, G.M.: *Oriented Matroids. Encyclopedia of Mathematics and Its Applications*, vol. 46, 2nd edn. Cambridge University Press, Cambridge (1999)
6. Bose, P., Hurtado, F.: Flips in planar graphs. *Comput. Geom.* **42**(1), 60–80 (2009)
7. Bose, P., Lubiw, A., Pathak, V., Verdonchot, S.: Flipping edge-labelled triangulations. *Comput. Geom.* **68**, 309–326 (2013)
8. Bose, P., Verdonchot, S.: Flips in edge-labelled pseudo-triangulations. *Comput. Geom.* **60**, 45–54 (2017)
9. Bryant, J.L.: Piecewise linear topology. In: Daverman, R., Sher, R. (eds.) *Handbook of Geometric Topology*, pp. 219–259. North-Holland, Amsterdam (2001)
10. Cano, J., Díaz-Báñez, J.M., Huemer, C., Urrutia, J.: The edge rotation graph. *Graphs Comb.* **29**(5), 1207–1219 (2013)
11. Danaraj, G., Klee, V.: Shellings of spheres and polytopes. *Duke Math. J.* **41**(2), 443–451 (1974)
12. Devadoss, S.L., O'Rourke, J.: *Discrete and Computational Geometry*. Princeton University Press, Princeton (2011)
13. Dyn, N., Goren, I., Rippa, S.: Transforming triangulations in polygonal domains. *Comput. Aided Geom. Design* **10**(6), 531–536 (1993)

14. Edelsbrunner, H.: *Geometry and Topology for Mesh Generation*. Cambridge University Press, Cambridge (2001)
15. Eppstein, D.: Happy endings for flip graphs. *J. Comput. Geom.* **1**(1), 3–28 (2010)
16. Hudson, J.F.P.: *Piecewise Linear Topology*. W.A. Benjamin, New York (1969)
17. Hurtado, F., Noy, M., Urrutia, J.: Flipping edges in triangulations. *Discrete Comput. Geom.* **22**(3), 333–346 (1999)
18. Ito, T., Demaine, E.D., Harvey, N.J.A., Papadimitriou, C.H., Sideri, M., Uehara, R., Uno, Y.: On the complexity of reconfiguration problems. *Theor. Comput. Sci.* **412**(12–14), 1054–1065 (2011)
19. Johnson, J.: Notes on piecewise-linear topology. <http://www.jessejohnson.me/notes.pdf> (2018). Accessed 20 July 2018
20. Kanj, I., Sedgwick, E., Xia, G.: Computing the flip distance between triangulations. *Discrete Comput. Geom.* **58**(2), 313–344 (2017)
21. Lawson, C.L.: Transforming triangulations. *Discrete Math.* **3**(4), 365–372 (1972)
22. Lawson, C.L.: Software for C^1 surface interpolation. *Mathematical Software III*. Publication of the Mathematics Research Center, vol. 39, pp. 161–194. Academic Press, New York (1977)
23. Lubiw, A., Pathak, V.: Flip distance between two triangulations of a point set is NP-complete. *Comput. Geom.* **49**, 17–23 (2015)
24. Lubiw, A., Pathak, V.: Reconfiguring ordered bases of a matroid (2016). [arXiv:1612.00958](https://arxiv.org/abs/1612.00958)
25. Molloy, M., Reed, B., Steiger, W.: On the mixing rate of the triangulation walk. In: Pardalos, P., Rajasekaran, S., Rolim, J. (eds.) *Randomization Methods in Algorithm Design*. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 43, pp. 179–190. American Mathematical Society, Providence (1999)
26. Munkres, J.R.: *Elements of Algebraic Topology*. Addison-Wesley, Menlo Park (1984)
27. Orden, D., Santos, F.: The polytope of non-crossing graphs on a planar point set. *Discrete Comput. Geom.* **33**(2), 275–305 (2005)
28. Pilz, A.: Flip distance between triangulations of a planar point set is APX-hard. *Comput. Geom.* **47**(5), 589–604 (2014)
29. Pournin, L.: The diameter of associahedra. *Adv. Math.* **259**, 13–42 (2014)
30. Seifert, H., Threlfall, W.: *A Textbook of Topology*. Pure and Applied Mathematics, vol. 89. Academic Press, New York (1980)
31. Sleator, D.D., Tarjan, R.E., Thurston, W.P.: Rotation distance, triangulations, and hyperbolic geometry. *J. Am. Math. Soc.* **1**(3), 647–681 (1988)
32. Stillwell, J.: *Classical Topology and Combinatorial Group Theory*. Graduate Texts in Mathematics, vol. 72, 2nd edn. Springer, New York (1993)
33. van den Heuvel, J.: The complexity of change. In: Blackburn, S.R., Gerke, S., Wildon, M. (eds.) *Surveys in Combinatorics*. London Mathematical Society Lecture Note Series, vol. 409, pp. 127–160. Cambridge University Press, Cambridge (2013)

Affiliations

Anna Lubiw¹ · Zuzana Masárová² · Uli Wagner²

Anna Lubiw
alubiw@uwaterloo.ca

Zuzana Masárová
zuzana.masarova@ist.ac.at

Uli Wagner
uli@ist.ac.at

¹ School of Computer Science, University of Waterloo, Waterloo, ON N2L 3G1, Canada

² IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria