

Flips in Combinatorial Pointed Pseudo-Triangulations with Face Degree at most Four

Oswin Aichholzer¹, Thomas Hackl¹, David Orden², Alexander Pilz¹, Maria Saumell³, and Birgit Vogtenhuber¹

¹ Institute for Software Technology, Graz University of Technology. Inffeldgasse 16b/II, Graz, Austria, [oaich|thackl|apilz|bvogt]@ist.tugraz.at

² Departamento de Física y Matemáticas, Universidad de Alcalá. Apdo. de Correos 20, E-28871. Alcalá de Henares, Spain, david.orden@uah.es

³ Département d'Informatique, Université Libre de Bruxelles. Boulevard du Triomphe - CP 212, 1050 Bruxelles, Belgium, maria.saumell.m@gmail.com

Abstract

In this paper we consider the flip operation for combinatorial pointed pseudo-triangulations where faces have size 3 or 4, so-called *combinatorial 4-PPTs*. We show that every combinatorial 4-PPT is stretchable to a geometric pseudo-triangulation, which in general is not the case if faces may have size larger than 4. Moreover, we prove that the flip graph of combinatorial 4-PPTs is connected and has diameter $O(n^2)$, even in the case of labeled vertices.

1 Introduction

Given a graph of a certain class, a *flip* is the operation of removing one edge and inserting a different one such that the resulting graph is again of the same class. An example of such a class is the class of maximal planar (simple) graphs, also called *combinatorial triangulations*, where any combinatorial embedding (clockwise order of edges around each vertex) has only faces of size 3. Flips in combinatorial triangulations remove the common edge of two triangular faces and replace it by the edge between the two vertices not shared by the faces, provided that these two vertices were not already joined by an edge. Combinatorial triangulations have a geometric counterpart in triangulations of point sets in the plane, which are maximal plane geometric (straight-line) graphs with predefined vertex positions. In this geometric setting there is also a flip operation, for which a different restriction applies: An edge can be flipped if and only if the two adjacent triangles form a convex quadrilateral (otherwise the new edge would create a crossing).

Flips in (combinatorial) triangulations have been thoroughly studied. See the survey by Bose and Hurtado [4]. A prominent question about flips is to study the *flip graph*. This is an abstract graph whose vertices are the members of a given graph class having the same number of vertices, and in which two graphs are neighbors if and only if one can be transformed into the other by a single flip. For both triangulations and combinatorial triangulations the flip graph is connected. Lawson [9] showed that the flip graph of triangulations of a point set is connected with quadratic diameter, which was later shown to be tight [6]. For combinatorial triangulations there are actually two classes to consider: those of labeled and unlabeled graphs, where in the latter class no two distinct elements are isomorphic. For unlabeled combinatorial triangulations on n vertices Wagner [14] proved connectedness of the flip graph, and Komuro [8] showed its diameter to be $\Theta(n)$. For the labeled setting Sleator, Tarjan, and Thurston [12] showed the diameter to be $\Theta(n \log n)$.

Triangulations have a natural generalization in pseudo-triangulations. They have become a popular structure in Computational Geometry within the last two decades, with applications in, e.g., rigidity theory and motion planning. See the survey by Rote, Santos, and Streinu [4]. A *pseudo-triangle* is a simple polygon in the plane with exactly three convex vertices (i.e., vertices whose interior angle is smaller than π). A *pseudo-triangulation* \mathcal{T} of a finite point set S in the plane is a partition of the convex hull of S into pseudo-triangles such that the union of the vertices of the pseudo-triangles is exactly S . Triangulations are a particular type of pseudo-triangulations, actually the ones with the maximum number of edges. Those with the minimum number of edges are the so-called *pointed pseudo-triangulations*, in which every vertex is *pointed*, i.e., incident to a reflex angle (an angle larger than π) [13].

Flips can also be defined for the class of pseudo-triangulations of point sets in the plane. The flip graph for general pseudo-triangulations is known to be connected [2], as well as the subgraph induced by pointed pseudo-triangulations [5]. The currently best known bound on the diameter is $O(n \log n)$ for both flip graphs [2, 3], where here and for the rest of the paper n denotes the number of vertices.

In a pseudo-triangulation, the pseudo-triangles can have linear size. Hence, in contrast to triangulations, the flip operation can no longer be computed in constant time. This fact led to the consideration of pseudo-triangulations in which the size of the pseudo-triangles is bounded by a constant. Kettner et al. [7] showed that every point set admits a pointed pseudo-triangulation with face degree at most four (except, maybe, for the outer face). We call such a pointed pseudo-triangulation a *4-PPT*.

On the one hand, 4-PPTs behave nicely for problems which are hard for general pseudo-triangulations. For instance, they are always properly 3-colorable, while 3-colorability is NP-complete to decide for general pseudo-triangulations [1]. On the other hand, known properties of general pseudo-triangulations remain open for 4-PPTs. For instance, it is not known whether the flip graph of 4-PPTs is connected, even for the basic case of a triangular convex hull.

The aim of this paper is to make a step towards answering this last question, by considering the combinatorial counterpart of 4-PPTs.

A *combinatorial pseudo-triangulation* [10] is a combinatorial embedding of a planar simple graph in the plane together with an assignment of marks *reflex/convex* to its angles such that (1) every interior face has exactly three angles marked convex, (2) all the angles of the outer face are marked reflex, and (3) no vertex is incident to more than one reflex angle. (These marks of the angles are called “labels” by Orden et al. [10], we use a different term to prevent confusion with the classic labels of the vertices.)

Note that the assignment of these marks fulfills the same properties as actual reflex/convex angles in a (geometric) pseudo-triangulation. This analogy with the geometric case goes on by calling *pointed* vertices in a combinatorial pseudo-triangulation those which, indeed, are incident to one angle marked reflex. Then, *combinatorial pointed pseudo-triangulations* are those in which every vertex is pointed. Combinatorial pointed pseudo-triangulations with face degree at most four (except, maybe, for the outer face), will be called *combinatorial 4-PPTs*.

As it has been done for combinatorial triangulations, we consider flip graph connectivity of the labeled and unlabeled graph; while we allow the outer face (predefined by the combinatorial embedding in the plane) to have an arbitrary number of vertices, we require that these vertices are the same in the source and the target graph.

2 Properties

In this section, we prove some properties of combinatorial 4-PPTs and, in particular, we show that every combinatorial 4-PPT is stretchable to a geometric pseudo-triangulation.

Lemma 2.1 *Let T be a combinatorial 4-PPT and H be a subgraph of T with $|V(H)| \geq 3$. Then H has at least 3 vertices whose reflex angle is contained in the outer face of H (called “corners of first type” by Orden et al. [10]).*

Proof. W.l.o.g., we may assume that H consists of a single connected component. Let H' be the maximal subgraph of T that has the same outer face as H . Hence, if the claim holds

for H' it also holds for H , and we only need to consider inner faces of size 3 or 4. For the subgraph H' , let us denote with n the number of vertices, e the number of edges, t the number of inner faces of size 3, q the number of inner faces of size 4, b the number of boundary angles and c the number of convex boundary angles in the outer face of H' . Note that $b \geq 3$ and that $b > n$ is possible.

Let us double-count the edges. On the one hand, the number of angles equals twice the number of edges; since there are n reflex angles and $3t + 3q + c$ convex angles, we get that $2e = 3t + 3q + c + n$. On the other hand, from Euler's formula we have $e = n + t + q - 1$. Eliminating e from these two equations, we get that the number of reflex angles is $n = t + q + 2 + c$. Now we can express the number n of reflex angles as $b - c + q$, to get that $b - c = t + 2 + c$, which is at least 3 if $c > 0$. Either in this case or if $c = 0$, we get that $b - c \geq 3$, as desired. \square

Corollary 2.2 *In any combinatorial 4-PPT of the interior of a simple cycle with b vertices, of which c have the reflex angle inside the cycle, the number t of triangular faces is given by $t = b - 2c - 2$.*

A combinatorial pseudo-triangulation has the *generalized Laman property* if every subset of x non-pointed vertices and y pointed vertices, where $x + y \geq 2$, induces a subgraph with at most $3x + 2y - 3$ edges. Both this property and the number of reflex angles from Lemma 2.1 are related to the stretchability of a combinatorial pseudo-triangulation into a geometric one. A face of a combinatorial pseudo-triangulation is called *degenerate* if it contains edges which appear twice on the boundary of this face. See Figure 1 (left). Note that in our setting this is equivalent to the definition by Orden et al. [10] where a face is non-degenerate if the edges incident to it form a simple closed cycle.

Proposition 2.3 (Orden et al. [10], Corollary 2) *The following properties are equivalent for a combinatorial pseudo-triangulation G :*

1. G can be stretched to become a pseudo-triangulation.
2. G has the generalized Laman property.
3. G has no degenerate faces and every subgraph of G with at least three vertices has at least three corners of first type.

Since, by definition, combinatorial 4-PPTs have no degenerate faces, we can use Proposition 2.3 to conclude the following.

Theorem 2.4 *Every combinatorial 4-PPT can be stretched to become a 4-PPT with the given assignment of angles. Furthermore, combinatorial 4-PPTs have the generalized Laman property.*

Note that there exist non-stretchable combinatorial pointed pseudo-triangulations with faces of size at most 5. See Figure 1 (right). There and in the forthcoming figures, circular arcs denote angles marked as reflex.



Figure 1: Left: A degenerate 5-face. Right: A non-stretchable combinatorial pointed pseudo-triangulation [10].

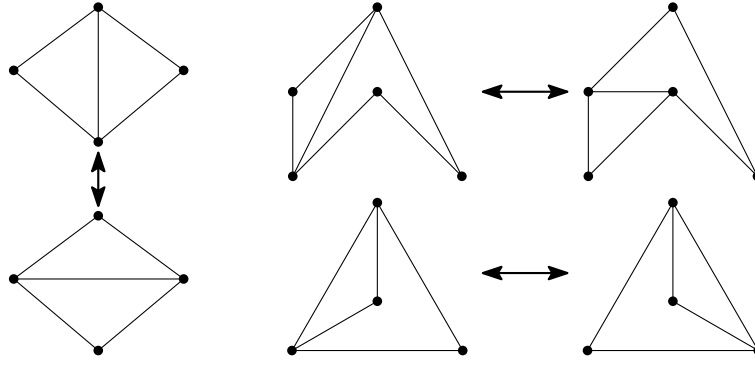


Figure 2: Geometric flip of an edge of a triangle. Left: Both faces are triangles. Right: One face is a quadrilateral and in the lower case the removal of the flipped edge gives a degenerate 5-face.

3 Flips

Before defining flips between combinatorial 4-PPTs, we make some observations about their geometric counterpart. For good visual distinction, we draw the edges of non-geometric graphs as non-straight Jordan arcs throughout this section.

In geometric 4-PPTs every edge of a triangle (except for those being part of the convex hull) is flippable [11]. Consider flipping an edge e which separates a triangle Δ from another face F in a geometric 4-PPT. If F is also a triangle, then removing e and inserting the other diagonal e' of the convex 4-face $\Delta \cup F$ is the well known “Lawson flip”. (Note that $\Delta \cup F$ has to be convex because of the pointedness of the 4-PPT.) If F is a 4-face, then the removal of e merges Δ and F into a 5-face, which might be degenerate if Δ and F share two edges. Note that this degenerate case is the only one in which Δ and F can share three vertices, as there are no multiple edges in geometric graphs. See Figure 2.

Similar to the geometric case, we consider flips of an interior edge e of an interior triangular face Δ in a combinatorial 4-PPT: Consider the face F , triangular or quadrangular, sharing e with Δ . A *flip* of e consists in replacing e by another edge e' such that (1) e' splits $(\Delta \cup F) \setminus e$ into a triangular face Δ' and a face F' , triangular or quadrangular, respectively, and (2) the result is a combinatorial 4-PPT. In particular, and in contrast to the geometric case, we have to explicitly avoid multiple edges in the combinatorial setting. Hence, we have to ensure that the edge e' that is inserted by the flip is not already contained in the combinatorial 4-PPT (as an edge outside $\Delta \cup F$). To emphasize that an exchange of two edges is a flip avoiding multiple edges, we sometimes call a flip *valid*. Further, to highlight that an exchange of two edges which would locally (inside $\Delta \cup F$) be a flip would introduce multiple edges, we call this an *invalid* flip. Recall though, that a flip is defined to be valid and we use this distinction only for emphasis in situations where we prove the existence of certain flips.

Observe that, in a (combinatorial) 4-PPT, if one face involved in a flip is triangular, then after the flip no face can have more than four vertices. Thus, we restrict ourselves to flips where at least one involved face is triangular. The following lemma shows that every interior edge of an interior triangular face can be flipped.

Lemma 3.1 *In a combinatorial 4-PPT, every edge e of an interior triangular face that is not an edge of the outer face is flippable. Furthermore: (1) If the removal of e results in a 4-face or a degenerate 5-face, then there is a unique valid flip for e . (2) If removing e results in a non-degenerate 5-face, then there are at least two valid flips for e .*

Proof. Let Δ be a triangular face and let F be the face that is separated from Δ by e . If F is also a triangular face, then $\Delta \cup F$ is a 4-face. In this case exchanging e by the unique other diagonal of $\Delta \cup F$ is a valid flip. See Figure 3 (left). If F is a face of size 4 we have to distinguish two cases. The first case is when $\Delta \cup F$ is degenerate. Then, there is only one choice of e' in order to split $\Delta \cup F \setminus e$ as required. Furthermore, the corresponding edge e' could not be already an edge, since it was not in the interior of $\Delta \cup F$ and it cannot go through

the exterior of $\Delta \cup F$ because of planarity. See Figure 3 (right). Hence, this choice is always valid.

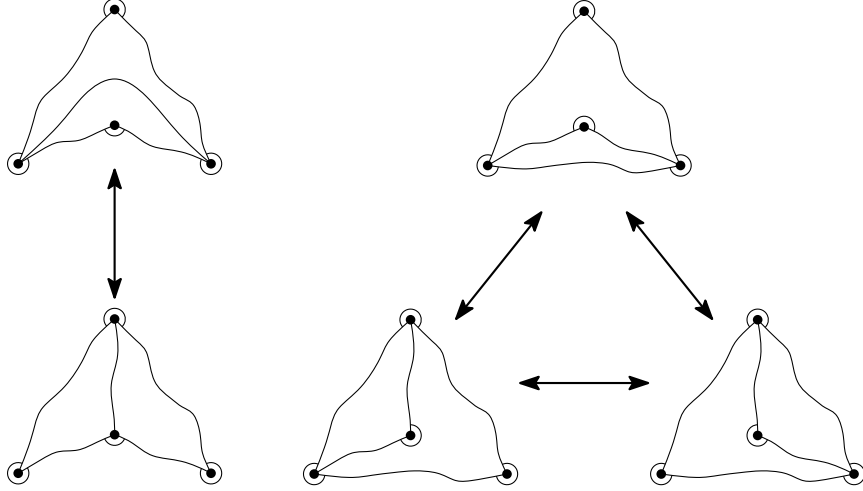


Figure 3: Combinatorial flips for an edge of a triangular face. Left: Both faces are triangular. Right: One face is a quadrangular and the 5-face is degenerate.

The second case is when $\Delta \cup F$ is non-degenerate. We show that flipping towards an edge using the reflex vertex is always valid. See Figure 5. Denote by v_1, \dots, v_5 the boundary vertices of $\Delta \cup F \setminus e$, in counterclockwise order and with v_1 being the reflex vertex. The edge e' we intend to insert is then either v_1v_3 or v_1v_4 . Let us focus on the first case, the other one being handled analogously. If $e' = v_1v_3$ is not valid, there has to be already an edge between v_1 and v_3 in the exterior of $\Delta \cup F$. But then at most two vertices of the 3-cycle $v_1v_2v_3$ have their reflex angle on the outside of that cycle, by this contradicting Lemma 2.1. See Figure 4.

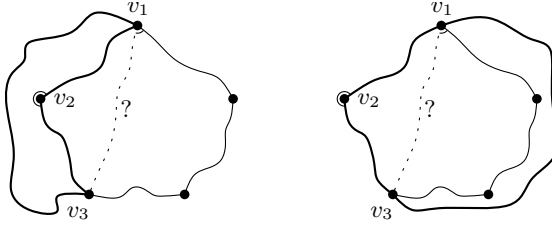


Figure 4: Flipping towards an edge incident to the reflex vertex is always valid.

It remains to prove that in the non-degenerate case there are at least two valid flips for e . Figure 5 shows the possible flips when $\Delta \cup F$ is non-degenerate, with solid arrows indicating always valid flips and dotted arrows indicating flips which might be valid or not.

If e is not incident to the reflex vertex, then there are two valid flips towards edges incident to that vertex. If e is incident to the reflex vertex, there is always a valid flip towards the other diagonal e' incident to that vertex. For a second valid flip, we show that the two remaining diagonals cannot simultaneously give invalid flips. Let the edge to flip be $e = v_1v_3$ (the other case is analogous).

In order for both v_2v_4 and v_3v_5 to give invalid flips, the combinatorial 4-PPT must have both edges v_2v_4 and v_3v_5 in the exterior of the 5-face v_1, \dots, v_5 . This is impossible since it would imply a crossing. See Figure 6. \square

Observe that, by Theorem 2.4, every combinatorial 4-PPT can be stretched into a geometric 4-PPT. Thus, the statement in Lemma 3.1 that every interior edge of a triangular face is flippable can also be seen from the geometric case. In contrast to the geometric case, where

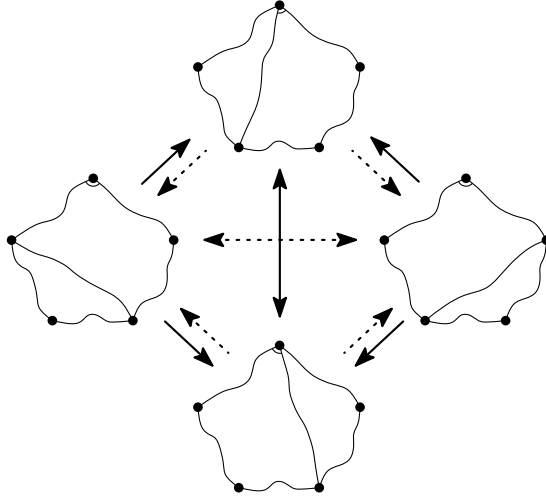


Figure 5: Combinatorial flips for an edge of a triangular face in a combinatorial 4-PPT, non-degenerate case.

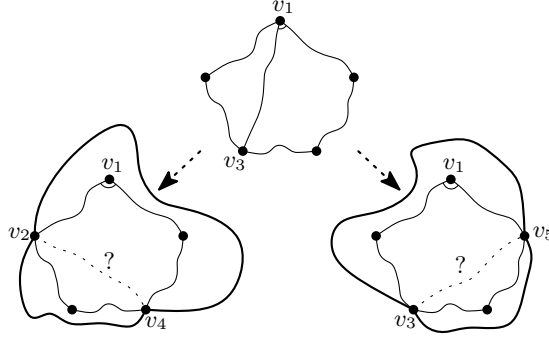


Figure 6: In a non-degenerate 5-face, the two diagonals not using the reflex vertex cannot both give invalid flips.

all valid flips are unique, the combinatorial case described in part (2) of the lemma has up to three possible flips, of which there are always at least two valid ones. This hints at another interesting observation. Given a combinatorial flip between two combinatorial 4-PPTs, by Theorem 2.4 we know that both of them can be stretched into geometric 4-PPTs with straight edges. However, it might not be possible to use the same geometric embedding for the vertices in both of them. See Figure 7 for an example where two different geometric embeddings are required.

4 Connectivity of the Flip Graph

A usual approach for proving connectivity of the flip graph and bounding its diameter is to define a special canonical graph and to show that there exists a sequence of a certain (bounded) number of flips from any graph to the canonical one. By the reversibility of the used flips, this proves connectivity of the flip graph and gives a bound on its diameter. In this section we deal with combinatorial 4-PPTs on unlabeled vertices with a fixed triangular outer face. In later sections we will extend this base case to labeled vertices (Section 5) and to the general case of combinatorial 4-PPTs with an arbitrarily sized outer face on unlabeled and labeled vertices (Section 6).

We define the unique *canonical* combinatorial 4-PPT with triangular outer face to be the

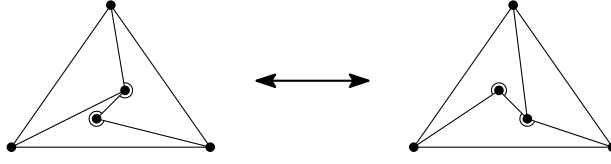


Figure 7: Combinatorial flips might need different geometric embeddings.

combinatorial 4-PPT where two of the vertices in the outer face are adjacent to all other vertices, while the third one has degree 2. An example can be found in Figure 12 (left). Observe that this canonical combinatorial 4-PPT is indeed unique as we consider unlabeled vertices for now. In the following we will show step by step how to build the flip sequence from any combinatorial 4-PPT to the canonical one. We only allow flips of interior edges of interior triangular faces, as by Lemma 3.1 these are always flippable. For a combinatorial 4-PPT with triangular outer face, Corollary 2.2 implies that there is only one interior triangular face. Hence, with the presented flip sequences we will “move” the single triangular face through the combinatorial 4-PPT.

Note that the following lemma is not restricted to a triangular outer face as the proof is (almost) the same for arbitrary sized outer faces, which will be needed also in Section 6.

Lemma 4.1 *Let T be a combinatorial 4-PPT with outer face o_1, \dots, o_h , $h \geq 3$. For any edge b of the outer face, there is a sequence of $O(n)$ flips resulting in a combinatorial 4-PPT with one interior triangular face incident to b .*

Proof. Consider the dual of T and choose a path $\Delta = F_0 \rightarrow F_1 \rightarrow F_2 \rightarrow \dots \rightarrow F_k = F_{\text{outer}}$ from an interior triangular face Δ to the outer face F_{outer} , such that F_{k-1} and F_k share the edge b . Let e_j , $1 \leq j \leq k$, be an edge separating the faces F_{j-1} and F_j in the path (note that there might be two edges shared by two faces). If there is a triangular face $F_{j'}$ ($1 \leq j' \leq k-1$) in the path, then choose the triangular face $F_{j'}$ with highest index j' among all triangular faces F_j , and replace Δ and the path in the dual of T by $F_{j'}$ and the shortest (sub)path starting from this new Δ . Note that this can only happen if $h > 3$, as with a triangular outer face there exists only one triangle.

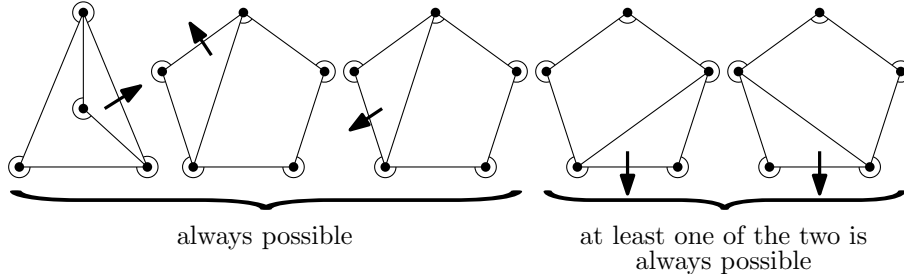


Figure 8: The region $\Delta \cup F_j$ after the flip. The arrows indicate the edge e_{j+1} to which the triangular face should be incident after the flip. Leftmost: $\Delta \cup F_j$ is degenerate. All but leftmost: The four different cases (depending on the relative position of e_{j+1} and the reflex interior angle) for non-degenerate $\Delta \cup F_j$. The two rightmost 4-PPTs show cases with two possible flips, of which at most one can lead to a double edge.

We define the sequence of flips in such a way that, after the j -th flip, Δ is incident to e_{j+1} , which then separates Δ from F_{j+1} . Thus, after $k-1 = O(n)$ flips Δ will be incident to $F_k = F_{\text{outer}}$ through $e_k = b$, as required.

At the j -th flip we consider the region $\Delta \cup F_j$ and we have to replace an edge e shared by Δ and F_j with a valid edge e' , such that the new triangular face Δ' is incident to e_{j+1} . If $\Delta \cup F_j$ is degenerate, then there are two edges shared by Δ and F_j . For each of these edges there exists a unique valid flip, by Lemma 3.1. Flipping the edge which does not share a vertex

with e_{j+1} yields the desired result. See Figure 8 (left). If $\Delta \cup F_j$ is non-degenerate, then the single edge shared by Δ and F_j is e_j . By Lemma 3.1 there exist two valid flips for e_j . At least one of these flips introduces an edge e' , such that the new triangular face Δ' is incident to e_{j+1} . See Figure 8 (all but leftmost). \square

Once the interior triangular face is incident to an edge b of the outer face, the next step will be flipping away interior edges incident to one endpoint of b .

Lemma 4.2 *Given a combinatorial 4-PPT with triangular outer face in which the interior triangular face Δ is incident to the edge b of the outer face, there is a sequence of flips resulting in a combinatorial 4-PPT in which the endpoint t of $b = rt$ has no interior incident edges.*

Proof. We describe a flip sequence that removes all inner edges incident to t . This flip sequence can be partitioned into two phases and some cases. Let the vertices neighbored to the vertex t be ordered radially around t , starting with r . In each case, let the vertices in that order be $r = w_0, \dots, w_k$.

Phase 1: During this phase, the inner triangular face Δ has rt as a side, i.e., $\Delta = trw_1$. We distinguish three different cases:

Case 1: tw_1 is the only inner edge incident to t , i.e., $k = 2$. If Δ is incident to only one 4-face F (i.e., $\Delta \cup F$ is degenerate), we can flip the edge tw_1 and we are done. Otherwise, let the 4-face F incident to tw_1 be tw_1uw_2 . See Figure 9. The reflex angle inside F is either at u or w_1 . If it is at u , we flip tw_1 to w_0u , obtaining the 4-face tw_0uw_2 . Otherwise, the reflex angle is at w_1 and we flip tw_1 to w_1w_2 , obtaining the 4-face $tw_0w_1w_2$. Either way, the degree of t is 2 and we are done.

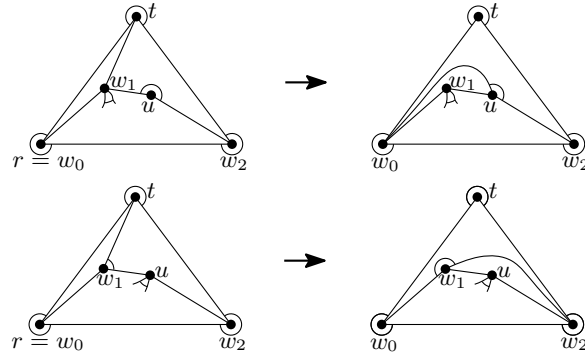


Figure 9: Phase 1, Case 1: Only one interior edge is incident to t .

Case 2: at least two inner edges are incident to t and there does not exist an edge w_0w_2 . See Figure 10. Since the reflex angle of t is at the outer face we can replace the edge tw_1 by w_0w_2 . This reduces the degree of t by one. The inner triangular face is again adjacent to w_0t , and we remain in Phase 1.

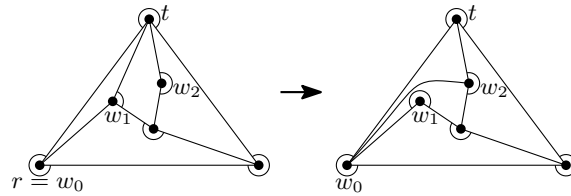


Figure 10: Phase 1, Case 2: Several interior edges are incident to t and w_0w_2 does not exist.

Case 3: at least two inner edges are incident to t and there exists an edge w_0w_2 . See Figure 11. If the two inner edges of Δ are incident to a single 4-face, we have a degenerate case; we flip the edge w_0w_1 to w_1w_2 , making tw_1w_2 the inner triangular face. Otherwise, let the 4-face F incident to tw_1 be tw_1uw_2 ; we flip tw_1 to tu (this is possible since if tu already existed, it would have to cross the cycle rw_2uw_1). Either way, the flip does not reduce the

degree of t , but the inner triangular face is now inside the 3-cycle tw_0w_2 . We switch to Phase 2.

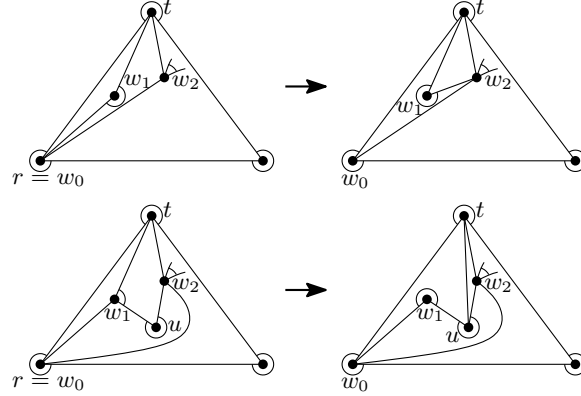


Figure 11: Phase 1, Case 3: The possible transitions to Phase 2.

Phase 2: During this phase, the inner triangular face is tw_1w_2 , and w_1 stays fixed for the whole phase. Further, we know that w_1 was enclosed by a 3-cycle (at the transition to this phase), which implies that there are no edges from w_1 to w_j for any $j \geq 2$. We decrease the degree of t in the following manner.

Case 1: there is a 4-face F incident to tw_2 . There cannot be an edge w_1w_3 since w_1 was enclosed by a 3-cycle. Further, the reflex angle of F is not at t . Hence, we can flip tw_2 to w_1w_3 , which reduces the degree of t and we remain in Phase 2, with tw_1w_3 being the new inner triangular face.

Case 2: there is no 4-face incident to tw_2 , i.e., $k = 2$. This case is symmetric to Case 1 of Phase 1. The edge tw_1 is flipped in one of the two described ways, reducing the degree of t to 2 and thus ending the process. \square

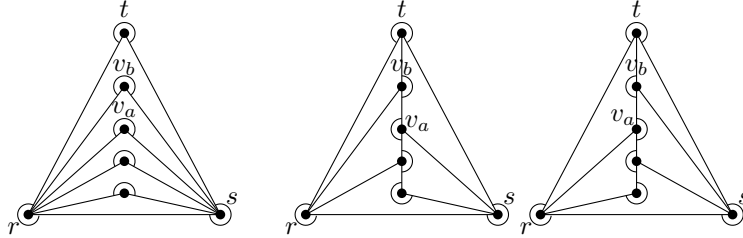


Figure 12: Left: The canonical combinatorial 4-PPT T , with outer face rst and base edge rs . Middle and right: The two spinal combinatorial 4-PPTs (from T) with tip t , s -spinal (middle) and r -spinal (right).

Theorem 4.3 *The graph of flips in combinatorial 4-PPTs with n vertices and triangular outer face is connected and has diameter $O(n^2)$.*

Proof. Given such a combinatorial 4-PPT, follow the steps in Lemmas 4.1 and 4.2, then use induction for the combinatorial 4-PPT obtained by removing t . This leads to the unique *canonical* combinatorial 4-PPT with triangular outer face, where two of the vertices in the outer face are adjacent to all other vertices, while the third one has degree 2. See Figure 12 (left).

Furthermore, the number of flips needed in Lemmas 4.1 and 4.2 is at most linear in the number of vertices of the combinatorial 4-PPT. \square

5 Connectivity for Labeled Combinatorial 4-PPTs

The canonical graph produced in the proof of Theorem 4.3 does not care about the order of the interior vertices with respect to the extremal vertices. When labeling the vertices of both the source and target graph accordingly, the two graphs produced are isomorphic and have the same combinatorial embedding, but might not be equivalent when arbitrary predefined labels are considered. In this section we describe how to flip between canonical combinatorial 4-PPTs with triangular outer face under consideration of the labels.

The canonical combinatorial 4-PPT, as exemplified in Figure 12 (left), induces a total order on the interior vertices (i.e., vertices that are not incident to the outer face) by inclusion of vertices in 3-cycles formed by one interior vertex and the two vertices r and s of high degree. Thus, given a canonical combinatorial 4-PPT, we say that an interior vertex v_b is *above* another interior vertex v_a (and v_a is *below* v_b) if and only if the 3-cycle defined by v_b (and r and s) contains v_a in its interior. Further, two vertices v_a and v_b are *neighbored* when they are neighbored in the total order. Let v_1, \dots, v_i be the $i = n - 3$ interior vertices in that order.

Besides the canonical form, a second special class of combinatorial 4-PPTs that we will use is the one of *spinal* combinatorial 4-PPTs, see Figure 12 (middle and right). In a spinal combinatorial 4-PPT the subgraph on $\{v_1, \dots, v_i\} \cup \{t\}$ is a path with t and v_1 as the end vertices. This path is called the *spine*. Further, $\{v_1, \dots, v_i\}$ are alternatingly (in the order on the spine) connected to r and s to complete a combinatorial 4-PPT. The reflex angle at v_k , $2 \leq k \leq i$ is the angle between the two edges of the spine (incident to v_k). The reflex angle at v_1 is inside the face defined by r, s, v_1 , and v_2 (if $i = 1$, then $v_2 = t$). Depending on whether v_1 is connected to r or s we distinguish between an r -spinal or s -spinal combinatorial 4-PPT, respectively.

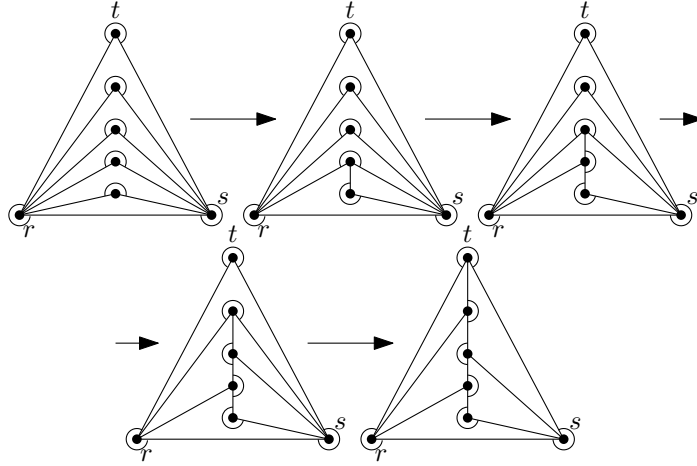


Figure 13: Flipping from the canonical combinatorial 4-PPT to the s -spinal combinatorial 4-PPT.

Observe that there exists a simple sequence of i flips to transform a canonical combinatorial 4-PPT to the r - or s -spinal combinatorial 4-PPT. See Figure 13 for an example of flipping to the s -spinal combinatorial 4-PPT. Flipping to the r -spinal combinatorial 4-PPT is analogous, but flipping the edge sv_1 in the first step. It is easy to see that the total order in the canonical combinatorial 4-PPT is the same as the one on the spine, for the two spinal combinatorial 4-PPTs: Two vertices are neighbored on the spine if and only if they are neighbored in the total order of the canonical combinatorial 4-PPT.

Observation 1 *For a triangular outer face rst and i interior vertices, flipping from a canonical combinatorial 4-PPT with base edge rs to an r - or s -spinal combinatorial 4-PPT can be done in i flips. The order in the canonical combinatorial 4-PPT equals the order on the spine.*

Let T be some canonical combinatorial 4-PPT with triangular outer face rst , base edge rs , and let the i interior vertices $\{v_1, \dots, v_i\}$ be labeled. If $i < 2$ then reordering of the interior

vertices is not necessary. For reordering the $i \geq 2$ labeled interior vertices we need to be able to exchange two neighbored labeled vertices (for which we will use the spine). We call the required sequence of flips a *swap*. For swapping two labeled vertices v_k and v_{k+1} in a spinal combinatorial 4-PPT (and thus also in a canonical combinatorial 4-PPT) we need to distinguish the three cases $k = 1$, $k = 2$, and $3 \leq k < i$. If $k \leq i - 2$, let $t' = v_{k+2}$; otherwise, let $t' = t$. For all cases we consider the subset $\{v_1, \dots, v_{k+1}\} \cup \{r\} \cup \{s\} \cup \{t'\}$. Further, we assume that we have already flipped to the spinal combinatorial 4-PPT with outer face rst' . Note that the subgraph in rst' is a spinal combinatorial 4-PPT. (By Observation 1, flipping from a canonical combinatorial 4-PPT to this situation takes $k + 1$ flips.)

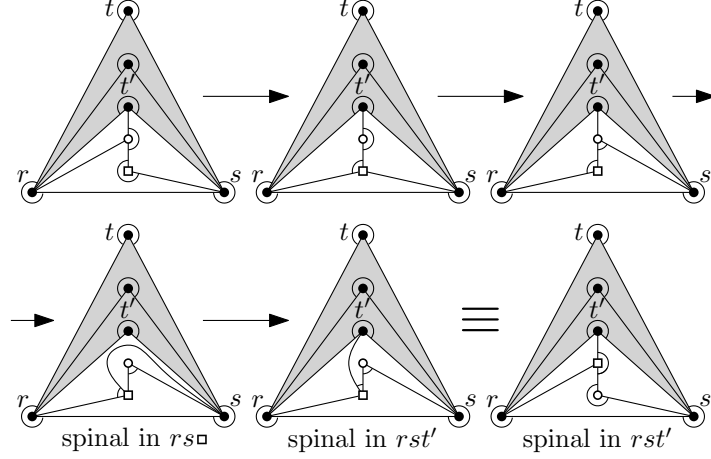


Figure 14: Swapping the position of v_1 and v_2 in a combinatorial 4-PPT that is spinal in rst' .

The flip sequence for the case $k = 1$ is depicted in Figure 14. The two vertices v_1 and v_2 are shown as a white square and a white dot, respectively, which depict the different labels of the vertices. After three flips we can reach a combinatorial 4-PPT that is spinal in rsv_2 , where the labeled vertices v_1 and v_2 have been swapped (Figure 14, second row, left). Recall that the numbering of the vertices v_k , $1 \leq k \leq i$, is defined by their position along the spine; e.g., v_1 is always the first vertex on the spine. Observe that this combinatorial 4-PPT is only one additional flip away from the canonical combinatorial 4-PPT (with swapped labels). Applying an additional flip, we can reach a combinatorial 4-PPT that is spinal in rst' , where the labeled vertices v_1 and v_2 have been swapped (Figure 14, second row, middle and right).

The flip sequences for the cases $k = 2$ and $3 \leq k < i$ are very similar. In fact, the case $k = 2$ is just a special case of the case $3 \leq k < i$. For completeness, the flip sequence for the case $k = 2$ is given in Figure 15.

Figure 16 exemplifies the flip sequence for swapping the neighbored vertices v_k and v_{k+1} for $3 \leq k \leq i - 1$. In the example $i = 5$ and the labeled vertices v_3 and v_4 should be swapped, i.e., $k = 3$. For larger values of i and k the remaining interior vertices will be placed in the interior of $rt'st$ (and rsv_1v_{k-1}). These areas (shown gray in Figure 14, 15, and 16) remain untouched throughout the whole swap operation. In this sense, all described swap operations (sequences of flips) are local.

Altogether, the presented flip sequences allow to reorder the labeled interior vertices in a canonical combinatorial 4-PPT with $O(n^2)$ flips.

Theorem 5.1 *The flip graph of labeled combinatorial 4-PPTs with n vertices and triangular outer face is connected with diameter $O(n^2)$.*

Proof. Let T_1 and T_2 be any two combinatorial 4-PPTs with n vertices and triangular outer face. By Theorem 4.3, flipping both T_1 and T_2 to a canonical combinatorial 4-PPT T'_1 and T'_2 , respectively, takes $O(n^2)$ flips. Except for the order of the labeled interior vertices, T'_1 and T'_2 are identical. Thus, to flip from T_1 to T_2 we need to reorder the labeled interior vertices of T'_1 to match their order in T'_2 and then reverse the flip sequence from T_2 to T'_2 . This reordering

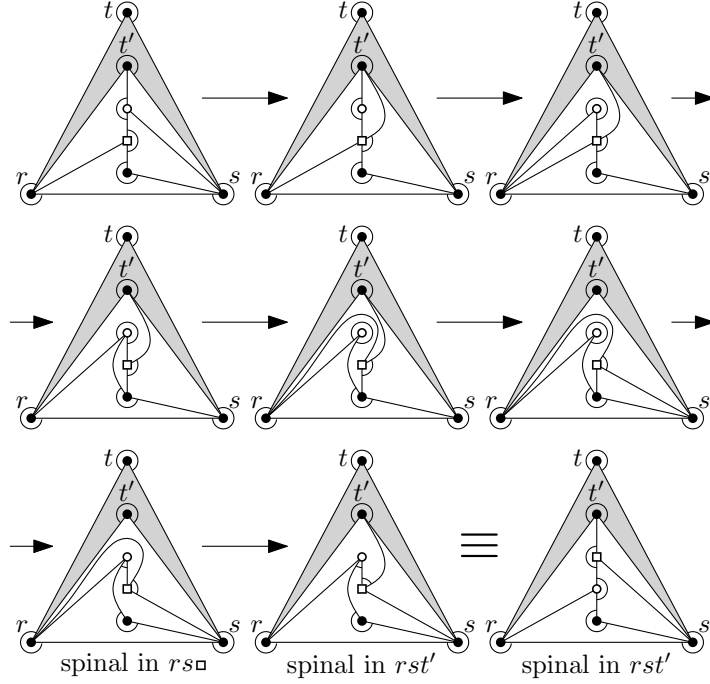


Figure 15: Swapping the position of v_2 and v_3 in a combinatorial 4-PPT that is spinal in rst' .

step can also be done in $O(n^2)$ flips, by using a sorting algorithm based on exchanging pairs of neighbors.

We use a “bubble sort” type algorithm: We flip the canonical combinatorial 4-PPT to the spinal combinatorial 4-PPT, i.e., a combinatorial 4-PPT that is spinal in rst . By Observation 1 this transformation takes $O(n)$ flips. We use the swap operation to exchange the two top inner vertices v_i and v_{i-1} if needed. Applying the swap operation results in a combinatorial 4-PPT that is spinal in rsv_i . If we do not need to exchange v_i and v_{i-1} , we need one flip to get a combinatorial 4-PPT that is spinal in rsv_i . Like in bubble sort, we continue to compare and possibly exchange the next pair of neighbored vertices until v_2 and v_1 have been processed. After every step k , k from $i-1$ down to 1, the two neighbored vertices v_{k+1} and v_k have been exchanged if needed and the resulting combinatorial 4-PPT is spinal in rsv_{k+1} .

After one such pass (from $k = i-1$ to $k = 1$) the vertex v_1 is at its final position (according to its label). Moreover, the combinatorial 4-PPT is just one flip away from the canonical one. Further, as in one pass we move from top to bottom on the spine, each swap operation needs only $O(1)$ flips (cf. Figure 14, 15, and 16). Hence, one pass needs $O(n)$ flips. It is easy to see that after $i-1 = O(n)$ such passes every labeled vertex has been moved to its required position. Therefore, with $O(n^2)$ flips we can reorder the labeled vertices in T'_1 to match their order in T'_2 . \square

6 Connectivity for Combinatorial 4-PPTs with Outer Face of Arbitrary Size and Labeled Vertices

So far we have proved connectivity of combinatorial 4-PPTs with the outer face restricted to be of size three. In this section we drop this restriction and allow outer faces of arbitrary size. We prove that for this general case the graph of combinatorial 4-PPTs stays connected, even for labeled interior vertices. The case with a triangular outer face will be a key ingredient for this proof. To this end we define a general canonical combinatorial 4-PPT and show how to reach it with $O(n^2)$ flips.

Let T be a combinatorial 4-PPT with h vertices on the outer face and i interior vertices ($n = h + i$). Let the h vertices on the outer face be o_1, \dots, o_h in counter-clockwise order. Then

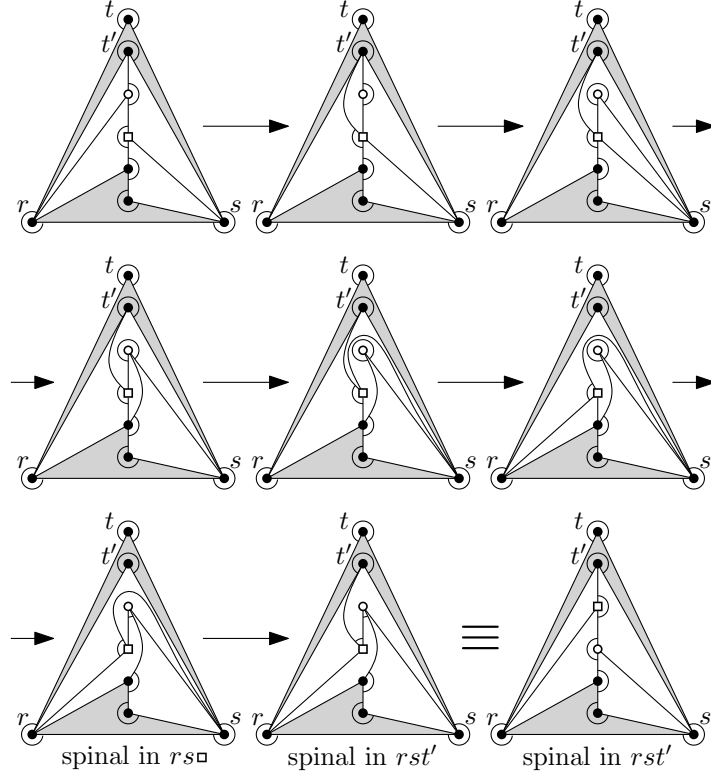


Figure 16: Swapping the position of v_k and v_{k+1} , $3 \leq k \leq i-1$, in a combinatorial 4-PPT that is spinal in rst' .

the general canonical combinatorial 4-PPT (for T) consists of a triangulation on o_1, \dots, o_h with diagonals $o_1 o_k$, $3 \leq k \leq h-1$, (a so-called *fan* at o_1) and a canonical combinatorial 4-PPT on all i interior vertices with triangular outer face $o_1 o_2 o_3$ and $o_1 o_2$ as the base edge. See Figure 17 for an example.

The flip sequence to obtain the general canonical combinatorial 4-PPT consists of three steps. Step 1): flipping to a combinatorial 4-PPT inducing the fan at o_1 ; Step 2): flipping to a canonical combinatorial 4-PPT inside each 3-cycle of that fan; and Step 3): moving all interior vertices into the 3-cycle $o_1 o_2 o_3$ such that $o_1 o_2 o_3$ is the “outer face” of a canonical combinatorial 4-PPT with base edge $o_1 o_2$. In the following we will present each step in detail and we will show that overall $O(n^2)$ flips are sufficient.

6.1 Step 1)

In a nutshell, we want to introduce one diagonal of the outer face after another, each time “cutting an ear”. In more detail, we first introduce the diagonal $o_1 o_{h-1}$ of the outer face o_1, \dots, o_h to cut off the ear $o_1 o_{h-1} o_h$. After that we will be ready to forget about the 3-cycle $o_1 o_{h-1} o_h$ for the moment, no matter whether its interior is empty of vertices or not. Then we recurse on the smaller outer face o_1, \dots, o_{h-1} . This way we get a combinatorial 4-PPT containing the diagonals $o_1 o_k$, $3 \leq k \leq h-1$, in its edge set.

It remains to show how to cut an ear in a combinatorial 4-PPT T if the diagonal $o_1 o_{h-1}$ does not already exist in T . This is very similar to the approach in Section 4. By Lemma 4.1 it is always possible to move a triangular face to an arbitrary edge of the outer face. Thus we can ensure that a triangular face Δ is incident to the edge $o_1 o_h$. If Δ is also incident to the edge $o_h o_{h-1}$, then we can cut off the ear $o_1 o_{h-1} o_h$ and iterate on the combinatorial 4-PPT with outer face o_1, \dots, o_{h-1} (note that this part contains $h-3$ triangular faces). Otherwise, we flip away all edges incident to o_h between $o_1 o_h$ and $o_h o_{h-1}$ (inside the area not containing the reflex angle) until we can introduce the diagonal $o_1 o_{h-1}$. We explain this process in the

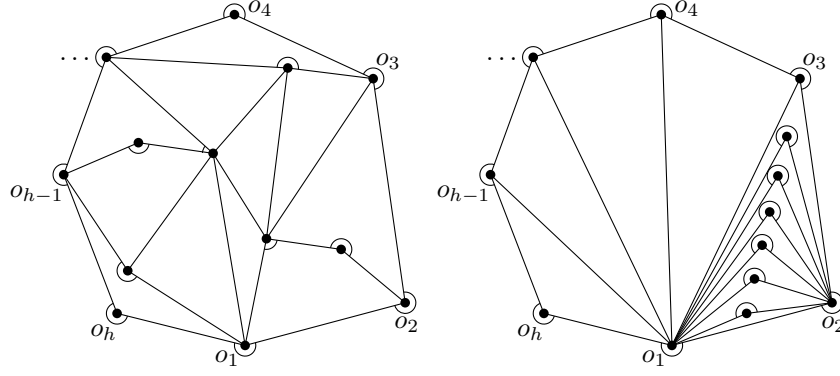


Figure 17: Left: An example combinatorial 4-PPT T . Right: The corresponding general canonical combinatorial 4-PPT.

next lemma, which is similar to Lemma 4.2 but differs at the end of the sequence.

Lemma 6.1 *Let T be a combinatorial 4-PPT with outer face o_1, \dots, o_h , $h \geq 4$, in which an interior triangular face Δ is incident to the edge $o_1 o_h$ and the diagonal $o_1 o_{h-1}$ is not an edge of T . There exists a sequence of $O(n)$ flips resulting in a combinatorial 4-PPT with $o_1 o_{h-1} o_h$ as a triangular face.*

Proof. Let e_0, \dots, e_{k+1} be the $k+2$ edges of T incident to o_h in the order of incidence, such that $e_0 = o_1 o_h$ and $e_{k+1} = o_h o_{h-1}$. Let $\Delta = F_0 \rightarrow F_1 \rightarrow F_2 \rightarrow \dots \rightarrow F_k$ be the path of faces (incident to o_h) such that e_j , $1 \leq j \leq k$, is an edge shared by the faces F_{j-1} and F_j . Note that F_k is incident to $o_h o_{h-1}$. With w_j we denote the end vertex of e_j that is not o_h .

We define a sequence of k flips. With the j -th flip we want to introduce a valid edge e' , such that a triangular face is incident to both e_{j+1} and o_1 . However, this is not always possible as the insertion of the edge e' might create a double edge. In this case, the triangular face is incident to e_{j+1} and a vertex named o'_1 , which will be used instead of o_1 for the remainder of the sequence of k flips. (See Cases 1 and 2 below.) Fortunately, this can happen only once and we will carefully distinguish the different cases. At the j -th flip, let Δ_j be the interior triangular face incident to o_h and o_1 (or o'_1). Consider the region $\Delta_j \cup F_j$.

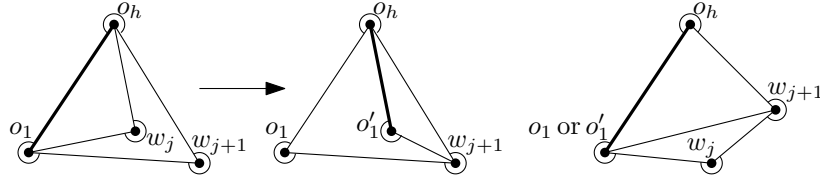


Figure 18: Left (before the flip) and middle (after the flip) show the flip in the degenerate case. Right shows the simple case of a flip between two triangular faces.

Case 1: $\Delta_j \cup F_j$ is a degenerate 5-face. A flip resulting in a new triangular face that is incident to both e_{j+1} and $o_1 o_h$ is clearly not possible. In this case, the vertex w_j is renamed o'_1 . For the edge $o_1 o'_1$ there exists a unique valid flip, by Lemma 3.1. This results in a new triangular face that is incident to e_{j+1} and $o'_1 o_h$. Observe, that o'_1 is *safe*, in the sense that for each future flip, resulting in an edge e' incident to o'_1 , e' will not create a double edge. Thus, this case can only occur once. See Figure 18 (left and middle).

Case 2: $\Delta_j \cup F_j$ is a non-degenerate 5-face and $o_1 w_{j+1}$ is already an edge of the combinatorial 4-PPT. See Figure 19 (a) and (b) for the two subcases, which differ in the two possible positions of the reflex angle inside $\Delta_j \cup F_j$. In both subcases we denote as o'_1 the vertex of $\Delta_j \cup F_j$ that is neither o_1 , o_h , w_j , nor w_{j+1} . Exchanging the edge e_j with the edge $o'_1 o_h$ is a valid flip. (See the proof of Lemma 3.1 for the non-degenerate case and use the fact that $o_1 w_{j+1}$ is assumed to be already an edge of the combinatorial 4-PPT.) Note that o'_1 is

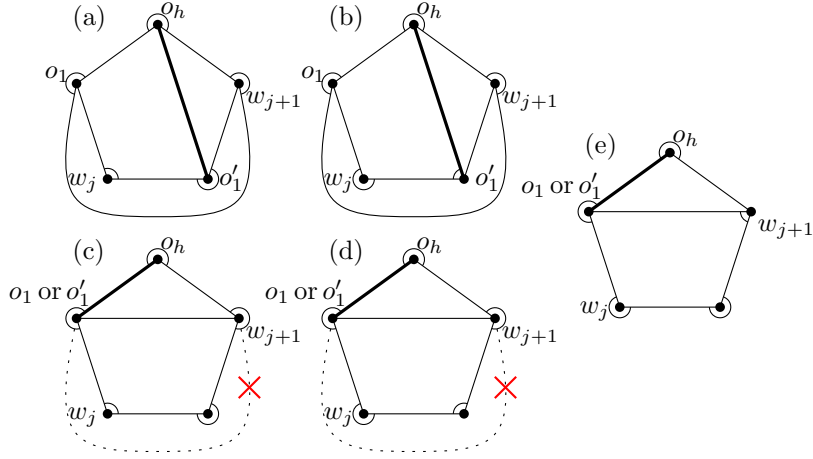


Figure 19: The five subcases of the non-degenerate case, depending on the relative position of the reflex interior angle. In (a) and (b) edge o_1w_{j+1} is assumed to be an edge of the combinatorial 4-PPT before the flip. The cases (c) and (d) are the same as (a) and (b), respectively, except that it is assumed that o_1w_{j+1} is not an edge of the combinatorial 4-PPT before the flip. The case (e) is a flip inserting an edge incident to the reflex interior angle, which is always valid.

also safe in this case. Therefore, Case 2 as well as Case 1 cannot occur as soon as the edge o'_1o_h has been introduced. That is, out of these two cases only one can occur at all and only at most once in total during the whole flip sequence.

Case 3: $\Delta_j \cup F_j$ is a non-degenerate 5-face and o_1w_{j+1} (or o'_1w_{j+1}) is not an edge of the combinatorial 4-PPT. See Figure 19 (c), (d), and (e) for the three subcases, which differ in the position of the reflex angle inside $\Delta_j \cup F_j$. In all three subcases there exists a valid flip introducing o_1w_{j+1} , such that the new triangular face is incident to both e_{j+1} and o_1o_h (or o'_1o_h).

Case 4: $\Delta_j \cup F_j$ is a quadrangular face. Hence, F_j is a triangular face. In this case, it is easy to see that there exists a valid flip introducing o_1w_{j+1} , such that one of the new triangular faces is incident to both e_{j+1} and o_1o_h (or o'_1o_h). See Figure 18 (right).

After k flips the sequence ends with a triangular face Δ' incident to $o_h o_{h-1}$ and either $o_1 o_h$ or $o'_1 o_h$. In the former case the resulting combinatorial 4-PPT has $o_1 o_{h-1} o_h$ as a triangular face, as required. In the latter case we need one more flip. The edge $o'_1 o_h$ separates Δ' from a face that is incident to $o_1 o_h$. We assumed the diagonal $o_1 o_{h-1}$ not to be an edge of the combinatorial 4-PPT and $o_1 o_{h-1}$ was also not introduced during the sequence of k flips. Therefore, replacing $o'_1 o_h$ by $o_1 o_{h-1}$ is a valid flip.

All in all, the sequence consists of k flips, one flip per edge e_j , plus possibly one additional flip if the edge $o'_1 o_h$ has been introduced during the flip sequence. Hence, $O(n)$ flips are sufficient. \square

Using Lemma 6.1, we can flip to a combinatorial 4-PPT that contains the diagonals of the fan at o_1 , by iteratively introducing the diagonals $o_1 o_j$, from $j = h - 1$ down to $j = 3$, whenever this diagonal is not already present.

Lemma 6.2 *Given a combinatorial 4-PPT with outer face o_1, \dots, o_h , $h \geq 4$, there exists a sequence of $O(n^2)$ flips resulting in a combinatorial 4-PPT, T , with outer face o_1, \dots, o_h , such that the diagonals (of the outer face) $o_1 o_j$, $3 \leq j \leq h - 1$, are in the set of edges of T .*

6.2 Step 2)

Let Δ_j be the 3-cycle $o_1 o_j o_{j+1}$, $2 \leq j \leq h - 1$. After Step 1) the edges of these 3-cycles are edges of the combinatorial 4-PPT. Let i_j be the number of interior vertices inside Δ_j . By

Theorem 4.3, $O(i_j^2)$ flips are sufficient to flip to the canonical combinatorial 4-PPT with outer face \triangle_j . Therefore, overall $O(n^2)$ flips are sufficient to flip to the canonical combinatorial 4-PPT inside each 3-cycle of the fan.

Lemma 6.3 *Let T be a combinatorial 4-PPT with outer face o_1, \dots, o_h , $h \geq 4$, such that the diagonals (of the outer face) $o_1 o_j$, $3 \leq j \leq h-1$, are in the set of edges of T . There exists a sequence of $O(n^2)$ flips resulting in a combinatorial 4-PPT, T' , with outer face o_1, \dots, o_h , such that 1) the diagonals (of the outer face) $o_1 o_j$, $3 \leq j \leq h-1$, are in the set of edges of T' , and 2) the subgraph of T' inside any three-cycle $o_1 o_j o_{j+1}$, $2 \leq j \leq h-1$, is a canonical combinatorial 4-PPT.*

6.3 Step 3)

After Step 2) the combinatorial 4-PPT with outer face o_1, \dots, o_h , $h \geq 4$, contains the diagonals $o_1 o_j$, $3 \leq j \leq h-1$, in its edge set and the subgraph inside $o_1 o_j o_{j+1}$, $2 \leq j \leq h-1$, is a canonical combinatorial 4-PPT with outer face $o_1 o_j o_{j+1}$. So it remains to move all interior vertices into $o_1 o_2 o_3$.

First consider two induced neighbored three-cycles $C_j = o_1 o_j o_{j+1}$ and $C_{j+1} = o_1 o_{j+1} o_{j+2}$, $2 \leq j \leq h-2$. Let $e = o_1 o_{j+1}$ be the diagonal that separates C_j and C_{j+1} . Assume that the canonical combinatorial 4-PPTs in C_j and C_{j+1} have both e as the base edge.

We want to move all interior vertices of C_{j+1} into C_j . It is easy to see that we can flip e (as e separates two triangular faces). This results in a combinatorial 4-PPT in $C_j \cup C_{j+1}$, exemplified in Figure 20 (a). To move one vertex from C_{j+1} to C_j there exists a very simple sequence of two flips, see Figure 20 (b) and (c). Repeatedly applying this sequence, until all vertices from C_{j+1} are moved, results in a combinatorial 4-PPT as the one exemplified in Figure 20 (e). We apply one more flip to reintroduce the diagonal e , and all vertices interior to C_{j+1} have been moved to C_j .

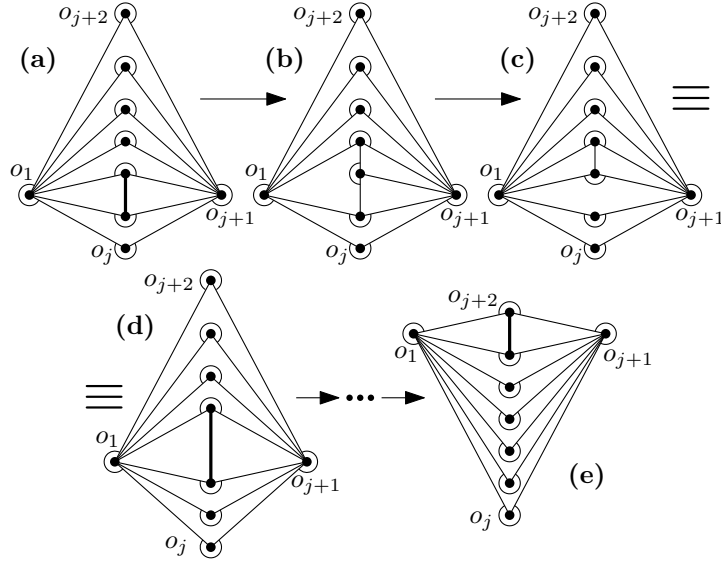


Figure 20: Moving the interior vertices between two neighbored three-cycles. (c) and (d) show different drawings of the same graph. Comparing (a) and (d), one vertex has been moved “down”.

Moving all interior vertices from C_{j+1} to C_j takes $O(i_{j+1})$ flips, with i_{j+1} being the number of vertices interior to C_{j+1} . For moving interior vertices between neighbored three-cycles we assumed that the canonical combinatorial 4-PPTs in C_j and C_{j+1} have both the same base edge. To fulfill this precondition we need a sequence of flips to rotate a canonical combinatorial 4-PPT, with triangular outer face and i interior vertices, in a linear number of flips.

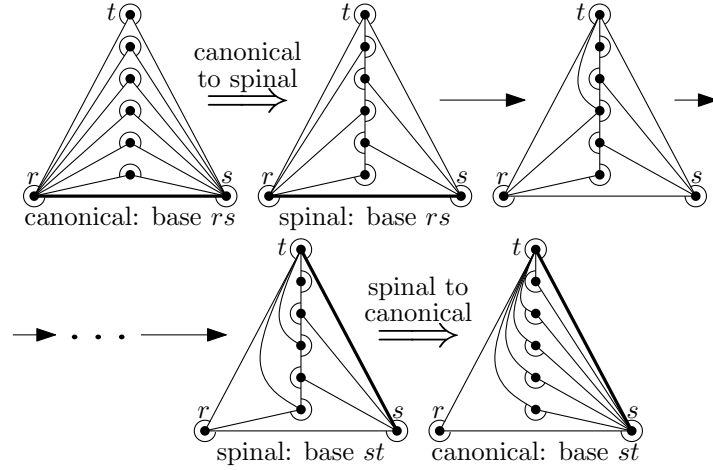


Figure 21: Flip sequence for “rotating” a canonical combinatorial 4-PPT with triangular outer face and i interior vertices. The bold segments indicate the base edge of the canonical and spinal 4-PPT, respectively.

This sequence consists of: (1) i flips to obtain the spinal combinatorial 4-PPT (by Observation 1); then (2) “rotating” the spine with $\lfloor \frac{i-1}{2} \rfloor$ flips by flipping every other non-spinal interior edge, as depicted in Figure 21 (flips indicated with single arrows); and (3) i flips from spinal back to the canonical combinatorial 4-PPT with new base edge. In Figure 21 the whole sequence of flips is exemplified for i being odd.

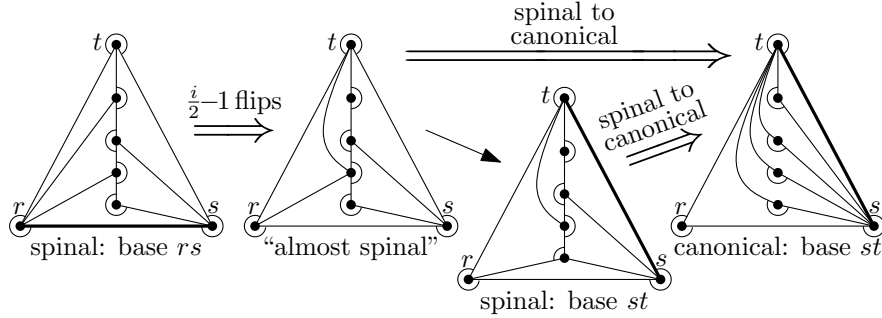


Figure 22: Flip sequence for “rotating the spine” for an even number of i interior vertices. The bold segments indicate the base edge of the canonical and spinal 4-PPT, respectively. The sequence of the first $\frac{i}{2} - 1$ flips is the same as for i is odd. The last flip to the rotated spinal 4-PPT can be eliminated, as the flip sequence from spinal to canonical can also be started from the graph named “almost spinal”.

For i being even, the flip sequence for “rotating the spine” is exemplified in Figure 22. The first part of the sequence equals that for i being odd. Note that we can eliminate the last flip to the rotated spinal combinatorial 4-PPT. The edge introduced by this last flip is the first edge that is removed in the flip sequence to the rotated canonical combinatorial 4-PPT.

Observe that we need to start with different spinal combinatorial 4-PPTs, depending on the new base edge and the parity of i . If we want to rotate the canonical combinatorial 4-PPT with outer face rst and base edge rs to the one with base edge st , then for i being odd the sequence starts with flipping to the r -spinal combinatorial 4-PPT, and for i being even the sequence starts with flipping to the s -spinal combinatorial 4-PPT. See again Figures 22 and 21.

Starting with $j = h - 2$ and stopping at $j = 2$, iteratively rotating the canonical combinatorial 4-PPTs of the two neighbored three-cycles $C_j = o_1 o_j o_{j+1}$ and $C_{j+1} = o_1 o_{j+1} o_{j+2}$, and moving the interior vertices from C_{j+1} to C_j , results in a combinatorial 4-PPT with all interior vertices inside the three-cycle $o_1 o_2 o_3$. As the number of flips for the sequences of both rotating

a canonical combinatorial 4-PPT (with triangular outer face) and moving the interior vertices is linear in the number of interior vertices, the overall sequence consists of $O(n^2)$ flips.

Finally, rotating the canonical combinatorial 4-PPT in $C_2 = o_1 o_2 o_3$ to the base edge $o_1 o_2$ can be done in another $O(n)$ flips. This results in a (general) canonical combinatorial 4-PPT with outer face o_1, \dots, o_h , $h \geq 4$.

Lemma 6.4 *Let T be a combinatorial 4-PPT with outer face o_1, \dots, o_h , $h \geq 4$, such that (1) the diagonals (of the outer face) $o_1 o_j$, $3 \leq j \leq h-1$, are in the set of edges of T , and (2) the subgraph of T inside a three-cycle $o_1 o_j o_{j+1}$, $2 \leq j \leq h-1$, is a canonical combinatorial 4-PPT. There exists a sequence of $O(n^2)$ flips resulting in a (general) canonical combinatorial 4-PPT with outer face o_1, \dots, o_h .*

6.4 General Connectivity

Summarizing over the presented three steps and Section 5 about labeled vertices, we can prove the following theorem.

Theorem 6.5 *The graph of flips in combinatorial 4-PPTs with n vertices, $h \geq 3$ of them on the outer face, is connected with diameter $O(n^2)$. This is still true for labeled vertices.*

Proof. Let T_1 and T_2 be two combinatorial 4-PPTs with n vertices, $h \geq 3$ of them on the outer face. Following the three steps summarized in Lemmas 6.2, 6.3, and 6.4 results in a sequence of $O(n^2)$ flips leading to the canonical combinatorial 4-PPTs T'_1 and T'_2 , respectively, with outer face o_1, \dots, o_h (see Figure 17).

In the unlabeled case $T'_1 = T'_2$. As all used flips are invertible this proves that the flip graph is connected with diameter $O(n^2)$.

In the case of labeled interior vertices, we can flip from T'_1 to T'_2 with $O(n^2)$ flips, by Theorem 5.1. Hence, the flip graph is connected with diameter $O(n^2)$ in the labeled case, too. \square

Acknowledgments

This work was initiated during the 9th European Research Week on Geometric Graphs and Pseudo-Triangulations, held May 14–18, 2012 in Alcalá de Henares, Spain. We thank Vincent Pilaud, Pedro Ramos, and André Schulz for helpful comments.

Oswin Aichholzer and Birgit Vogtenhuber are partially supported by the ESF EUROCORES programme EuroGIGA – CRP ‘ComPoSe’, Austrian Science Fund (FWF): I648-N18. Thomas Hackl is supported by the Austrian Science Fund (FWF): P23629-N18 ‘Combinatorial Problems on Geometric Graphs’. David Orden is partially supported by MICINN Project MTM2011-22792, ESF EUROCORES programme EuroGIGA - ComPoSe IP04 - MICINN Project EUI-EURC-2011-4306 and Junta de Castilla y León Project VA172A12-2. Alexander Pilz is recipient of a DOC-fellowship of the Austrian Academy of Sciences at the Institute for Software Technology, Graz University of Technology, Austria. Maria Saumell is supported by ESF EuroGIGA project ComPoSe as F.R.S.-FNRS - EUROGIGA NR 13604 and by ESF EuroGIGA project GraDR as GAČR GIG/11/E023.

References

- [1] O. Aichholzer, F. Aurenhammer, T. Hackl, C. Huemer, A. Pilz, and B. Vogtenhuber. 3-colorability of pseudo-triangulations. In *EuroCG 2010*, pages 21–24, 2010.
- [2] O. Aichholzer, F. Aurenhammer, H. Krasser, and P. Braß. Pseudotriangulations from surfaces and a novel type of edge flip. *SIAM J. Comput.*, 32(6):1621–1653, 2003.
- [3] S. Bereg. Transforming pseudo-triangulations. *Inf. Process. Lett.*, 90(3):141–145, 2004.
- [4] P. Bose and F. Hurtado. Flips in planar graphs. *Comput. Geom.*, 42(1):60–80, 2009.

- [5] H. Brönnimann, L. Kettner, M. Pocchiola, and J. Snoeyink. Counting and enumerating pointed pseudotriangulations with the greedy flip algorithm. *SIAM J. Comput.*, 36(3):721–739, 2006.
- [6] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete Comput. Geom.*, 22:333–346, 1999.
- [7] L. Kettner, D. Kirkpatrick, A. Mantler, J. Snoeyink, B. Speckmann, and F. Takeuchi. Tight degree bounds for pseudo-triangulations of points. *Comput. Geom.*, 25(1–2):3–12, 2003.
- [8] H. Komuro. The diagonal flips of triangulations on the sphere. *Yokohama Mathematical Journal*, 44(2):115–122, 1997.
- [9] C. L. Lawson. Transforming triangulations. *Discrete Math.*, 3(4):365–372, 1972.
- [10] D. Orden, F. Santos, B. Servatius, and H. Servatius. Combinatorial pseudo-triangulations. *Discrete Math.*, 307(3–5):554–566, 2007.
- [11] G. Rote, F. Santos, and I. Streinu. *Pseudo-triangulations — a survey*. Contemp. Math. AMS, 2008.
- [12] D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Short encodings of evolving structures. *SIAM J. Discrete Math.*, 5(3):428–450, 1992.
- [13] I. Streinu. A combinatorial approach to planar non-colliding robot arm motion planning. In *FOCS*, pages 443–453. IEEE Computer Society, 2000.
- [14] K. Wagner. Bemerkungen zum Vierfarbenproblem. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 46:26–32, 1936. In German.