



HAL
open science

Practical Reduction of Edge Flip Sequences in Two-Dimensional Triangulations

Jérémy Espinas, Raphaëlle Chaine, Pierre-Marie Gandoin

► **To cite this version:**

Jérémy Espinas, Raphaëlle Chaine, Pierre-Marie Gandoin. Practical Reduction of Edge Flip Sequences in Two-Dimensional Triangulations. 2013. hal-00871087

HAL Id: hal-00871087

<https://hal.science/hal-00871087v1>

Submitted on 9 Oct 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Practical Reduction of Edge Flip Sequences in Two-Dimensional Triangulations

Jérémy Espinas

Université de Lyon, CNRS
Université Lyon 1, LIRIS
France
jeremy.espinas@liris.cnrs.fr

Raphaëlle Chaine

Université de Lyon, CNRS
Université Lyon 1, LIRIS
France
raphaelle.chaine@liris.cnrs.fr

Pierre-Marie Gandoin

Université de Lyon, CNRS
Université Lyon 2, LIRIS
France
pierre-marie.gandoin@liris.cnrs.fr

Abstract

The development of laser scanning techniques has popularized the representation of 3D shapes by triangular meshes with a large number of vertices. Compression techniques dedicated to such meshes have emerged, which exploit the idea that the connectivity of a dense mesh does not deviate much from the connectivity that can be constructed automatically from the vertex positions (while possibly being guided by additional codes). The edge flip is one of the tools that can encode the differences between two meshes, and it is important to control the length of a sequence of flips that transform one triangulation into another. This paper provides a practical solution to this problem. Indeed, the problem of determining a minimal sequence of edge flips between two triangulations is NP-complete for some types of triangulations including manifold triangulations of surfaces, so that it is necessary to develop heuristics. Moreover, it is sometimes difficult to establish a first sequence of flips between two meshes, and we advocate a solution based on the reduction of an existing sequence. The new approach we propose is founded on the assignment of labels to identify the edges, with a property of label transfer during a flip. This gives a meaning to the tracking of an edge in a sequence of flips and offers the exploitation of very simple combinatorial properties. All the operations are performed directly on the sequence of labels denoting the edges to be flipped, almost regardless of the underlying surface, since only local temporary connectivity is involved:

- Given an initial triangulation T , and a sequence of edge flips ϕ , we introduce three simple moves to generate new sequences γ such that ϕ and γ transform T into a similar triangulation, up to a permutation of its labels. For the case where T is a triangulation of a convex n -gon, we prove that these three moves could be used to turn ϕ into any other sequence γ , such that ϕ and γ transform T into a similar triangulation, up to a permutation of its labels.
- On the basis of these three moves, we propose an efficient algorithm to reduce the input sequence of flips ϕ , without ensuring the minimization of the sequence length. This algorithm is polynomial and can be used in a purely combinatorial setting as well as in a geometric setting. The effectiveness of our approach is illustrated by practical benchmarks and comments.

1 Introduction

In this paper, we consider triangulations of dimension 2, possibly with boundaries, in which each face is bounded by a cycle of 3 edges and each edge is incident to one or two faces. Besides, the vertices are assumed to be labelled. If an edge (bc) is incident to faces (abc) and (dcb) , the flip of edge (bc) removes (abc) and (dcb) and constructs (abd) and (adc) which are incident to the created edge (ad) . There is no flip of an edge belonging to the boundary of a triangulation. The edge flipping is an operation widely used in computational geometry, especially to enumerate triangulations or to transform a mesh so that it satisfies certain properties.

The edge flip operation cannot be applied to any edge and some flips will be prohibited to maintain the structural properties of the triangulation. These structural properties may be purely combinatorial or they can also include geometric fold over free predicates determined by a particular embedding [1]. For example, this is the case if we restrict ourselves to planar straight-edge triangulations where each vertex has a geometric location. There is a specific case where all the edges can be flipped even if one considers geometric intersection-free properties: the triangulations of a convex n -gon.

Given two triangulations of the same domain with a common set of vertices, algorithms have been derived to generate a sequence of flips that transforms the first triangulation into the second. There are many ways to construct this sequence when dealing with a convex n -gon [2]. As far as planar straight-edge triangulations are concerned, since a sequence is reversible, it is possible to find a solution using Delaunay triangulation as a pivot between two triangulations, using the Lawson's algorithm [3]. In the combinatorial case, if we restrict ourselves to triangulations of genus 0 surfaces and impose that any two faces share at most one edge, the canonical form proposed by Wagner [4] can be used in turn. Espinas's *et al.* [5] work within a different class of triangulations, where any two faces can share the same triplet of vertices, and their algorithm produces a sequence of flips between two triangulations sharing the same topological genus without using any pivots.

Finding a minimal sequence of flips is more complicated. The only known method consists in building the *flip graph* of a given set of vertices S , i.e. the graph whose vertices are the possible triangulations of S and where two vertices are connected if a single flip exists to transform the first triangulation into the second. Minimal flip sequences between two triangulations are the shortest paths that connect the two triangulations in the flip graph. However, it is difficult to compute these sequences since the number of triangulations grows exponentially with the number of vertices, and the flip graph quickly becomes impossible to build. However, it is known that in the case of a convex n -gon, the flip graph is Hamiltonian [6, 7]. There are also a number of studies giving approximations of the flip distance, i.e. the number of flips in a minimal sequence. In the triangulations of a convex n -gon, the question of whether the determination of the flip distance in polynomial-time is still an open problem [8]. However, Pournin [9] showed that for a convex n -gon with $n \geq 13$, $2n - 10$ flips are sufficient and sometimes necessary to transform one triangulation into another. Regarding more general planar straight-edge triangulations, computing the flip distance is an NP-complete problem [10, 11], as is the case for simple polygons [12]. However, Lawson [3] showed that for any two planar triangulations consisting of n vertices, $O(n^2)$ flips are sufficient, and Hurtado *et al.* [13] proved that this is a lower bound for the number of flips. Hanke *et al.* [14] showed that the flip distance between two planar triangulations is at most the number of edge intersections obtained by superimposing the two triangulations. If the point set does not have five vertices forming an empty pentagon, Eppstein [15] showed that the flip distance can be computed in time $O(n^2)$. For the general case, he gave a polynomial-time algorithm for computing a lower bound. It can be observed that the more configurations are authorized for the flips, the better bounds are. In the combinatorial case, determining the flip distance between two triangulations of a genus 0 surface whose flips are conditioned by the criteria proposed by Wagner [4], the best upper bound for now is $5.3n - 24.4$ edge flips [16]. Regarding the lower limit, pairs of triangulations exist that require $2n - 15$ edge flips [17]. In this paper, we work on triangulations whose edges are *labelled*. This idea has already been used for non-crossing spanning trees to which an *edge move* operation is applied. This is a generalization to trees of the edge flip and it is used to modify a tree or to swap its edge labels [18–20].

In this paper, we will illustrate that the labelling of edges is also advantageous when dealing with triangulations. After introducing our framework with a number of definitions and properties, we underline that some basic flip sequences correspond to transpositions of labels. We also introduce the notions of weak and strong equivalence between the flip sequences that turn a given triangulation into a target one, depending on whether or not the resulting labels are the same (section 2). We then introduce three moves to generate flip sequences that are weakly equivalent to a given sequence. In the case of triangulations of a convex n -gon, these three moves correspond to a generating set of the equivalence classes (section 3). These results are then used to develop an algorithm for the reduction of flip sequences which imposes no edge labelling on the initial or resulting triangulations, is adaptable to all types of triangulations and works in polynomial time (section 4). The behaviour of this algorithm is illustrated in a practical situation.

2 Sequences of edge flips in triangulations with labelled edges

2.1 Triangulations with labelled edges

In this work, we exploit the advantages of identifying each edge of a triangulation with a label, rather than using its endpoints, as generally done. Each edge is given a different label, which will persist throughout the edge flips, except for boundary edges that cannot be flipped. Of course, there are other cases where a labelled edge cannot be flipped, because the very definition of the flip operation is structurally conditioned by a combinatorial or a geometric intersection-free criterion. Actually, we construct the set of triangulations with labelled edges \mathcal{T}_{lab} on top of a class of triangulations \mathcal{T} with an edge flip operation defined on it, and the accompanying conditions required for flipping an edge. In the following, we will only impose this class of triangulations to guarantee that every face is incident to three different edges and that each edge is incident to different vertices. Then we can choose to use purely combinatorial triangulations or remain within planar triangulations, manifold triangulations of sampled surfaces or even triangulations of the convex n -gon.

2.2 Flip of a labelled edge: definitions and properties

Let $T \in \mathcal{T}_{lab}$ be a triangulation with a flippable edge i between two vertices a and b . i is incident to the faces (abc) and (adb) . The edge flip of i , denoted by \mathcal{F}_i , suppresses the edge (ab) and its label i is transferred onto the new created edge (dc) . The *support* of the flip \mathcal{F}_i on T , denoted as $supp(i, T)$, consists of the two faces (abc) and (adb) incident to i in T . $\mathcal{F}_i(T)$ denotes the triangulation that would result from the edge flip i on T .

Involution: If i is a flippable edge in T , then i is also flippable in $\mathcal{F}_i(T)$ and $\mathcal{F}_i(\mathcal{F}_i(T)) = T$, which means that \mathcal{F}_i is an involution.

Let $T \in \mathcal{T}_{lab}$ be a triangulation with labelled edges $1, \dots, n$. We define the composition of two edge flips as $(\mathcal{F}_{i_2} \circ \mathcal{F}_{i_1})(T) = \mathcal{F}_{i_2}(\mathcal{F}_{i_1}(T))$, which means that it can be performed on T if i_1 is flippable in T and i_2 is flippable in $\mathcal{F}_{i_1}(T)$. The *sequence of flips* of the edges i_1, \dots, i_f of T corresponds to $\phi = (\mathcal{F}_{i_n} \circ \mathcal{F}_{i_{n-1}} \circ \dots \circ \mathcal{F}_{i_1})$.

Restricted commutativity: Let i and j be two flippable edges of T , such that no face of T is incident to both i and j (we say that the cardinality of $supp(i, T) \cap supp(j, T)$ is null in terms of the number of common faces and note $|supp(i, T) \cap supp(j, T)| = 0$). We can observe that $\mathcal{F}_i \circ \mathcal{F}_j$ can be performed on T and that $(\mathcal{F}_i \circ \mathcal{F}_j)(T) = (\mathcal{F}_j \circ \mathcal{F}_i)(T)$ and we say that \mathcal{F}_i and \mathcal{F}_j commute on T .

From this, we can deduce rules of commutativity between two consecutive edge flips within a sequence: let ϕ , μ and γ be three sequences of flips such that $\phi(T) = (\mu \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \gamma)(T)$. We have $\phi(T) = (\mu \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \gamma)(T)$ if $|supp(i, \gamma(T)) \cap supp(j, \gamma(T))| = 0$.

We now introduce $\mathcal{P}_{(i,j)}$ as an operation of *label transposition* between two edges i and j of a triangulation T . When being composed within a sequence, edge flips and transpositions can commute in the following way:

$$(\mathcal{P}_{(i,j)} \circ F_k)(T) = (F_{\theta_{(i,j)}(k)} \circ \mathcal{P}_{(i,j)})(T)$$

where $\theta_{(i,j)}(k)$ corresponds to j if $k = i$, i if $k = j$ and k otherwise.

In fact, the main reason for introducing the transposition of labels is the relationship between this operation and specific sequences of edge flips.

Proposition 1. Let T_5 be a triangulation of a 5-gon in which any possible configuration of an inner edge is flippable, and let i and j be the two inner edges of T_5 . We have:

$$\mathcal{P}_{(i,j)}(T_5) = (\mathcal{F}_i \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \mathcal{F}_i)(T_5) = (\mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j)(T_5)$$

Proof. Figure 1 illustrates that $(\mathcal{F}_i \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \mathcal{F}_i)$ amounts to the transposition of the labels i and j when performed on a triangulation of the 5-gon. This proves our result since the other triangulations of the 5-gon are similar to this one up to rotation. In particular, this result is true when dealing with a convex n -gon. \square




Figure 1: Illustration that $\mathcal{P}_{(i,j)}(T_5) = (\mathcal{F}_i \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \mathcal{F}_i)(T_5)$ on the 5-gon.

More generally, this result illustrates that the flip can be used to change the connectivity of a triangulation, but also to move its edges. This remark is the starting point of a process to generate sequences of flips that are equivalent when performed on a triangulation, and to develop an algorithm to reduce sequences.

2.3 Equivalence between sequences of edge flips

We first introduce two equivalence relations between sequences of edge flips performed on a labelled triangulation T .

Let ϕ_1 and ϕ_2 be two sequences of edge flips. $\phi_1(T)$ and $\phi_2(T)$ are two triangulations of the same set of vertices.

- ϕ_1 is *strongly equivalent* to ϕ_2 when performed on T , if $\phi_1(T) = \phi_2(T)$,
- ϕ_1 is *weakly equivalent* to ϕ_2 when performed on T , if $\phi_1(T)$ and $\phi_2(T)$ correspond to the same triangulation up to a permutation of the edge labels; we denote $\phi_1(T) \simeq \phi_2(T)$ in that case.

In the following, we are interested in generating sequences that are either strongly or weakly equivalent to a given sequence ϕ , when performed on T . Let us observe that the strong equivalence implies weak equivalence and that a minimal sequence of edge flips between unlabelled triangulations T and $\phi(T)$ is weakly equivalent to ϕ on T .

3 Generation of equivalent sequences of edge flips

3.1 Introduction of three moves

The three moves we introduce are directly derived from the properties of the flip operation performed on a triangulation with labelled edges. They can be used to generate sequences that are either strongly or weakly equivalent to an input sequence ϕ , when performed on an input triangulation $T \in \mathcal{T}_{lab}$. Furthermore, if T is a triangulation of a convex n -gon, we will show that one can generate any sequence that is weakly equivalent to ϕ on T using only these three moves. In the following, all the sequences we consider are composed of flippable edges, and the operations we introduce transform a sequence of flippable edges into another one where the flips remain valid.

1st move (exploiting commutativity): If there exists two sequences μ and ν such that $\phi = (\mu \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \nu)$ and that $|supp(i, \nu(T)) \cap supp(j, \nu(T))| = 0$, then the sequence $(\mu \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \nu)$ is strongly equivalent to ϕ on T . Therefore, if the sequence ϕ applied to T contains flips that can be commuted, this move can be used to generate new sequences that are strongly equivalent to ϕ . In the following, if two sequences can be transformed one into another, using only this first move, we say that they are *strongly equivalent by commutativity*.

2nd move (exploiting involution):

We have the property that $(\mathcal{F}_i \circ \mathcal{F}_i)(T) = T$ for any flippable edge i in T . Therefore, one can generate strongly equivalent sequences by inserting $(\mathcal{F}_i \circ \mathcal{F}_i)$ wherever i is flippable in the intermediate triangulations produced by ϕ (expanding version of the 2nd move). Conversely, if ϕ contains several occurrences of \mathcal{F}_i , and that there exists two sequences μ and ν such that ϕ is strongly equivalent by commutativity to the sequence $(\mu \circ \mathcal{F}_i \circ \mathcal{F}_i \circ \nu)$ when performed on T , then the first two moves can be used to generate the sequence $(\mu \circ \nu)$ which is strongly equivalent to ϕ but contains two fewer edge flips (reducing version of the 2nd move).

3rd move (exploiting transposition of labels):

Proposition 2. If there exists μ and ν such that $\phi = (\mu \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \nu)$ with $|supp(i, \nu(T)) \cap supp(j, \nu(T))| = 1$ and that i and j are also flippable in $\nu(T)$ and $(\mathcal{F}_i \circ \nu)(T)$ respectively, then $(\mu \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \nu)(T) = (\mu \circ P_{(i,j)} \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \nu)(T)$.

Proof. Since $|supp(i, \nu(T)) \cap supp(j, \nu(T))| = 1$ and we can flip i in $\nu(T)$ and j in $(\mathcal{F}_i \circ \nu)(T)$, the local configuration of the two supports is a 5-gon where it is possible to flip these two inner edges successively. Let us consider $T' = (\mathcal{F}_j \circ \mathcal{F}_i \circ \nu)(T)$. Following the same reasoning as for Proposition 1, we have $(\mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j)(T') = P_{(i,j)}(T')$. We now replace T' with $(\mathcal{F}_j \circ \mathcal{F}_i \circ \nu)(T)$ in this equality, which amounts to $(\mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \nu)(T) = (P_{(i,j)} \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \nu)(T)$ i.e. $(\mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \nu)(T) = (P_{(i,j)} \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \nu)(T)$ (see Figure 2). \square

Let $P_{(i,j)}(\mu)$ denote the sequence μ in which every occurrence of \mathcal{F}_i has been replaced with \mathcal{F}_j and conversely.

The above property implies that $\phi = (\mu \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \nu)$ is weakly equivalent to $(P_{(i,j)}(\mu) \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \nu)$ when performed on T (reducing version of the 3rd move corresponding to one flip fewer).

Conversely, if there exists μ and ν such that $\phi = (\mu \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \nu)$ with $|supp(i, \nu(T)) \cap supp(j, \nu(T))| = 1$, that j is flippable in $\nu(T)$ and $(\mathcal{F}_i \circ \mathcal{F}_j \circ \nu)(T)$, and i is flippable in $(\mathcal{F}_j \circ \nu)(T)$, then $\phi = (\mu \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \nu)$ is weakly equivalent to $(P_{(i,j)}(\mu) \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \nu)$ when performed on T (expanding version of the 3rd move corresponding to one more flip).




Figure 2: In the triangulated pentagon, $(\mathcal{F}_i \circ \mathcal{F}_j \circ \mathcal{F}_i)(T) = (\mathcal{P}_{(i,j)} \circ \mathcal{F}_i \circ \mathcal{F}_j)(T)$

3.2 Particular case of the convex n -gon

When dealing with triangulations of a convex n -gon, any inner edge is flippable, whatever its support. Furthermore, we have the following result, regarding the equivalent sequences that can be generated with the three moves introduced in the previous section:

Theorem 1. Let $T \in \mathcal{T}_{lab}$ be a labelled triangulation of a convex n -gon and ϕ_1 a sequence of edge flips. The three moves exploiting the commutativity and the involution of edge flips, in combination with the transposition of labels, can be used to generate any sequence of edge flips that is weakly equivalent to ϕ_1 on T . This means that the three moves are a generating set of the equivalence class associated with ϕ_1 .

The proof of Theorem 3.2 can be found in the appendix. We believe that it constitutes a step forward in the understanding of edge flips, but the problem of exhibiting the minimal sequences remains complicated. Let us consider the example presented in Figure 3 that illustrates how the transformation of the initial sequence in a) into the minimal sequence in b) cannot easily be obtained by a direct algorithm :

$$\begin{aligned}
 (\mathcal{F}_4 \circ \mathcal{F}_3 \circ \mathcal{F}_1 \circ \mathcal{F}_2 \circ \mathcal{F}_3 \circ \mathcal{F}_4)(T) &= (P_{(1,2)} \circ \mathcal{F}_4 \circ \mathcal{F}_3 \circ \mathcal{F}_1 \circ \mathcal{F}_2 \circ \mathcal{F}_1 \circ \mathcal{F}_3 \circ \mathcal{F}_4)(T) \\
 &= (P_{(1,2)} \circ \mathcal{F}_4 \circ \mathcal{F}_1 \circ \mathcal{F}_3 \circ \mathcal{F}_2 \circ \mathcal{F}_3 \circ \mathcal{F}_1 \circ \mathcal{F}_4)(T) \\
 &= (P_{(1,2)} \circ P_{(2,3)} \circ \mathcal{F}_4 \circ \mathcal{F}_1 \circ \mathcal{F}_3 \circ \mathcal{F}_2 \circ \mathcal{F}_1 \circ \mathcal{F}_4)(T) \\
 &= (P_{(1,2)} \circ P_{(2,3)} \circ \mathcal{F}_1 \circ \mathcal{F}_4 \circ \mathcal{F}_3 \circ \mathcal{F}_4 \circ \mathcal{F}_2 \circ \mathcal{F}_1)(T) \\
 &= (P_{(1,2)} \circ P_{(2,3)} \circ P_{(3,4)} \circ \mathcal{F}_1 \circ \mathcal{F}_4 \circ \mathcal{F}_3 \circ \mathcal{F}_2 \circ \mathcal{F}_1)(T) \\
 &\simeq (\mathcal{F}_1 \circ \mathcal{F}_4 \circ \mathcal{F}_3 \circ \mathcal{F}_2 \circ \mathcal{F}_1)(T)
 \end{aligned}$$




Figure 3: a) and b) correspond to two weakly equivalent sequences.

3.3 General case

When working within a class of triangulations that is broader than the triangulations of the convex n -gon (possibly with internal vertices or even with concavities on boundaries, if in a geometric setting), the previous result probably does not directly generalize and we cannot ensure that a minimal sequence of flips could be reached by using only the three moves in their expanding and reducing versions.

4 Reducing sequences of edge flips

Based on the previous results and on the three moves introduced in section 3.1, we have devised an efficient algorithm for determining a reduced sequence that is weakly equivalent to an input sequence. It provides a heuristic alternative to the usually NP-complete problem of generating a sequence of minimal

length between two triangulations. This algorithm illustrates the possibilities offered by our results. It can be used in a purely combinatorial setting or in combination with geometric intersection free predicates. The reduced sequence obtained is locally minimal with regards to the other sequences that could be reached using any single occurrence of the expanding or reducing moves 2 and 3, in combination with the first move, but it is not necessarily a globally shortest edge flip sequence.

4.1 Algorithm

Given a triangulation T and a sequence ϕ with f flips, the algorithm attempts to eliminate each of the edge flips, considered one after the other, into a constructive order.

4.1.1 Attempt to reduce ϕ by moving the flip \mathcal{F}_i within ϕ :

Given an occurrence of \mathcal{F}_i in ϕ , we investigate whether the sequence ϕ contains another occurrence of \mathcal{F}_i among the following flips. This requires a search among at most $n - 1$ elements, with a complexity in $O(n)$.

If there is no occurrence of \mathcal{F}_i among the following flips in ϕ , we say that \mathcal{F}_i cannot be removed and nothing is done. Otherwise, we denote \mathcal{F}'_i the next occurrence of \mathcal{F}_i in ϕ , and we try to bring the two edge flips together in ϕ : \mathcal{F}_i is moved to the left, towards \mathcal{F}'_i , as long as \mathcal{F}_i can be commuted with the next edge flip on the corresponding intermediate triangulation. Similarly, \mathcal{F}'_i is moved to the right, towards \mathcal{F}_i , as long as \mathcal{F}_i can be commuted with the previous edge flip on the corresponding intermediate triangulation. A move to the left denotes a move amongst the flips that were initially planned to be performed afterwards. Similarly, a move to the right is a move amongst the flips that were planned to be performed beforehand.

More formally, let ϕ_{ev} denote the evolving sequence of flips which is initialized to ϕ . At each step of the displacement of \mathcal{F}_i , let μ_1 , μ_2 and μ_3 denote the three sequences such that $\phi_{ev} = (\mu_3 \circ \mathcal{F}'_i \circ \mu_2 \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mu_1)$. If we have $|supp(i, \mu_1(T)) \cap supp(j, \mu_1(T))| = 0$ and the edges can be flipped, \mathcal{F}_i is moved to the left; ϕ_{ev} is updated to $(\mu_3 \circ \mathcal{F}'_i \circ \mu_2 \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \mu_1)$ and $\mu_1(T)$ is updated to $(\mathcal{F}_j \circ \mu_1)(T)$. The displacement is stopped otherwise. The check of the independence between the supports of i and j is performed on $\mu_1(T)$, which means that we update $\mu_1(T)$ throughout the displacement of \mathcal{F}_i . It should be noted that each single displacement on the left corresponds to a flip in $\mu_1(T)$, and thus is performed in constant time. The total cost of displacing \mathcal{F}_i on the left is $\Theta(l)$ with $l \leq f$.

Afterwards, we similarly displace \mathcal{F}'_i on the right towards \mathcal{F}_i as long as it is possible and the two occurrences of \mathcal{F}_i do not cross. Each single displacement on the right is checked and performed in constant time, except for the first displacement of \mathcal{F}'_i , which accounts for $\Theta(k)$, where k is the number of flips between the two occurrences of \mathcal{F}_i ($k \leq f - l$). Therefore, the total cost of displacing \mathcal{F}'_i on the right is $\Theta(k)$.

When the displacement of \mathcal{F}'_i and \mathcal{F}_i is over, three cases can be encountered:

- If $\phi_{ev} = (\nu \circ \mathcal{F}'_i \circ \mathcal{F}_i \circ \mu)$, then ϕ_{ev} is reduced to $(\nu \circ \mu)$ (two removal of \mathcal{F}_i , that is two edge flips fewer in constant time)
- If $\phi_{ev} = (\nu \circ \mathcal{F}'_i \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mu)$ and $|supp(i, \mu(T)) \cap supp(j, \mu(T))| = 1$ with each edge i and j flippable within the sequence $(\mathcal{F}_i \circ \mathcal{F}_j)$ performed on $\mu(T)$, then ϕ_{ev} is reduced to $(T_{(i,j)}(\nu) \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \mu)$ (one removal of \mathcal{F}_i , that is one flip fewer in $O(n)$, since the sequence ν which may be transformed contains fewer than f edge flips)
- If neither of the previous cases is encountered, we say that ϕ cannot directly be reduced using \mathcal{F}_i . The first occurrence of \mathcal{F}_i is reset to its initial position which is performed in constant time.

In the end, the cost of an attempt to reduce ϕ by moving the flip \mathcal{F}_i within ϕ is $O(f)$ where f is the number of edge flips in ϕ .

4.1.2 Algorithm for reducing sequences of edge flip:

The algorithm iterates over all the edge flips in ϕ , from the beginning to the end of the sequence (right to left), until it finds a flip \mathcal{F}_i whose displacement reduces ϕ . The removal of an occurrence of \mathcal{F}_i makes it necessary to reconsider a previously checked edge flip, in order to reduce the updated sequence. Therefore, after each reduction of ϕ , the iteration is restarted from the beginning of the updated ϕ .

After each reduction, the algorithm iterates over a subset of the $l \leq f$ remaining edges in ϕ_{ev} until it identifies the next reduction, if any. The complexity of finding the next reduction is $O(l^2)$. If no more edge can lead to a sequence reduction, the algorithm iterates over the $l \leq f$ edges of ϕ_{ev} to search for a last attempt at reduction with each remaining edge ($O(l^2)$ complexity). The sequence is strictly decreased in length after each reduction, so that there can be no more than f reductions. Therefore the overall complexity of the algorithm is $O(f^3)$. There are several variations around this basic version of the algorithm, but they do not affect the overall complexity of the algorithm.

It can be noted that when applied to a sequence that is nearly optimal, the algorithm performs in $O(f^2)$ or even better ($O(f)$) if most of the edge flips are present only once and if we trigger the displacement of a couple of flips only when identifying a second occurrence of a flip. The algorithm reduces the number of edge flips in a sequence and exhibits a local minimum in terms of the length of the sequence. If it is possible to reduce the sequence further, we would have to go through temporary sequences whose length is longer. Ideas in this direction can be suggested by Theorem 1, but their integration in an effective algorithm is not that simple.

	Convex n -gon			Geometrical Setting			Combinatorial Setting			
Nb of edges	2 000	10 000	30 000	2 000	10 000	30 000	2 000	10 000	30 000	100 000
Initial size ($r=1.1$)	200	500	2 000	200	500	2 000	200	500	2 000	10 000
Gain (%)	27 (13)	64 (12)	216 (10)	2 (1)	7 (1)	70 (3)	8 (4)	31 (6)	157 (7)	1 043 (10)
Time	1ms	13ms	76ms	1ms	4ms	31ms	5ms	24ms	541ms	2s
Initial size ($r=2$)	2 000	6 000	10 000	2 000	6 000	10 000	2 000	6 000	10 000	20 000
Gain (%)	973 (48)	2 950 (49)	4 964 (49)	596 (29)	1 062 (17)	2698 (26)	854 (42)	2 288 (38)	4 608 (46)	9 152 (45)
Time	276ms	7s	1min22s	101ms	588ms	1s	3s	57s	1m28s	5min36s
Initial size ($r=10$)	6 000	10 000	20 000	6 000	10 000	20 000	6 000	10 000	20 000	30 000
Gain (%)	5 277 (87)	8 666 (86)	17 759 (88)	4 212 (70)	5 686 (43)	10 353 (51)	4 617 (76)	5 804 (58)	15 060 (75)	22 530 (75)
Time	14s	1min15s	3min53s	281ms	946ms	3s	8s	5min52s	20min14s	1h32min

Table 1: Experimental results : in the geometrical and combinatorial settings, we have used the same flip sequences, ensuring that all edge flips verify the constraints of each setting. This helps emphasize the effect of the choice of a class on the reduced flip sequence. The last column corresponds to a sequence generated and reduced on a purely combinatorial case.

4.2 Experimental results

The experimental results in Table 1 were calculated on a PC with an Intel Core 2 Duo 2 GHz and 4 Gb of RAM. We worked within three different classes of triangulations: convex n -gon, straight-edge triangulations (geometric setting) of a sphere (but an underlying surface with a different genus could also be used) and a purely combinatorial setting where all the connectivity requirements prescribed in this paper are satisfied and each vertex is associated with a single topological disk (as in [5]). The sequences were randomly generated such that the triangulation remained valid after each flip and there is at least the flip of one neighbouring edge between two occurrences of the same edge flip. Flip sequences are

characterized by their length f as well as their redundancy r which is the ratio of the sequence length to the number of different edges involved. In particular, if the sequence has at most one occurrence of each flip, r is equal to 1. Also, $f(1 - 1/r)$ gives an estimate of the number of flips that could be removed by the simplification.

It is worthwhile noting that the observed complexity of the algorithm is directly related to the redundancy r : linear in the length of the original sequence for r around 1, it becomes quadratic or even cubic when r grows. In practice, we observe on the experimental results presented in 1 that the overall complexity approximation performs in $O(f^{3-2/r})$. Likewise, the gain essentially depends on the triangulation class and the redundancy of the initial sequence rather than its length. The resulting lengths are comparable to the theoretical bounds on flip distances. In particular, it is instructive to see that the same flip sequence generated in a geometric setting can be reduced further when being handled within a combinatorial setting, where we benefit from more possible intermediate triangulations. This reduces the number of flips removed but also the algorithm's running time.

To our knowledge, there are no other algorithms to reduce sequences of edge flips, and we cannot compare our results to others. However, we note that the reduced sequences applied on a convex n -gon, contains less than $2n - 10$ flips, i.e. the length is less than the upper bound introduced by Pournin [9].

Finally, we note that the progress of a triangulation on which we perform a sequence of flips is similar before and after the reduction (see Figure 4).




Figure 4: a) Progress of a sequence composed of 40 000 edge flips with redundancy $r = 1.5$ (combinatorial setting). b) The reduced sequence composed of 29 942 edge flips. For each sequence, the figure illustrates the initial triangulation, the triangulation where $1/3$ of the sequence has been applied, then $2/3$ of the sequence, and the final triangulation.

5 Conclusion

We have presented an efficient algorithm for reducing flip sequences based on the edge move study. We note that the gain is very close to the redundancy of the sequence, especially in the case of the n -gon, which demonstrates the practical value of the method for the efficient encoding of a triangulation connectivity. Indeed, several compression methods ([21, 22] for instance) are able to automatically recover much of the connectivity of a triangulation from its geometry information and could benefit, for the remaining connectivity, from a differential coding expressed as a simplified flip sequence. Another application could consist in using the labels to encode additional information. For example, we could assign colors to edges and transport them by successive permutations. We could also encode the transformation of a non-triangular graph by embedding it in a triangulation and marking the edges to be removed after the transformation.

References

- [1] Bose, P., Hurtado, F.: Flips in planar graphs. *Comput. Geom. Theory Appl.* **42** (January 2009) 60–80
- [2] De Loera, J., Rambau, J., Santos, F.: *Triangulations: Structures for Algorithms and Applications. Algorithms and Computation in Mathematics.* Springer (Novembre 2010)
- [3] Lawson, C.L. In: *Software for C^1 surface interpolation.* Academic Press (1977) 161–194
- [4] Wagner, K.: Bemerkungen zum vierfarbenproblem. *Jahresber. Deutsch. Math.-Verein.* **46** (1936) 26–32
- [5] Espinas, J., Chaine, R., Gandoin, P.M.: A simple yet efficient algorithm to turn one oriented triangular mesh connectivity into another. Technical report. hal-00798335 (2012)
- [6] Lucas, J.M.: The rotation graph of binary trees is hamiltonian. *Journal of Algorithms* **8**(4) (1987) 503 – 535
- [7] Hurtado, F.: Graph of triangulations of a convex polygon and tree of triangulations. *Computational Geometry* **13**(3) (1999) 179–188
- [8] Sleator, D., Tarjan, R., Thurston, W.: Rotation distance. In: *Open problems in communication and computation,* Springer-Verlag (1987)
- [9] Pournin, L.: The diameters of associahedra. arXiv preprint arXiv:1207.6296 (2012)
- [10] Lubiw, A., Pathak, V.: Flip distance between two triangulations of a point-set is np-complete. *CoRR* **abs/1205.2425** (2012)
- [11] Pilz, A.: Flip distance between triangulations of a planar point set is np-complete. arXiv preprint arXiv:1206.3179 (2012)
- [12] Aichholzer, O., Mulzer, W., Pilz, A.: Flip distance between triangulations of a simple polygon is np-complete. *CoRR* **abs/1209.0579** (2012)
- [13] Hurtado, F., Noy, M., Urrutia, J.: Flipping edges in triangulations. In: *Proceedings of the twelfth annual symposium on Computational geometry. SCG '96, New York, NY, USA, ACM* (1996) 214–223
- [14] Hanke, S., Ottmann, T., Schuierer, S.: The edge-flipping distance of triangulations. *Journal of Universal Computer Science* **2**(8) (1996) 570–579
- [15] Eppstein, D.: Happy endings for flip graphs. *Journal of Computational Geometry* **1** (2010) 3–28
- [16] Bose, P., Jansens, D., van Renssen, A., Saumell, M., Verdonchot, S.: Making triangulations 4-connected using flips. *CoRR* **abs/1110.6473** (2011)
- [17] Komuro, H.: The diagonal flips of triangulations on the sphere. *The Yokohama Mathematical Journal* **44**(2) (1997) 115–122
- [18] Holzmann, C., Harary, F.: On the tree graph of a matroid. *SIAM Journal On Applied Mathematics* **22** (1972) 187–193
- [19] Liu, G.: On connectivities of tree graphs. *Journal of Graph Theory* **12**(3) (1988) 453–459
- [20] Hernando, C., Hurtado, F., Marquez, A.: Grafos de árboles etiquetados y grafos de árboles geométricos etiquetados. In *Proc. X Encuentros de Geometría Computacional* (2003) 13–19

- [21] Devillers, O., Gandoïn, P.M.: Geometric compression for interactive transmission. In: Visualization 2000. Proceedings, IEEE (2000) 319–326
- [22] Valette, S., Chaine, R., Prost, R.: Progressive lossless mesh compression via incremental parametric refinement. In: Proceedings of the Symposium on Geometry Processing. SGP '09, Aire-la-Ville, Switzerland, Eurographics Association (2009) 1301–1310

6 Appendix : proof of the theorem

Theorem 1. Let $T \in \mathcal{T}_{lab}$, a labelled triangulation of a convex n -gon and ϕ_1 a sequence of edge flips. The three moves exploiting the commutativity and the involution of edge flips, in combination with the transpositions of labels, can be used to generate any sequences of edge flips that is weakly equivalent to ϕ_1 on T . This means that the three moves are a generating set of the equivalence class associated to ϕ_1 .

There may be other ways of proving this theorem, but we have chosen to demonstrate it with labelled edge flips, while introducing new results on triangulations with labelled edges.

6.0.1 Preliminary definitions

Let ϕ_1 and ϕ_2 denote two sequences of edge flips. For the needs of the demonstration of Theorem 1, we say that ϕ_2 is a *direct reduction* of ϕ_1 if edge i exists (and possibly an edge j) and two sequences μ and ν that satisfy one of the two following conditions:

- ϕ_1 is strongly equivalent by commutativity to $(\nu \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mathcal{F}_j \circ \mu)$ with $|supp(i, T) \cap supp(j, T)| = 1$ and $\phi_2 = (T_{(i,j)}(\nu) \circ \mathcal{F}_j \circ \mathcal{F}_i \circ \mu)(T)$ (weak equivalence of ϕ_1 and ϕ_2 on T).
- ϕ_1 is strongly equivalent by commutativity to $(\nu \circ \mathcal{F}_i \circ \mathcal{F}_i \circ \mu)$ and $\phi_2 = (\nu \circ \mu)$ (strong equivalence of ϕ_1 and ϕ_2 on T).

We say that a sequence is *reduced* if there is no direct reduction of it. By iterating direct reductions on an input sequence ϕ , we can construct a *monotonically reduced* version of it. There is weak equivalence between ϕ and its *monotonic reductions*. It is noted that monotonically reduced versions of an input sequence are not unique and that monotonically reduced sequences are not necessarily minimal (see Figure 3 where a) corresponds to a non-minimal reduced sequence, since the sequence in b) contains one flip fewer).

In order to demonstrate theorem 1, we need the following lemmas:

Lemma 1. Let $T \in \mathcal{T}_{lab}$ be a triangulation of a convex n -gon. If a sequence of flips ϕ contains no double occurrence of the same edge flip, then the sequence is minimal. Any other minimal sequence that is weakly equivalent to ϕ on T is composed of the same edge flips (up to their order in the sequence).

Proof. The result is clear when ϕ is empty. We now consider the case where the length of ϕ is not zero, and there is no double occurrence of the same edge flip in ϕ . If ϕ is not minimal, there exists a sequence ϕ_{min} that is weakly equivalent to ϕ on T but whose length is strictly lower. Let \mathcal{F}_i be a flip that is present in ϕ but not in ϕ_{min} (\mathcal{F}_i necessarily exists due to the length difference). Let v_1 and v_2 be the two vertices incident to i in T (but also in $\phi_{min}(T)$). Let v_3 and v_4 be the two vertices incident to i in $\phi(T)$ (they are the two opposite vertices incident to the support of i just before its flip). In the initial triangulation T , the edge i splits the n -gon into two disconnected parts. They contain the vertices v_3 and v_4 , respectively. This means that v_3 and v_4 are separated by i and they cannot be connected without flipping i . However, v_3 and v_4 should be connected in $\phi_{min}(T)$, which contradicts the very existence of ϕ_{min} since i is not flipped in ϕ_{min} . Therefore ϕ is minimal.

Let us now assume that ϕ' is another minimal sequence such that $\phi'(T) \simeq \phi(T)$, but ϕ' does not flip the same edges as ϕ . Let \mathcal{F}_i be a flip present in ϕ but not in ϕ' . With the same reasoning as above, we deduce that ϕ and ϕ' cannot be equivalent on T , which is a contradiction. \square

We note that the Lemma 6.0.1 is not valid when there are internal vertices. Figure 5 illustrates this point, with a counter-example of a non-minimal sequence without any redundancy in the edges that are flipped.

Lemma 2. Let $T \in \mathcal{T}_{lab}$ be a triangulation of a convex n -gon. If a minimal sequence of flips ϕ_1 contains no double occurrence of the same edge flip, then any minimal sequence ϕ_2 weakly equivalent to ϕ_1 on T is strongly equivalent to ϕ_1 by commutativity.




Figure 5: The left and right triangulations are composed of the same edges and oriented faces but the labeling of the edges is different. Therefore, the sequence $(\mathcal{F}_4 \circ \mathcal{F}_3 \circ \mathcal{F}_2 \circ \mathcal{F}_1)$ is weakly equivalent to the identity when performed on the left triangulation. This sequence is not minimal, but it contains no double occurrence of a same edge flip.

Proof. Lemma 1 ensures that ϕ_1 and ϕ_2 are composed of the same edge flips. Let us seek to reorder, by commutativity, the flips of the sequence ϕ_2 according to their position in ϕ_1 . If the first edge to be flipped in ϕ_1 is \mathcal{F}_i , then we move \mathcal{F}_i within ϕ_2 towards the first position, as long as \mathcal{F}_i can commute with the edge flip that precedes. If \mathcal{F}_i can be moved to the first place, we repeat the process by moving the edge flip to be put in the position $p = 2$, and so on. This operation is continued until one meets an edge flip \mathcal{F}_j that cannot be moved to the desired position p by commutativity. Therefore, there exist two sequences μ_1 and ν with $|\nu| = p - 1$, such that $\phi_1(T) = (\mu_1 \circ \mathcal{F}_j \circ \nu)(T)$, and there exist μ_2 and μ_3 such that $\phi_2(T) = (\mu_3 \circ \mathcal{F}_j \circ \mathcal{F}_k \circ \mu_2 \circ \nu)(T)$ and \mathcal{F}_j and \mathcal{F}_k do not commute on $(\mu_2 \circ \nu)(T)$. This means that j and k are incident to a common face (see Figure 6 a). Let v_1 and v_2 be the two vertices incident to k in the triangulation $(\mu_2 \circ \nu)(T)$. k is incident to the faces $(v_1 v_2 v_3)$ and $(v_2 v_1 v_4)$. Without loss of generality, we can assume that the edge j is incident to the vertices v_1 and v_3 and to the faces $(v_3 v_1 v_2)$ and $(v_1 v_3 v_5)$.

Therefore, the edge k is incident to v_3 and v_4 in $(\mathcal{F}_j \circ \mathcal{F}_k \circ \mu_2 \circ \nu)(T)$, but also in $\phi_2(T)$ since it is flipped only once (see Figure 6 a). It can be noted that j is already incident to v_1 and v_3 in $\nu(T)$, since it is not flipped between $\nu(T)$ and $(\mu_2 \circ \nu)(T)$. After being flipped in ϕ_1 , j splits the n -gon into two disconnected parts of $(\mathcal{F}_j \circ \mu)(T)$. The first part contains v_3 and the second part contains all together v_1 , v_4 and k (see Figure 6 b). This prevents the subsequent creation of an edge between v_3 and v_4 in $\phi_1(T)$, since v_3 and v_4 are separated by the already flipped j , which is a contradiction. \square




Figure 6: a) Tracking the position of k and j in $(\mathcal{F}_j \circ \mathcal{F}_k \circ \mu_2 \circ \nu)(T)$; b) evolution of j in $(\mathcal{F}_j \circ \nu)(T)$. It can be noted that v_3' can coincide with v_2 and that the position of v_5' with regards to v_5 is not specified.

Lemma 3. Let $T_F \in \mathcal{T}_{lab}$ be a triangulation of a convex n -gon that corresponds to a fan (all the edges are incident to a vertex v_0). We consider ϕ a sequence of edge flips on T_F . If ϕ is reduced on T_F then ϕ is a minimal sequence. Furthermore, any other minimal sequence that is weakly equivalent to ϕ on T_F is indeed strongly equivalent to ϕ by commutativity.

Proof. Let us demonstrate that whatever the reduced sequence of flips in T_F , there is no double occurrence of the same edge flip. This can be done by induction on the number n of vertices in the fan.

- The result is trivial when $n = 3$ or $n = 4$. Let us note that when $n = 5$, reduced sequences are Id , \mathcal{F}_1 , \mathcal{F}_2 , $\mathcal{F}_1 \circ \mathcal{F}_2$ and $\mathcal{F}_2 \circ \mathcal{F}_1$ and none of them contains more than one occurrence of \mathcal{F}_1 and \mathcal{F}_2 .

- Let us assume that the property is true on a convex n -gon, and let us extend it to a convex $(n + 1)$ -gon.

Let ϕ be a reduced sequence in a convex $(n+1)$ -gon triangulated as a fan T_F and k ($0 < k < n-2$) the first edge to be flipped in ϕ . v_0 denotes the apex of the fan. There exists a sequence ϕ' such that $\phi = (\phi' \circ \mathcal{F}_k)$. ϕ' is reduced on $\mathcal{F}_k(T_F)$ since ϕ is reduced on T_F .

It can be observed that flipping an interior edge of a fan has the effect of decreasing the number of edges in this fan. In $\mathcal{F}_k(T_F)$, k splits the $(n+1)$ -gon into a single face and a convex n -gon triangulated as a fan of apex v_0 (see Figure 7).




Figure 7: For $n > 0$, if T_F is a convex $(n+1)$ -gon with a fan triangulation, $\mathcal{F}_k(T_F)$ contains a convex n -gon with a fan triangulation.

Let us assume that ϕ' contains an occurrence of \mathcal{F}_k . Let μ and ν be two sequences such that $\phi = (\nu \circ \mathcal{F}_k \circ \mu \circ \mathcal{F}_k)$ and μ contains no occurrence of \mathcal{F}_k . \mathcal{F}_k does not commute with μ since ϕ is reduced.

Therefore, there exists $\mathcal{F}_{k'}$ such that $\phi = (\nu \circ \mathcal{F}_k \circ \mu_2 \circ \mathcal{F}_{k'} \circ \mu_1 \circ \mathcal{F}_k)$ and $\mathcal{F}_{k'}$ is the first flip of the sequence that does not commute with the first occurrence of \mathcal{F}_k , ie $(\mathcal{F}_{k'} \circ \mu_1 \circ \mathcal{F}_k)(T_F) = (\mathcal{F}_{k'} \circ \mathcal{F}_k \circ \mu_1)(T_F)$ and k and k' are incident to a common face in $\mu_1(T_F)$.

The reduced sequence $\mu = (\mu_2 \circ \mathcal{F}_{k'} \circ \mu_1)$ is restricted to the fan included in $\mathcal{F}_k(T_F)$ since it contains no occurrence of \mathcal{F}_k . By induction hypothesis, there is no double occurrence of the same flip in $(\mu_2 \circ \mathcal{F}_{k'} \circ \mu_1)$.




Figure 8: k' splits the triangulation $(\mathcal{F}_{k'} \circ \mu_1 \circ \mathcal{F}_k)(T_F)$ into two disconnected parts: a p -gon triangulated as a fan and a part where k is an inner edge.

Furthermore, k' splits $(\mathcal{F}_{k'} \circ \mu_1 \circ \mathcal{F}_k)(T_F) = (\mathcal{F}_{k'} \circ \mathcal{F}_k \circ \mu_1)(T_F)$ into two disconnected parts (see Figure 8). Let T_{F_p} denote the part that corresponds to a p -gon triangulated as a fan with $p < n$, and T_k denotes the second part that contains k as an inner edge. It can be noted that all the edges in T_k that are different from k have already been flipped in μ_1 . Therefore, the edges flipped in μ_2 are located inside T_{F_p} exclusively and \mathcal{F}_k can be commuted with μ_2 in $(\mathcal{F}_{k'} \circ \mathcal{F}_k \circ \mu_1)(T_F)$. Thus $\phi(T_F) = (\nu \circ \mu_2 \circ \mathcal{F}_k \circ \mathcal{F}_{k'} \circ \mathcal{F}_k \circ \mu_1)(T_F)$ with $|\text{supp}(k, \mu_1)(T_F) \cap \text{supp}(k', \mu_1)(T_F)| = 1$, which is not possible since ϕ is reduced. This ensures that ϕ' does not contain an occurrence of \mathcal{F}_k . Since ϕ' contains no double occurrence of the same edge split, by induction hypothesis, we have $\phi = (\phi' \circ \mathcal{F}_k)$ which satisfies the same property. Afterwards, Lemma 1 and 2 are used to close the demonstration. \square

We can now prove Theorem 1:

Proof. Let ϕ_1 and ϕ_2 two sequences of flips such that $\phi_1(T) \simeq \phi_2(T)$, we propose to show that one can generate a third sequence that is weakly equivalent to both ϕ_1 and ϕ_2 by using the three moves introduced in section 3.1.

Let T_F be a fan triangulation of the input convex n -gon and μ a sequence of flips such that $\mu(T) = T_F$. We have $\phi_1(T) = (\phi_1 \circ \mu^{-1} \circ \mu)(T)$ and $\phi_2(T) = (\phi_2 \circ \mu^{-1} \circ \mu)(T)$. Let γ_1 be a reduced sequence of $(\phi_1 \circ \mu^{-1})$ and γ_2 be a reduced sequence of $(\phi_2 \circ \mu^{-1})$. Following Lemma 3, γ_1 and γ_2 are minimal and they are strongly equivalent by commutativity. Thus γ_1 can be generated from γ_2 using the move

based on commutativity. This ensures that a combination of the three moves can be used on ϕ_1 and ϕ_2 to generate a sequence $(\gamma_1 \circ \mu)$ that is weakly equivalent to both ϕ_1 and ϕ_2 on T . This also means that the three moves can be used to construct ϕ_2 from ϕ_1 . \square