

Simultaneous edge flipping in triangulations

J. Galtier¹ F. Hurtado² M. Noy² S. Pérennes³ J. Urrutia⁴

Abstract

We generalize the operation of flipping an edge in a triangulation to that of flipping several edges simultaneously. Our main result is an optimal upper bound on the number of simultaneous flips that are needed to transform a triangulation into another. Our results hold for triangulations of point sets and for polygons.

1 Introduction

Given a triangulation T of a set P of points in the plane, an edge e of T is *flippable* if it is incident to two triangles whose union is a convex quadrilateral C . By *flipping* e we mean the operation of removing e from T and replacing it by the other diagonal of C . In this way we obtain a new triangulation T' of P , and we say that T' has been obtained from T by means of a *flip*.

Local transformations on triangulations such as flips have been used in various fields starting with a simple greedy algorithm by Lawson that constructs the Delaunay triangulation of a point set by successive flips from an arbitrary initial triangulation of the point set (see [9]). Their main characteristic is to change incrementally the mesh towards something of better quality. For instance, local changes are used in compression (or simplification) techniques for visualization [17]. An evolving triangulation may be associated with different types of physical data, that require to be smoothly adapted to a new topology [6]. When one of the fields associated to the vertices is height, then changes in the triangulation may be suitable in order to keep some quality measures of the mesh [7]. On the other hand, finite-element applications often associate some aspect ratio to a mesh, that can be improved by flipping operations [12, pages 240-247]. Topology changes occur when the measure taken is non-isotropic [11] and the computation is adaptive [3].

From a different point of view, we remark the existence of a bijection between triangulations of a convex $(n + 2)$ -gon and binary trees with n internal nodes. Under this bijection, flipping an edge in a triangulation corresponds precisely to a *rotation* in the corresponding binary tree [18, 13]. Finally, we note that the flip operation is also applied to triangulations in higher dimensions or to topological triangulations [1, 2, 8, 15, 16, 19].

In a previous paper [14], the authors studied several questions about flips in triangulations, mainly the question of how many flips are needed to transform a triangulation of a plane point set (or of a simple polygon) into another triangulation. Among other results, it was shown that two triangulations of a set of n points (or of a simple polygon with n vertices) in the plane are at most $O(n^2)$ flips apart, and that this bound can be made sensitive to some geometric characteristics. Moreover, pairs of triangulations were produced where $\Omega(n^2)$ flips are necessary, both for simple polygons and for point sets.

¹France Télécom R& D and Projet Mascotte (INRIA / UNSA / CNRS).

²Departament de Matemàtica Aplicada II, Universitat Politècnica de Catalunya, Barcelona, Spain. Partially supported by Projects DGES-SEUID PB98-0933, MCYT-BFM2001-2340, MCYT-FEDER BFM2002-0557 and Gen. Cat 2001SGR00224.

³SLOOP Project (CNRS I3S-INRIA) and CNET, Sophia Antipolis, France.

⁴Instituto de Matemáticas, Universidad Nacional Autónoma de México. Supported by CONACYT grant 37540-A and grant PAPIIT.

It seems natural to go one step further and to allow several edges to be flipped simultaneously, or in parallel. For this operation to make sense the edges must be independent in the sense that no two of them can be sides of the same triangle. Let us call this operation a *simultaneous flip* or a *parallel flip*. Then it is reasonable to expect that allowing parallel flips as a primitive operation the above $O(n^2)$ bound can be decreased substantially. The main purpose of this paper is to show that this is indeed the case.

The parallel flip problem can be addressed in two different ways. In the first one, a centralized controller computes the flips that have to be performed. The cost of this computation is considered as irrelevant and we measure the complexity only in terms of the number of parallel flips required. In the second one there is no central control, the triangulation evolves in a distributed manner, and the decision to perform a flip or not must be driven by local information. In this paper we show how to solve optimally the controlled problem.

We first prove that one can transform any triangulation of a convex n -gon into any other one with at most $O(\log n)$ parallel flips, and that this bound is tight. We remark that when translated in terms of binary trees, this result means that one can transform any binary tree into any other one using at most $O(\log n)$ “parallel rotations”, a fact of independent interest.

For triangulations of general polygons and point sets we obtain an optimal upper bound of $O(n)$ parallel flips. The lower bound $\Omega(n)$ will follow from a construction in [14].

Our last result is that every triangulation of a set of n points contains a set of $(n-4)/6$ edges that can be flipped in parallel, and that there are triangulations in which at most $(n-4)/5$ edges can be flipped in parallel. In [14] it was proved that every triangulation contains at least $(n-4)/2$ flippable edges.

The organization of the paper is as follows. In Section 2 we state and discuss in detail the results of the paper. Proofs are given in Section 3.

2 Results

The following definitions apply both to triangulations of polygons and to triangulations of point sets. An edge e and a triangle t of a triangulation are said to be incident if e is one of the sides of t . Two edges e and f of a triangulation T are called *flip-independent* if no triangle of T is incident to both e and f . A set E of edges is called *flippable* if every edge of E is flippable (in the ordinary sense) and if they are pairwise flip-independent. Finally, given a flippable set E , the operation of *flipping* E consists of flipping simultaneously all the edges of the set. This gives rise to a new triangulation T' , and we say that T' has been obtained from T by means of a *simultaneous flip* or a *parallel flip* (see Figure 1 for an example). Observe that because of the assumption of flip-independence, the operation is well defined. In particular, it does not matter in which order the edges of E are flipped when considered as sequential flips. Observe also that ordinary flips are a particular case of parallel flips.

By *introducing* a diagonal e into a triangulation T we mean to perform a sequence of flips starting from T until e is one of the edges in the current triangulation.




Figure 1: Flipping in parallel the set of thick edges of T produces T' .

Our first result deals with convex polygons. Let us mention that, in the case of a convex n -gon, $O(n)$ ordinary flips always suffice to transform one triangulation into another (this is easy to prove, the hard problem is to determine the exact value for the maximum number of ordinary flips required [18]).

Theorem 1 *Any triangulation of a convex n -gon can be transformed into any other triangulation using at most $O(\log n)$ parallel flips, and this bound is tight.*

Besides its geometric content, this result can be translated into the language of binary trees, where parallel flips correspond to *parallel rotations*, that is, several rotations that take place without conflict simultaneously. While $\Omega(n)$ (sequential) rotations are eventually necessary to transform a binary tree into another one [18], Theorem 1 means that $O(\log n)$ parallel rotations are always sufficient.

Next we turn to triangulations of arbitrary polygons. We have already mentioned the example given in [14] of two triangulations of a polygon that are $\Omega(n^2)$ ordinary flips apart. Since a triangulation has a linear number of edges, this implies that the two triangulations are at least $\Omega(n)$ parallel flips apart. In order to obtain a matching upper bound the main ingredient is the following result, whose proof requires a series of technical lemmas. We recall that a vertex v of a polygon is called *convex* when the internal angle at v is smaller than π , and is called *concave* or *reflex* otherwise.

Proposition 1 *Let T be a triangulation of a simple polygon Q_n and let e be a diagonal not in T . Then e can be introduced in T with at most $O(n \log c)$ parallel flips, where c is the number of convex vertices in Q_n .*

Now come the two main results of the paper.

Theorem 2 *Any triangulation T of a simple polygon Q_n can be transformed into any other triangulation T' using at most $O(n)$ parallel flips, and this bound is tight.*

From this we can deduce the same result for triangulations of points sets. We remark that again the example of triangulations at (sequential) quadratic distance shows that $\Omega(n)$ is a lower bound.

Theorem 3 *Any triangulation T of a set P_n of n points on the plane can be transformed into any other triangulation T' using at most $O(n)$ parallel flips, and this bound is tight.*

Finally, in the above mentioned paper [14], it was proved that any triangulation of a set of n points contains at least $(n - 4)/2$ edges that can be flipped. Here we present an analogous result for parallel flips.

Theorem 4 *Every triangulation T of a set P_n of n points on the plane contains a set of at least $(n - 4)/6$ edges of T that can be flipped in parallel. Also, for every n there exists a triangulation of a set of n points in which at most $(n - 4)/5$ edges can be flipped in parallel.*

3 Proofs

The proof of Theorem 1 requires some preliminaries. Given a triangulation T of a convex polygon, its *dual tree* \widehat{T} is defined as the dual graph of T excluding the unbounded face (see Fig. 2, top left, where the dual tree is shown with dashed edges). The leaves of \widehat{T} correspond to the *ears* of T , that is, to vertices of the polygon such that its two neighbors are adjacent in T . The diameter of \widehat{T} is the maximum distance between two nodes of \widehat{T} . Finally, a triangulation is called a *fan* if there is a vertex v of the polygon adjacent to all the other vertices; v is called the *apex* of the fan.

Lemma 1 *Let T be any triangulation of a convex polygon such that the diameter of \widehat{T} is k , and let v be any vertex. Then T can be transformed into the fan with apex v using at most k parallel flips.*

Proof. Let e be an edge of T not incident with v . The distance from v to e is defined to be the number of edges of T intersected by any line segment joining any interior point of e to v , plus one (see Fig. 2, top left), and is at most k by hypothesis. Equivalently, consider the diagonals not incident with v : those which are *visible* from v get distance 1, when they are all flipped (and hence become incident with v) the new set of visible diagonals get distance 2, and so on. By construction the edges at distance i can be flipped in parallel once all edges at distances $1, \dots, i - 1$ have been previously flipped in this order, therefore if we flip at step i all the edges at distance i , we reach the fan with apex v using at most k parallel flips (see Fig. 2). \square

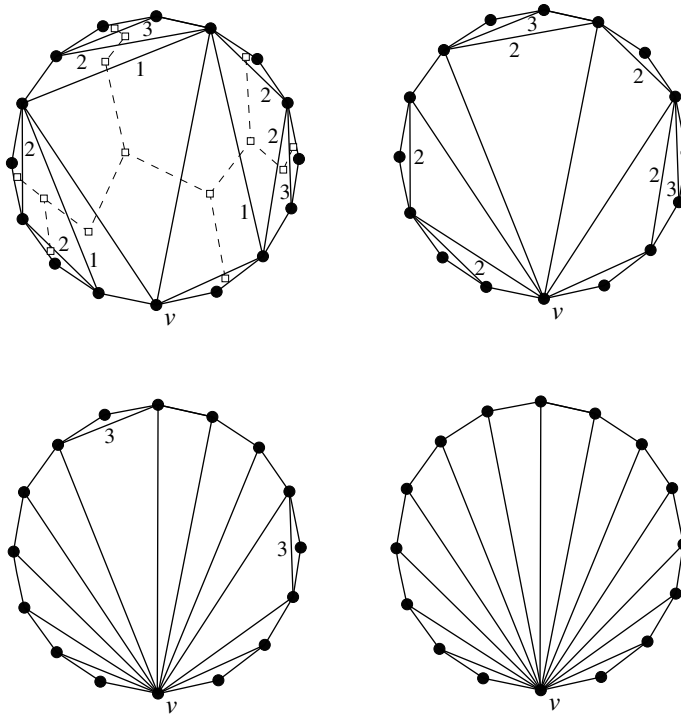


Figure 2: Transforming a triangulation into a fan.

Lemma 2 *Let T be a triangulation of a convex n -gon. Then with at most two parallel flips one can transform T into a triangulation having at least $n/6$ ears.*

Proof. The vertices of the dual tree of a triangulation have degrees 1, 2 or 3. The goal is to reduce the number of vertices of degree 2 thus increasing those of degree 1, i.e. the leaves, that correspond to the ears. Consider three consecutive vertices of degree 2 in \widehat{T} . As shown in Fig. 3 with at most two flips one of them can be transformed into a leaf. If \widehat{T} contains long paths consisting of vertices of degree 2, subdivide them into subpaths of length three, plus a possible remainder of length one or two. If the remainder is one, it is given to one of the extremes of the path; if it is two, each one of them is given to one of the two extremes. From the previous remark, using at most two parallel flips we can eliminate these paths and obtain a new triangulation T_1 such that \widehat{T}_1 contains no paths of length greater than two consisting of vertices of degree 2.

Finally we show that T_1 has at least $n/6$ ears. Substitute every path of vertices of degree 2 in \widehat{T}_1 by a single edge. Then we obtain a full binary tree having m internal nodes and $m + 1$ leaves. By construction, every edge of this tree gives rise to at most two vertices of degree 2 in \widehat{T}_1 . Then \widehat{T}_1 has $m + 1$ leaves out of a total of at most $m + (m + 1) + 2 \cdot 2m = 6m + 1$ nodes. \square




Figure 3: With 1 or 2 flips an ear (dashed) and a triangle of degree 3 are created.

Proof of Theorem 1. Let T be a triangulation of a convex n -gon P and let v be any vertex of P . By Lemma 2 we can transform T into another triangulation T' having at least $n/6$ ears. Let P' be the polygon obtained from P by removing all the ears of T' (except those incident with v). By recursive application of Lemma 2 to the triangulation induced on P' by T' we obtain a triangulation T'' of P whose dual tree has a logarithmic diameter. By Lemma 1, T'' can be transformed into the fan with apex v . As a consequence T can be transformed with $O(\log n)$ parallel flips into the fan with apex v and the upper bound follows.

Finally we show that the bound is tight. Let S be a triangulation having v as an ear and let T_v be the fan with apex v . A parallel flip applied to T_v can remove at most half of the edges incident with v , and this implies that $\Omega(\log n)$ parallel flips are required to reach S . \square

Our next result, Proposition 1, is technically the most complex and requires a series of three lemmas.

Lemma 3 *Let Q_n be a polygon with vertices $v_1, \dots, v_k, \dots, v_n$ such that v_1, v_k, v_{k+1} and v_n are the only convex vertices, and such that there exists a line r that crosses the boundary of the polygon Q_n only at sides $v_k v_{k+1}$ and $v_n v_1$. Then $O(n)$ parallel flips are sufficient to transform any triangulation of Q_n into any other one.*

Proof. For ease of exposition we assume that r is horizontal. Let u_0, u_1, \dots, u_m be the vertices on the upper chain, and let l_0, l_1, \dots, l_p be those on the lower chain, in both cases in clockwise order (see Figure 4a), so that u_0, u_m, l_0, l_p are the convex vertices. Any triangulation of Q_n consists of $m+p$ triangles, m of them having one side on the upper chain and the third vertex on the lower chain, and p of them in the opposite way. By Δ_i ($i = 1, \dots, p$) we denote the triangle having base $l_{i-1}l_i$. If in some triangulation, the triangle Δ_i has vertices $l_{i-1}l_i u_t$, then we say that u_t is the *top vertex* of Δ_i .

Define a *target triangulation* T^* as follows. Join vertex u_m with all the vertices seen by it from the lower chain, next join u_{m-1} with all the vertices it sees from the lower chain not joined previously with u_m , and so on. T^* can be described as the triangulation in which the top vertices of $\Delta_1, \dots, \Delta_p$ are as much to the right as possible (Figure 5). If two triangles $\Delta_{i+1}\Delta_i$ share a diagonal, then this diagonal cannot be flipped, and the same applies to two triangles with base on the upper chain. It may also happen that visibility restrictions do not allow to flip a diagonal shared by two triangles of different kind.

We define a *forward flip* as the simultaneous (parallel) flip of all the diagonals that can be flipped so that the top vertex of every Δ_i advances to the right, if this is possible, or remains




Figure 4: a) Triangulation of Q_n . b) Result after applying a forward flip.

where it is if no advance is possible. Figure 4b shows the result of applying a forward flip to the triangulation in Figure 4a. Observe that a forward flip is uniquely defined at every step, and that it can always be performed unless the target T^* has been reached. This is because if $\Delta_1, \dots, \Delta_{i-1}$ are already in their final position, and Δ_i is not, the top vertex of Δ_i can be advanced.

Next we show that $m + p - 1$ forward flips are always enough to reach T^* from any triangulation, and this will prove the lemma. More precisely, we show that if initially u_t is the top vertex of Δ_i then, with at most $m + i - t - 1$ forward flips, $\Delta_1, \dots, \Delta_i$ reach their final positions. The proof is by induction on i .

For $i = 1$ the result is clear, since the top vertex u_t of Δ_1 can always advance up to its final position, which is u_m in the worst case, and this would require $m - t = m + 1 - t - 1$ flips. Now assume that the initial top vertex of Δ_i is u_t and that of Δ_{i-1} is u_s , and that by induction the first $m + i - s - 2$ forward flips send $\Delta_1, \dots, \Delta_{i-1}$ to their final positions. If these flips also take Δ_i to its final position we are done, otherwise let's see how many additional forward flips are required.

Consider first the case $s > t$. Let us perform the $m + i - s - 2$ forward flips which take the top vertices of $\Delta_1, \dots, \Delta_{i-1}$ to their final positions. At each of them three possibilities arise:

- Both the top vertex of Δ_i and the top vertex of Δ_{i-1} advance; in this case the number of edges which separate them in the upper chain remains the same as it was before the forward flip;
- the top vertex of Δ_i advances one edge but the top vertex of Δ_{i-1} doesn't; in this case the separation decreases by one;
- the top vertex of Δ_{i-1} advances one edge but the top vertex of Δ_i doesn't. For this to happen it is necessary that Δ_{i-1} and Δ_i share an edge before the flip, hence in this case the separation between their top vertices will change from 0 to 1.

Initially the separation between the top vertices of Δ_{i-1} and Δ_i in the upper chain consists of $s - t$ edges, where $s - t \geq 1$. According to the above cases after $m + i - s - 2$ forward flips the separation between them is at most $s - t$. Therefore with at most $s - t$ additional flips Δ_i reaches its final position. The total number of forward flips would be $(m + i - s - 2) + (s - t) = m + i - t - 2 < m + i - t - 1$.

In the case $s = t$, that is, when Δ_i and Δ_{i-1} have the same initial top vertex, we have just seen that the separation between the top vertices is at most 1 after $m + i - s - 2$ forward flips. Hence, at most $(m + i - s - 2) + 1 = m + i - s - 1 = m + i - t - 1$ forward flips suffice. \square

Observation 1. We remark that if in the above lemma $v_k = v_{k+1}$ (or $u_0 = l_p$ in the proof), so that there are only three convex vertices, the result still applies.




Figure 5: Target triangulation.

Lemma 4 Let Q_n be a polygon with vertices $v_1, \dots, v_p, \dots, v_q, \dots, v_n$ such that a) v_1, v_p, v_q, v_{q+1} and v_n are the only convex vertices; b) v_1 and v_q see each other; c) Q_n admits a triangulation such that all the diagonals cross v_1v_q . Then $O(n)$ parallel flips are sufficient to introduce the diagonal v_1v_q in T .

Proof. For ease of exposition we assume that v_1v_q is horizontal, v_p is above v_1v_q , and v_{q+1} and v_n are below it (see Figure 6). At each step of the process we call *active polygon* the subpolygon of Q_n formed by the union of triangles which cross v_1v_q in the current triangulation; at the beginning the active polygon is the whole Q_n . Flips will be performed always inside the active polygon, i. e., at each step the diagonals in the current triangulation of Q_n external to the active polygon are frozen and will not be flipped in the sequel.

All the vertices in the chain $v_{q+1} \dots v_n$ see either v_q or v_1 . Assume that v_{q+1}, \dots, v_m see v_q , and v_m, \dots, v_n see v_1 . In order to introduce v_1v_q it is enough to introduce first v_1v_m and v_mv_q , since then in the polygon $v_mv_1 \dots v_p \dots v_qv_m$ the diagonal v_1v_q can be introduced with $O(n)$ flips (even sequential flips, as shown in [14, Lemma 3.1]).




Figure 6: Two possible initial situations in Lemma 4.

If v_m is joined in T to v_i and v_j with $1 \leq i \leq p$ and $p \leq j \leq q$ (see Fig. 6, left), the claim is clear since then Lemma 3 can be applied to polygons $v_1 \dots v_iv_mv_m \dots v_nv_1$ and $v_qv_{q+1} \dots v_mv_j \dots v_q$, and v_mv_1 and v_mv_q can be introduced with $O(n)$ parallel flips.

Let us assume that v_m is connected to v_i ($1 \leq i \leq p$), but not to any vertex in the chain $v_{i+1} \dots v_p \dots v_q$. Let v_t be such that the triangles with bases $v_{q+1}v_{q+2}, \dots, v_{t-1}v_t$ have their top vertices in the chain $v_p \dots v_q$, while triangles with bases $v_tv_{t+1}, \dots, v_{m-1}v_m$ have their top

vertices in $v_i \cdots v_p$ (Figure 6, right). Let v_s be the vertex connected to v_{t-1} and v_t ; then Lemma 3 applies to polygons $v_i v_m \cdots v_n v_1 \cdots v_i$ and $v_s \cdots v_q v_{q+1} \cdots v_t v_s$ introducing $v_t v_q$ and $v_m v_1$ with $O(n)$ parallel flips (Figure 7, left).

As in the proof of the previous lemma, let u_t, \dots, u_{m-1} be the top vertices of triangles with bases $v_t v_{t+1}, \dots, v_{m-1} v_m$, respectively. After the step in the preceding paragraph, we have $u_t = v_h$ and $u_{m-1} = v_i$ (Figure 7, right), and the active polygon is $v_1 \cdots v_p \cdots v_q v_t \cdots v_m v_1$. Our final goal is to have $u_t = u_{t+1} = \cdots = u_{m-1} = v_q$.




Figure 7: Situation corresponding to Lemma 4 after the first round of flips.

Inside polygon $F_t = v_t v_h \cdots v_p \cdots v_q v_t$ we introduce next the tangent $v_h v_r$ from $u_t = v_h$ to the chain $v_p \cdots v_q$. Let us show that this can be done sequentially with a cost linear in the number of vertices that end up above this tangent. Vertices v_p and v_t are convex in F_t , hence the diagonal joining them can be flipped; in this way we obtain the diagonal $v_{p-1} v_{p+1}$ which we call the current *bridge*. We update F_t to be now the polygon $v_t v_h \cdots v_{p-1} v_{p+1} \cdots v_q v_t$. The flip just performed is called a *bridge-flip*, and it is charged to the vertex v_p . Observe that at least one of the vertices v_{p-1} and v_{p+1} must be a vertex of the convex hull of the polygon $v_1 \cdots v_{p-1} v_{p+1} \cdots v_q v_1$; hence the diagonal joining that vertex to v_t can be flipped, say to $v_{p-2} v_{p+1}$. This is again a bridge-flip, which we charge to v_{p-1} ; the new bridge is $v_{p-2} v_{p+1}$ and F_t is updated to be $v_t v_h \cdots v_{p-2} v_{p+1} \cdots v_q v_t$. The process is iterated until the tangent $v_h v_r$ from $u_t = v_h$ to the chain $v_p \cdots v_q$ is introduced. Notice that the vertices above the bridge are no longer part of the active polygon and therefore will never be charged again. An example of this process is shown in Figure 8.




Figure 8: A sequence of bridge-flips. Current bridges are the thicker lines, and the vertices which get charged at each step are encircled.

The situation is now as in the left part of Figure 9. Inside polygon $v_q v_t v_{t+1} v_h v_r \cdots v_q$ the only diagonal that can be flipped is $v_t v_h$ and thus u_t moves to the upper left chain $v_p \cdots v_q$ with

this flip, which we call *jump-flip* because one may think that u_t “jumps” from the upper right chain to the upper left chain. A jump-flip is shown in Figure 9, right.




Figure 9: A jump-flip.

We can now imagine that u_t starts “travelling” towards v_q , with a total of $q - r$ required flips, advancing one position at each step: for this reason we denote this kind of flips as *advance-flips*. Nevertheless, let us consider the moment when u_t is still at v_r . Inside the polygon $F_{t-1} = v_r v_{t+1} u_{t+1} \cdots v_h v_r$ we can perform a bridge-flip simultaneously with the first advance of u_t towards v_q , the polygon loses one triangle, which goes above the bridge, and gains one triangle (see Figure 10). We keep doing parallel advance-flips and bridge-flips until either u_t reaches v_q (in this case we would continue just doing bridge-flips) or we introduce the tangent from u_{t+1} to the upper left chain. In the latter situation we would perform a jump-flip for u_{t+1} ; then a new bridge-flip would be performed in parallel with the simultaneous advance-flips of both u_t and u_{t+1} , and the process would be iterated.




Figure 10: An advance-flip and a bridge-flip which are done in parallel. It is also shown how the polygon F_{t-1} (shaded) is updated.

Two more details have to be specified before we can analyze the cost of this process. First, notice that we are doing advance-flips simultaneously with bridge-flips; this is true until the jump of u_{m-1} , because from then on no more bridge-flips are performed. Second, a particular situation arises when several top vertices, say u_{t_1}, \dots, u_{t_r} , are at the same position in the upper right chain (as is the case of u_{t+1} and u_{t+2} in Figure 10). In this case we perform sequentially all the corresponding jump-flips, one by one. After that u_{t_1}, \dots, u_{t_r} are at the same position in the upper left chain. Only u_{t_1} is now able to advance towards v_q , but once this flip is done, both u_{t_1} and u_{t_2} may advance simultaneously in a single parallel flip. After this second flip u_{t_1} , u_{t_2} and u_{t_3} are able to advance simultaneously in a single parallel flip, and so on. We don’t consider these $r - 1$ flips as advance-flips: as they are done in order to separate the top vertices,

we call them *split-flips*, and count them separately as well. Once $u_{t_{r-1}}$ and u_{t_r} are at different positions in the upper left chain we may resume on doing advance-flips. An example is shown in Figure 11.




Figure 11: The flips a , b and c are jump-flips, while d and e are split-flips.

The advance-flips after the jump of u_{m-1} are at most $q - p$. All other advance-flips are done in parallel with bridge-flips, and these are charged once and only once to a vertex, which gives at most q flips. The jump-flips are $m - t$ and, finally, at most $m - t$ split-flips are done. This gives a total of at most $2n$ parallel flips, because

$$(q - p) + q + (m - t) + (m - t) \leq 2q + 2(m - t) = 2(q + m - t) \leq 2n.$$

□

Observation 2. We remark that if in the above lemma $v_q = v_{q+1}$, the result still applies.




Figure 12: Situation corresponding to Lemma 5.

Lemma 5 *Let Q_n be a polygon such that x, u, y and v are the only convex vertices in counter-clockwise order, and such that uv and xy are diagonals of Q_n . Then any triangulation of Q_n containing xy can be transformed into another triangulation containing uv with $O(n)$ parallel flips.*

Proof. Let T be a triangulation containing xy . As x and y cannot both have degree one in T , we can assume that y has neighbors different from x in the chain $x \dots u$, as in Fig. 12a.

Let now x_1 be the last vertex in the chain $x \dots u$ connected to y , and let $x_1 w$ be the tangent from x_1 to $v \dots x$. Let also w' be the neighbor of w in the chain $w \dots x$, and let y_1 be the vertex connected to w and w' . Use Lemma 4 in the polygon $x \dots x_1 y \dots y_1 w \dots x$ to introduce $x_1 w$

and x_1y_1 with cost proportional to the total length of chains xx_1, yy_1 and wx . Similarly, next introduce the tangent y_1z from y_1 to $u \cdots y$, and the diagonal y_1t , where the triangle $zz't$ is defined as $wy'y_1$ (see Figure 12b). The cost is proportional to the total length of $x_1 \cdots t, z \cdots y$ and $y \cdots y_1$.

Now repeat the process in the new polygon $u \cdots zy_1 \cdots v \cdots wx_1 \cdots u$, with x_1 and y_1 playing the role of x and y in the starting polygon. If $t = x_1$ we have a smaller polygon in the situation of the lemma, in which no cost has been charged to the edges in the previous step. If $t \neq x_1$ then y_1 has neighbors different from x_1 in the chain $x_1 \cdots u$ and we are in the previous situation, with the additional fact that the chain $x_1 \cdots t$ has already contributed once to the cost. Nevertheless the chain $x_1 \cdots t$ plays the role of the chain $x \cdots x_1$ which pays once and disappears.

At the end the diagonal uv has been introduced; the edges in chains $v \cdots x$ and $u \cdots y$ have been charged once, and edges in $y \cdots v$ and $x \cdots u$ have been charged twice. The total number of flips is thus $O(n)$. \square

Equipped with the above three lemmas, we proceed to the proof of Proposition 1.

Proof of Proposition 1. Assume that the diagonal $e = uv$ to be inserted is horizontal, and let P_e be the polygon formed by all the triangles of T that cross e . The polygon P_e consists of two chains from u to v , let us call them the upper and the lower chains (see Figure 13, top). Let c be the number of convex vertices of P_e ; our goal is, with at most $O(n)$ parallel flips, to reduce the number of convex vertices in P_e from c to $3c/4$. This is done with the help of the above three lemmas. Thus iterating the procedure e will be introduced into T in $O(\log c)$ steps, each of them requiring $O(n)$ parallel flips, and the result will follow.




Figure 13: Illustrating the proof of Proposition 1.

Let v_i, v_j, v_k be three consecutive convex vertices in P_e , i.e. all the vertices between v_i and v_j , and between v_j and v_k are reflex (Figure 13, top). Then we define the *tangent* τ_j associated to v_j as the segment tangent to the two reflex chains (possibly reduced to v_i or v_k) neighboring v_j . Assume that at least $c/2$ convex vertices of P_e are in the lower chain. Now the goal is to introduce $c/4$ tangents, essentially one for every two consecutive convex vertices in the lower chain, the updated P_e will have at most $3c/4$ convex vertices (Figure 13, bottom).

Let $y = v_j$ be a convex vertex in the lower chain and let $ab = \tau_j$ the associated tangent. Consider the polygon P_{ab} formed by all the triangles that cross ab , and let k be the number of convex vertices in the upper chain (excluding the two extremes) belonging to P_{ab} . Several cases arise depending on the value of k ; in some of them we are introducing the tangent ab ,

which is always the initial target, but in some cases we are introducing instead a tangent to the upper chain, also contained in P_{ab} . In all these situations a tangent is introduced in a number of parallel flips which is linear in the size of P_{ab} . As the polygons P_{ab} associated with different tangents have disjoint interiors the total number of parallel flips is $O(n)$.

We show next how to proceed case by case, depending on the values of k , which will conclude the proof. For the sake of clarity we assume that the line ab is horizontal.




Figure 14: The cases $k = 0$ and $k > 2$ in the proof of Proposition 1.

Case $k = 0$. In this situation we simply apply Lemma 4 in order to introduce ab (see Figure 14 left).

Case $k = 1$. Let x be the opposite convex vertex in the upper chain. We can assume that x is joined to y : otherwise x is joined to some z between, say, a and y ; then we introduce the tangent from z to the chain $y \cdots b$ using Lemma 4 applied to the union of triangles crossed by this tangent, and z becomes y (Figure 15, left). Once we are in the situation where x is joined to y (see Figure 15, right), we apply Lemma 3 and place the tangents from a and b to the corresponding upper reflex chains, and we are in the situation covered by Lemma 5, allowing us to introduce ab .




Figure 15: The case $k = 1$ in the proof of Proposition 1.

Case $k > 2$. Let x be one of the convex vertices in the upper chain, different from the leftmost one and the rightmost one, and let γ be the tangent associated with x (Figure 14, right). Instead of introducing ab we introduce in this case γ , this will also make one convex vertex from P_e disappear, namely x . Notice that once γ plays the role of ab , i.e., we only consider the union of triangles that cross γ . Now, from the viewpoint of γ , we are either in the case $k = 0$ or in the case $k = 1$, already considered, and thus we are done.

Case $k = 2$. Let x and z be the internal convex vertices of the upper chain.

If the tangent γ in P_e connecting the reflex chain between x and z with one of the neighboring reflex chains lies entirely inside P_{ab} , instead of introducing ab we introduce γ , and notice that from the viewpoint of γ we are in the situation $k \leq 1$ (see Figure 16, left).

Otherwise let q be the rightmost vertex of the upper chain, and let $\gamma = pq$ be the tangent from q to the reflex chain between x and z (see Figure 16, right). We first introduce pq ; from the viewpoint of this diagonal we are in situation $k \leq 1$. When the polygon $z \cdots pq \cdots z$ is cut from P_e the convex vertex z will be gone; nevertheless we cannot stop here, as we did in some previous cases, because it is possible that q were reflex in the initial P_e but becomes convex in the updated P_e . Now in a second round we turn back to introducing ab , which at this moment is in the situation $k = 1$. \square




Figure 16: The case $k = 2$ in the proof of Proposition 1.

Proof of Theorem 2. Our proof is by induction on n . We denote by $T(n)$ an upper bound on the number of parallel flips needed to transform any two triangulations of a simple n -gon into each other. Let $e = uv$ be a *long diagonal* of T' , that is, a diagonal that splits the boundary of Q_n into two chains, each of them of length at most $2n/3$. The existence of this diagonal (in fact a much stronger result) is proven in [5], where constructive algorithms are also given. Let $A(n)$ be an upper bound on the number of parallel flips needed to introduce e in T . Then clearly

$$T(n) \leq T(\lfloor 2n/3 + 1 \rfloor) + A(n).$$

Let P_e be the polygon formed by the union of all the triangles of T that cross e (see Fig. 17). Since the diagonals of the triangulation T_e induced by T on P_e are linearly ordered by their crossings with e , we can find three diagonals $e_i = x_i y_i$, $i = 1, 2, 3$, of T_e that split P_e into four subpolygons R_1, R_2, R_3, R_4 of size at most $\lfloor n/4 + 2 \rfloor$.




Figure 17: Splitting the polygon P_e associated to the diagonal $e = uv$.

We next transform the triangulation T_i induced by T in R_i into a triangulation T_i^0 including two concave chains that do not cross the edge e ; one of them links x_{i-1} to x_i and the other links y_{i-1} to y_i , where $x_0 = y_0 = u$ and $x_4 = y_4 = v$. They are simply shortest paths inside R_i . By induction, all the T_i^0 can be reached using at most $T(\lfloor n/4 + 2 \rfloor)$ parallel flips.

Now the polygon P_e^0 whose boundary is the set of concave chains has at most 8 convex vertices. By Proposition 1, e can be introduced using αn parallel flips for some $\alpha > 0$. Therefore

$$A(n) \leq T(\lfloor n/4 + 2 \rfloor) + \alpha n.$$

Finally,

$$T(n) \leq T(\lfloor 2n/3 + 1 \rfloor) + T(\lfloor n/4 + 2 \rfloor) + \alpha n,$$

which gives $T(n) \in O(n)$. \square

Observation 3. Let us extend the class of polygons we have been considering by requiring the interior to remain simply connected but allowing the boundary to touch itself externally. In Figure 18 (right) some ways in which this fact may happen are shown. Notice that these polygons are not properly simple as their boundary is not a Jordan curve.

If the simple polygon in Figure 18 (left) differs only infinitesimally from the polygon at the right, in the sense that the internal visibility is exactly the same, both polygons have the same triangulations and these behave identically. Hence the preceding results in this section also apply to the extended class of polygons. This fact is used in the proof of the next theorem, because certain polygons are constructed by gluing adjacent triangles and situation (a) in Figure 18 might arise.




Figure 18: The simple polygon to the left and the externally self-touching polygon to the right have identical internal visibility.

Proof of Theorem 3. We first define a target triangulation T^* , and then show how to reach the target from any given triangulation.

Performing a rotation if necessary, we can assume that no two points of P_n have the same y -coordinate. Let h be a horizontal line dividing P_n into two balanced subsets, P_n^+ and P_n^- , the points above and below h , respectively (see Figure 19a). Let us consider the polygonal region R of the convex hull $CH(P_n)$ of P_n , comprised between the lower hull of P_n^+ and the upper hull of P_n^- . Let $T^*(R)$ be any fixed triangulation of R : this will be the portion of T^* inside R . Finally, we construct T^* recursively inside $CH(P_n^+)$ and $CH(P_n^-)$.

Let T be any given triangulation of P_n and let u_{i-1}, u_i, u_{i+1} be consecutive vertices of the lower hull of P_n^+ (traversed counterclockwise). Observe that, as the angle between the vectors $\overrightarrow{u_i u_{i-1}}$ and $\overrightarrow{u_i u_{i+1}}$ exceeds π , at least one edge e in T must cross h and have one extreme at u_i and the other one at some point in P_n^- . A similar reasoning can be applied to the vertices in the upper hull of P_n^- .

Let us consider the union Q of all triangles in T crossing h (see Figure 19b). This is a polygon which contains R and includes among its vertices all the vertices in the lower hull of P_n^+ and the upper hull of P_n^- , as observed above. By Theorem 2 (Q is not necessarily simple but see Observation 3 after the proof of the theorem) we can perform $O(n)$ parallel flips and obtain inside Q all the edges of the lower hull of P_n^+ , the edges of the upper hull of P_n^- , and the edges of $T^*(R)$.




Figure 19: a) Defining T^* . b) Polygonal region Q (shaded) of T .

We can now repeat the process in parallel inside $CH(P_n^+)$ and $CH(P_n^-)$. In this way we reach T^* with a total of

$$O\left(n + \frac{n}{2} + \frac{n}{4} + \dots\right) = O(n)$$

parallel flips, and this proves the result. \square

Proof of Theorem 4. Let E be a set of flippable edges from T with the maximum cardinality. We know that $|E| \geq (n-4)/2$ as at least this number of edges are flippable in any triangulation, as shown in [14]. In general, not all of them can be flipped in parallel, since they may not be flip-independent. Complete T combinatorially by adding a point w in the external face and connecting w to all the points in the convex hull of P_n , and let \hat{T} be the dual graph of the completed triangulation. As a bridgeless planar cubic graph, \hat{T} can be edge-coloured with 3 colours [4, p. 254]. The edges of \hat{T} are in a one-to-one correspondence with the edges of T ; let \hat{E} be the set of edges associated with E . Any set of edges in \hat{E} pairwise not sharing any node corresponds to a flip-independent set in E . If we take the subset of edges in \hat{E} with the most frequent colour of the three, we obtain a flip-independent set of at least $(n-4)/6$ edges.

To construct a triangulation in which at most $(n-4)/5$ edges can be flipped in parallel we proceed as follows. Take a triangulation T of a convex polygon whose dual tree τ is a complete balanced binary tree of odd height h (here the tree is rooted at an internal triangle and the root has out-degree 3). In every triangle of T corresponding to a node in τ at an odd level, i.e. at levels $1, 3, \dots, h$, place a new vertex and connect it to the three vertices of the corresponding triangle (this is illustrated in Figure 20 for $h = 3$). A simple counting shows that the triangulation so obtained has $n = 10 \cdot 4^{(h-1)/2} - 1 = 10 \cdot 2^{h-1} - 1$ vertices, and only $2 \cdot 4^{(h-1)/2} - 1 = 2 \cdot 2^{h-1} - 1 = (n-4)/5$ diagonals can be flipped in parallel, one for every empty triangle of the original convex polygon. \square

4 Conclusions and open problems

In this paper we have shown that $O(n)$ parallel flips are sufficient to transform any triangulation of a simple polygon or of a point set into any other triangulation, and that this bound is tight. From the algorithmic point of view this number of flips is the right measure of complexity assuming that a centralized controller computes, with cost considered as irrelevant, the flips that have to be performed in every step.




Figure 20: The six lunes on the left are as the drawing on the right. Only $7 = (39 - 4)/5$ edges can be flipped in parallel, one in the root triangle and one in each lune.

The main open question is the uncontrolled problem in which the triangulation evolves in a distributed manner and decisions are driven by local information. Some heuristics for this problem are described in [10]. Another open problem appears if we restrict the simultaneous flip operation by requiring that all the single flips involved are locally Delaunay. Then it is natural to ask how many of these “parallel Delaunay flips” are necessary in order to reach the Delaunay triangulation from any given starting triangulation.

From a different point of view, it would be interesting to find more precise bounds, as was done in [14] for sequential flips, depending on the number of reflex vertices in a simple polygon, or on the number of convex layers of a point set. Finally, it is an open question to close the gap between the fractions $1/6$ and $1/5$ in Theorem 4.

References

- [1] D. Avis, Generating rooted triangulations without repetitions, *Algorithmica* 16 (1996), 618-632.
- [2] D. Avis and K. Fukuda, Reverse search for enumeration, *Discrete Applied Math.* 65 (1996), 21-46.
- [3] N. Di Cesare, O. Pironneau and E. Polak, Consistent approximation of an optimal design problem, In *Publications du Laboratoire d'Analyse Numérique*, volume 17, page R98005. Université Pierre et Marie Curie, 1998.
- [4] G. Chartrand and L. Lesniak, *Graphs and digraphs, 3rd edition*, Chapman and Hall (1996).
- [5] B. Chazelle, A theorem on polygon cutting with applications, Proc. 23rd Annu. IEEE Sympos. Found. Comput. Sci., pp. 339-349, 1982.
- [6] D. Clark and M. Bailey, Visualization of height field data with physical models and texture photomapping, In *Proceeding of the IEEE Visualization'97*, pp. 89-94, Phoenix, Arizona, October 1997.
- [7] P. Desnougès and O. Devilliers, A locally optimal triangulation of the hyperbolic paraboloid, In *Proceedings of the seventh Canadian Conference on Computational Geometry*, pp. 49-54, Laval, Quebec, August 1995.
- [8] H. Edelsbrunner and N.R. Shah, Incremental Topological Flipping Works for Regular Triangulations, *Algorithmica* 15 (1996), 223-241.
- [9] S. Fortune, Voronoi diagrams and Delaunay triangulations, in *Computing in Euclidean geometry* 2nd edition (1992), D.Z. Du and F.K. Hwang eds., World Scientific, 225-265.

- [10] J. Galtier and S. Perennes, Parallel flips on planar triangulations, in *Proceedings of the 15th European Workshop on Computational Geometry*, pp. 81–86, Antibes, March 1999.
- [11] P.L. George, Automatic 3d mesh generation conforming a prescribed size map, In M. Cross., B. K. Soni, J. F. Thompson, J. Hauser, and P. R. Eiseman, editors, *Numerical Grid Generation in Computational Field Simulations - Proceedings of the 6th International Conference, held at the University of Greenwich*, pp. 21–32, July 1998.
- [12] P.L. George and H. Borouchaki. *Triangulation de Delaunay et maillage – applications aux éléments finis*. Hermes, 1997.
- [13] F. Hurtado and M. Noy, Graph of Triangulations of a Convex Polygon and Tree of Triangulations, *Computational Geometry: Theory and Applications* 13 (1999) 179–188.
- [14] F. Hurtado, M. Noy and J. Urrutia, Flipping edges in triangulations, *Discrete and Computational Geometry* 22 (1999) 333–346.
- [15] B. Joe, Construction of three-dimensional Delaunay triangulations using local transformations, *Computer Aided Geom. Design* 8 (1991), 123–142.
- [16] M. Pocchiola and G. Vegter, Topologically sweeping visibility complexes via pseudotriangulations, *Discrete Comput. Geom.* 16 (1996), 419–453.
- [17] W.J. Schroeder, A topology modifying progressive decimation algorithm, In *Proceeding of the IEEE Visualization'97*, pp. 205–213, Phoenix, Arizona, October 1997.
- [18] D.D. Sleator, R.E. Tarjan and W.P. Thurston, Rotation distance, triangulations and hyperbolic geometry, *J. Amer. Math. Soc.* 1 (1988), 647–682.
- [19] K. Wagner, Bemerkungen zum Vierfarbenproblem, *Jahresber. Deutsch. Math.-Verein.* 46 (1936), 26–32.