Adapting (Pseudo)-Triangulations with a Near-Linear Number of Edge Flips

Oswin Aichholzer[1], Franz Aurenhammer, Hannes Krasser[2]

1 Introduction

In geometric data processing, structures that partition the geometric input, as well as connectivity structures for geometric objects, play an important role. A versatile tool in this context are triangular meshes, often called triangulations; see e.g., the survey articles [6, 12, 5]. A *triangulation* of a finite set S of points in the plane is a maximal planar straight-line graph that uses all and only the points in S as its vertices. Each face in a triangulation is a triangle spanned by S.

In the last few years, a relaxation of triangulations, called *pseudo-triangulations* (or geodesic triangulations), has received considerable attention. Here, faces bounded by three concave chains, rather than by three line segments, are allowed. The scope of applications of pseudo-triangulations as a geometric data stucture ranges from ray shooting [10, 14] and visibility [25, 26] to kinetic collision detection [1, 21, 22], rigidity [32, 29, 15], and guarding [31]. Still, only very recently, results on the combinatorial properties of pseudo-triangulations have been obtained. These include bounds on the minimal vertex and face degree [20] and on the number of possible pseudo-triangulations [27, 3].

The usefulness of (pseudo-)triangulations partially stems from the fact that these structures can be modified by constant-size combinatorial changes, commonly called *flip operations*. Flip operations allow for an adaption to local requirements, or even for generating globally optimal structures [6, 12]. A classical result states that any two triangulations of a given planar point set can be made to coincide by applying a quadratic number of edge flips; see e.g. [16, 19]. A similar result has been proved recently for the class of *minimum* pseudo-triangulations [8, 29].

Results and outline The present paper demonstrates that the quadratic bound for the number of required flip operations can be beaten drastically. We will provide two main results – for minimum pseudo-triangulations when using traditional flips operations, as well as for triangulations when a novel and natural edge flip operation is included into the repertoire of admissible flips. Extending the set of flips allows to transform pseudo-triangulations of arbitrary edge rank into each other, using a near-linear number of flips, and without changing the underlying set of vertices. A tool for rapidly adapting pseudo-triangulations (and in particular, triangulations) becomes available, using constant-size combinatorial changes. Various applications may be expected, in areas where (pseudo-)triangular meshes are of importance. In particular,

[1]Institute for Software Technology, Graz University of Technology, Graz, Austria. Work done while this author was with the Institute for Theoretical Computer Science, Graz University of Technology, Graz, Austria. Research partially supported by APART [Austrian Programme for Advanced Research and Technology] of the Austrian Academy of Sciences.

[2]Institute for Theoretical Computer Science, Graz University of Technology, Graz, Austria. Research supported by the FWF [Austrian Fonds zur Förderung der Wissenschaftlichen Forschung]

the new flip type is indispensable for a proper treatment of spatial embeddings of pseudo-triangulations; see [2].

Section 2 starts with revisiting edge flip operations in pseudo-triangulations. The *edge-removing flip* is introduced, and is shown to have a geodesics interpretation consistent with classical flip types. Section 3 demonstrates that when edge-removing flips are admitted, the flip distance between any two pseudo-triangulations (and especially, triangulations) of a set of n points is reduced to $O(n \log n)$. In particular, any two triangulations of a simple polygon with n vertices can be transformed into each other with at most $2n-6$ flips. This substantially improves over the situation without using the new type, where an $\Omega(n^2)$ lower bound for triangulations holds [16, 19]. We also show that every given triangulation can be made minimum (i.e. pointed, see below) using $O(n)$ flips. In Section 4, we derive an $O(n \log^2 n)$ bound on the flip distance *within* the class of minimum pseudo-triangulations, that is, without applying the new flip type. This improves previous bounds of $O(n^2)$ in [8, 29], and shows that the diameter of the high-dimensional polytope in [29] is $O(n \log^2 n)$. Our results partially rely on new partitioning results for pseudo-triangulations, in Section 5, which may be of separate interest. Section 6 discusses relations of the edge-removing flip to known existing types. In view of the lack of non-trivial lower bounds, a reduction of flip distances to $O(n)$ is left as an open problem.

Basic properties of pseudo-triangulations This is a brief review of basic notions and properties concerning pseudo-triangulations. For more details, see e.g. [20, 8, 29]. For a (simple) polygon P in the plane, let $vert(P)$ denote the set of vertices of P. A *corner* of P is a vertex with internal angle less than π. The other vertices of P are called *non-corners*. The chain of edges between two consecutive corners of P is called a *side chain* of P. The *geodesic* between two points $x, y \in P$ is the shortest curve that connects x and y and lies inside P. A *pseudo-triangle* is a polygon with exactly three corners.

Let S be a finite set of points in the plane. We will assume, throughout this paper, that S is in general position, i.e., no three points in S are collinear. Let $conv(S)$ denote the convex hull of S. A *pseudo-triangulation* of S is a partition of $conv(S)$ into pseudo-triangles whose vertex set is exactly S. A pseudo-triangulation is a face-to-face two-dimensional cell complex. The intersection of two faces (pseudo-triangles) may consist of up to two edges, however. In case of such *double-adjacencies*, the union of the two adjacent pseudo-triangles is a pseudo-triangle itself.

Let $\mathcal{PT}$ be some pseudo-triangulation of S. A vertex of $\mathcal{PT}$ is called *pointed* if its incident edges lie in an angle smaller than π. Note that all vertices of $conv(S)$ are pointed. The more pointed vertices there are in $\mathcal{PT}$, the less edges and faces it has. In particular, $\mathcal{PT}$ contains exactly $3n - p - 3$ edges and $2n - p - 2$ pseudo-triangles, if $|S| = n$ and there are $p \leq n$ pointed vertices in $\mathcal{PT}$.

We define the *edge rank* of $\mathcal{PT}$ as $n - p$. The minimum edge rank is zero, where $\mathcal{PT}$ is commonly called a *minimum* (or a *pointed*) pseudo-triangulation. $\mathcal{PT}$ then is a maximal planar straight-line graph on S where all vertices are pointed; see e.g. [32]. It contains exactly $2n-3$ edges and $n-2$ pseudo-triangles. The edge rank expresses the excess in edges, compared to a minimum pseudo-triangulation. Its value is at most $n - |vert(conv(S))|$, which is attained if and only if $\mathcal{PT}$ is a triangulation.

2 Flips in pseudo-triangulations revisited

2.1 Classical flips

So-called *flips* are operations of constant combinatorial complexity which are commonly used to modify triangulations. The standard edge flip, also called *Lawson flip* [23], takes two triangles Δ_1 and Δ_2 whose union is a convex quadrilateral and exchanges its diagonals e and e'. To generalize to pseudo-triangulations, a different view of this edge flip is of advantage: Take the vertex of Δ_1 and Δ_2, respectively, that lies opposite to e and replace e by the geodesic between these two vertices. The geodesic is just a line segment e' in this case.

The geodesics interpretation above has been used in [25, 32] to define flips in minimum pseudo-triangulations. Let ∇_1 and ∇_2 be two adjacent pseudo-triangles, and let e be an edge they have in common. A flip replaces e by the part contributed by the geodesic inside $\nabla_1 \cup \nabla_2$ that connects the two corners of ∇_1 and ∇_2 opposite to e. In a minimum pseudo-triangulation each vertex is pointed, so the geodesic indeed contributes a line segment e' which is no edge of ∇_1 or ∇_2. See Figure 1(a) and (b), where the edge e to be flipped is shown in bold. Note that the flipping partners e and e' may cross or not. In either case, the flip creates two valid pseudo-triangles. We refer to such flips as *exchanging* flips. Each internal edge in a minimum pseudo-triangulation is flippable in this way. In a pseudo-triangulation of non-zero edge rank, however, edges incident to non-pointed vertices may be *non-flippable* in this sense. In particular, in a full triangulation, an internal edge is non-flippable if and only if its two incident triangles form a non-convex quadrilateral; see Figure 1(c). Non-flippable edges have been the source for the theoretically poor behavior of certain flipping algorithms, concerning the flip distance [16, 19] as well as the non-existence of flip sequences [11].

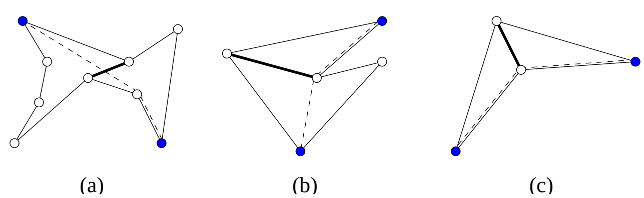

Figure 1: Exchanging flips and non-flippable edge

2.2 A novel flip type

We wish to generalize the edge flip so as to cover the situation in Figure 1(c) as well. In fact, when being consistent with the geodesics rule above, flipping a non-flippable edge $e = \nabla_1 \cap \nabla_2$ means removing e because its substitute is empty. A pseudo-triangle $\nabla_1 \cup \nabla_2$ is obtained. We include this *edge-removing* flip (and its inverse, the *edge-inserting* flip) into the repertoire of admissible flips. By definition, an edge-removing flip is applicable only if a valid pseudo-triangle is created. That is, a single non-pointed vertex of the pseudo-triangulation is made pointed by the flip.

This simple modification makes each internal edge in every pseudo-triangulation (and in particular, in every triangulation) flippable. Note that edge-removing flips decrement the edge rank, whereas edge-inserting flips increment it. This allows for 'surfing' between pseudo-triangulations of different edge ranks. Several interesting consequences will be discussed, including a reduction of flip distances in Section 3, and relations to other flip types in Section 6.

Remarks Edge-removing flips arise implicitly in a greedy flip algorithm for pseudo-triangulations of *convex objects*, in [26]. Certain flips that exchange bitangents of such objects cause an edge removal (or insertion) in the corresponding pseudo-triangulation for the object centers.[3]

3 Reducing the flip distance

Let S be a set of n points in the plane. It is well known that $\Theta(n^2)$ Lawson flips may be necessary, and are also sufficient, to transform two given triangulations of S into each other; see e.g. [16, 19]. The upper bound also applies to exchanging flips in minimum pseudo-triangulations, see [8, 29], but no non-trivial lower bounds are known in this case. For our admissible set of flip operations, several results will be shown in this section.

3.1 Simple polygons

We start with proving that flip distances become linear between pseudo-triangulations of (simple) polygons, when edge-removing flips and their inverses are allowed. Consider a polygon P in the plane. The *shortest-path tree* of P with root $v \in vert(P)$ is the union of all geodesics in P from $vert(P)$ to v. Let $\pi_v(P)$ denote this structure. It is well known [17] that $\pi_v(P)$ is a tree that partitions P into pseudo-triangles in a unique way.

Lemma 1 *Let P be a polygon with n vertices, and let $v \in vert(P)$. The shortest-path tree $\pi_v(P)$ can be constructed by triangulating P arbitrarily, and applying at most $n-3$ exchanging or edge-removing flips.*

Proof. Fix some triangulation $\mathcal{T}$ of P. We prove the assertion by induction on the number of triangles of $\mathcal{T}$. As an induction base, let Q be the union of all triangles of $\mathcal{T}$ incident to the vertex v. Clearly, the restriction of $\mathcal{T}$ to Q just gives $\pi_v(Q)$. We show that this invariant can be maintained by flipping, when an adjacent triangle Δ of $\mathcal{T}$ is added to Q.

Let u be the vertex of Δ that does not belong to Q. Consider the unique edge $e = Q \cap \Delta$ (which is a diagonal of P). If e belongs to $\pi_v(Q \cup \Delta)$ then an edge of Δ connects u to $\pi_v(Q)$, and $\pi_v(Q \cup \Delta)$ is already complete. No flip is performed. Else let ∇ denote the unique pseudo-triangle in $\pi_v(Q)$ that is adjacent to Δ at e. There are two cases. If $\nabla \cup \Delta$ is a pseudo-triangle then, again, u is connected to $\pi_v(Q)$ by Δ. Perform a flip that removes e, which restores $\pi_v(Q \cup \Delta)$. Otherwise, let w be the corner of ∇ opposite to e. Apply an exchanging flip to e. The new edge e' lies on the geodesic between u and w. Thus e' connects u to $\pi_v(Q)$, which constructs $\pi_v(Q \cup \Delta)$ in this case.

The total number of flips is at most $n - 3$, because each flip can be charged to the triangle of $\mathcal{T}$ that is added. ◻

[3] We recently learned that Orden and Santos [24] also considered this type of flip, to obtain a polytope representation of all possible pseudo-triangulations of a given point set.

Corollary 1 *Any two triangulations of a polygon P with n vertices can be flipped into each other by at most $2n-6$ exchanging, edge-removing, or edge-inserting flips.*

Proof. Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two triangulations of P. Choose some $v \in vert(P)$ and flip $\mathcal{T}_1$ to $\pi_v(P)$. Then flip $\pi_v(P)$ to $\mathcal{T}_2$ by reversing the sequence of flips that transforms $\mathcal{T}_2$ to $\pi_v(P)$. This is possible and takes at most $2n-6$ flips, by Lemma 1. ◻

Corollary 1 implies a flip distance of $O(n)$ between *any* two pseudo-triangulations $\mathcal{PT}_1$ and $\mathcal{PT}_2$ of a given polygon P, because $\mathcal{PT}_1$ and $\mathcal{PT}_2$ can be completed to triangulations of P with $O(n)$ edge-inserting flips.

3.2 Planar point sets

We continue with pseudo-triangulations of planar point sets. In fact, we choose a slightly more general scenario, namely a point set enclosed by an arbitrary simple polygon (a so-called augmented polygon). This setting will turn out to be more appropriate for our developments, as it arises naturally from constraining the pseudo-triangulated domain. We will show how to flip any given pseudo-triangulation into a canonical one, by splitting the underlying augmented polygon in a balanced way, until empty polygons are obtained and Corollary 1 applies.

Let P be a polygon, and consider a finite point set $S \subset P$ with $vert(P) \subseteq S$. We call the pair (P, S) an *augmented* polygon. A pseudo-triangulation $\mathcal{PT}$ of (P, S) is a partition of P into pseudo-triangles whose vertex set is exactly S. It contains exactly $3n-m+k-p-3$ edges and $2n-m+k-p-2$ pseudo-triangles if $|S| = n$, P is an m-gon with k corners, and p counts the pointed vertices of $\mathcal{PT}$. The maximum edge rank of $\mathcal{PT}$ is $n-k$. In the special case $P = conv(S)$, we have $m = k$ and deal with pseudo-triangulations of the point set S. Below is another corollary to Lemma 1.

Corollary 2 *Let $\mathcal{T}$ be a (full) triangulation of an augmented polygon (P, S). Let e be some line segment spanned by S, which lies inside P and crosses $\mathcal{T}$ at $j \geq 1$ edges. Then $\mathcal{T}$ can be modified to a triangulation that contains e by applying $O(j)$ exchanging, edge-removing, or edge-inserting flips.*

Proof. Let Q be the union of the triangles of $\mathcal{T}$ that are crossed by e. Note that Q may contain points of S in its interior, or even may contain holes, namely if Q contains internal edges which do not cross e. In this case, we cut Q along these edges, and move Q apart infinitesimally at the cuts, to obtain a simple polygon empty of points in S. This is possible because general position is assumed for S. Now choose any triangulation $\mathcal{T}_e$ of Q which includes the edge e to be integrated. By Corollary 1, the part of $\mathcal{T}$ inside Q can be flipped to $\mathcal{T}_e$ by $O(j)$ flips. ◻

We are now prepared to prove the following general assertion on flip distances.

Theorem 1 *Any two pseudo-triangulations of a given planar point set S (or more generally, of a given augmented polygon (P, S)) can be transformed into each other by applying $O(n \log n)$ flips of the types exchanging, edge-removing, and edge-inserting, for $n = |S|$.*

Proof. The two pseudo-triangulations of the augmented polygon (P, S) in question can be completed to triangulations by applying $O(n)$ edge-inserting flips. We show how to transform two arbitrary triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ of (P, S) into the same, using

$O(n \log n)$ flips. Let P be an m-gon. If $m = n$ then $O(n)$ flips suffice by Corollary 1. Else we partition P into subpolygons, each containing at most $\frac{2}{3}(n-m)$ points of $S \setminus vert(P)$. A constant number of line segments spanned by S suffice for this purpose, by Theorem 4(1) in Section 5. Incorporate these segments into $\mathcal{T}_1$ and $\mathcal{T}_2$, respectively, in $O(n)$ flips, which is possible by Corollary 2. Treat the obtained $O(1)$ augmented polygons recursively. This yields a polygonal partition of P whose vertex set is exactly S, and two triangulations thereof, in $O(n \log n)$ flips. By Corollary 1, another $O(n)$ flips let these two triangulations coincide. $\square$

Remarks Theorem 1 demonstrates that flip distances are substantially reduced when using the new flip type. 'Shortcuts' via pseudo-triangulations with varying edge rank become possible. The interested reader may check that the constant involved in the $O(n \log n)$ term is small (less than 6). We conjecture that Theorem 1 can be improved to $O(n)$ flips, because the $\Omega(n^2)$ worst-case examples for Lawson flips in triangulations are based on (non-convex) polygons without internal points [19], an instance covered by Corollary 1 in $O(n)$ flips.

All flips used in Theorem 1 are constant-size combinatorial operations, which can be carried out in $O(\log m)$ time each, if the size of the two pseudo-triangles involved is at most m; see e.g. [13]. This implies that any two (pseudo-)triangulations of a given set of n points can be adapted by local operations in $O(n \log^2 n)$ time – a result we expect to have various applications.

It is well known that not every pseudo-triangulation can be made minimum by removing edges. It can only be made *minimal* in edge rank, and is termed thereafter in [1, 20, 30]. In particular, a minimal pseudo-triangulation may be a full triangulation, even when its vertices are not in convex position [30]. We can show the following:

Lemma 2 *Let $\mathcal{PT}$ be any pseudo-triangulation of a planar n-point set S. Then $\mathcal{PT}$ can be transformed into a minimum pseudo-triangulation of S with $O(n)$ exchanging, edge-removing, or edge-inserting flips.*

4 Minimum pseudo-triangulations

Our next aim is to provide a stronger version of Theorem 1, namely for minimum pseudo-triangulations and *without* using the new flip type. That is, we restrict ourselves to staying within the class of minimum pseudo-triangulations, and use exchanging flips exclusively. By extending and modifying the arguments in Subsections 3.1 and 3.2, we will arrive at a flip distance bound of $O(n \log^2 n)$.

4.1 Two basic tools

Let P be a polygon. The *minimum shortest-path tree* $\mu_c(P)$, for a fixed corner c of P, is the union of all geodesics inside P that lead from c to the *corners* of P. Observe that $\mu_c(P)$ defines a minimum pseudo-triangulation for P, which is a subset of the shortest-path tree $\pi_c(P)$. The proof of Lemma 1 now can be adapted easily, to show that every minimum pseudo-triangulation of P can be transformed into $\mu_c(P)$ by at most $n-3$ *exchanging* flips. The new flip type is not used here, because each edge is flippable in the classical sense. We obtain:

Lemma 3 *Let P be a polygon with k corners. Any two minimum pseudo-triangulations of P are transformable into each other by applying at most $2k-6$ exchanging flips.*

The following assertion (which we state here without proof) is a variant of Corollary 2, for minimum pseudo-triangulations.

Lemma 4 *Let $\mathcal{MPT}$ be a minimum pseudo-triangulation of an augmented polygon (P, S), and let G be some pointed planar straight-line graph on S and in P. Then $\mathcal{MPT}$ can be made to contain G by applying $O(nj)$ exchanging flips, if S has n vertices and $G \setminus P$ has j edges.*

4.2 Exchanging-flip distance bound

Lemma 4 implies an $O(n^2)$ bound on the exchanging-flip distance in minimum pseudo-triangulations. The following theorem shows that this can be improved.

Theorem 2 *Let S be a set of n points in the plane, and let $\mathcal{MPT}_1$ and $\mathcal{MPT}_2$ be two minimum pseudo-triangulations of S. Then $\mathcal{MPT}_1$ can be transformed into $\mathcal{MPT}_2$ by applying $O(n \log^2 n)$ exchanging flips. No other flip types are used. The same result holds for augmented polygons (P, S).*

Proof. Consider an augmented polygon (P, S). We recursively split (P, S) in a balanced way, by applying Theorem 4(2) from Section 5. This constructs a polygonal partition Π of P whose vertex set is S, and where all vertices are pointed. Π is obtained by introducing $O(\log n)$ edges internal to P in each recursive step, and the number of recursive steps is $O(\log n)$ as well.

By Lemma 4, $O(n \log^2 n)$ exchanging flips are sufficient to make $\mathcal{MPT}_1$ and $\mathcal{MPT}_2$ contain all the edges of Π. Finally, Lemma 3 allows for adapting pseudo-triangles within the polygons of Π in $O(n)$ such flips. $\square$

Remarks Theorem 2 improves the recent bound of $O(n^2)$ in [8] for minimum pseudo-triangulations of point sets. Again, we conjecture that the truth is $O(n)$ flips.[4]

In [29], the polytope, $\mathcal{M}(S)$, of minimum pseudo-triangulations of a point set S has been introduced. $\mathcal{M}(S)$ is a high-dimensional convex polytope. Its vertices correspond to all the minimum pseudo-triangulations of S, and its edges represent all possible exchanging flips. By Theorem 2, the diameter of $\mathcal{M}(S)$ is bounded by $O(n \log^2 n)$.

There are examples where the transformation between two given minimum pseudo-triangulations is speeded up by using intermediate edge-inserting and edge-removing flips; see [18]. This indicates that flexibility of pseudo-triangulations does not only come from low edge rank, but also stems from the ability to change this parameter – using the new flip type.

5 Partitioning results

This section presents some partitioning results concerning pseudo-triangulations, which have been referred to in Sections 3 and 4. The theorems in Subsections 5.1 and 5.2 might be of separate interest.

[4]By a very recent result in [7], a flip distance of $O(n \log n)$ for minimum pseudo-triangulations of point sets is obtainable, using a different divide-and-conquer approach. This approach does not carry over to the more general case of augmented polygons, however.

5.1 Pseudo-triangulations with small cut

Let P be a simple polygon. Consider a pseudo-triangle $\nabla \subset P$ with vertices from $vert(P)$. ∇ is called *nice* if its three corners are corners of P. We define the *cut* of ∇ as the number of diagonals of P on ∇'s boundary.

A polygon P is *pseudo-convex* if every geodesic inside P is a convex chain. A *corner tangent* is an inner tangent of P incident to at least one corner of P. A pseudo-convex polygon P is *strongly pseudo-convex* if no corner tangent exists for P. We state the following fact without proof.

Lemma 5 *Let P be a strongly pseudo-convex polygon with k corners. There exists a nice pseudo-triangle for P with cut $O(\log k)$.*

We are now ready to prove the following structural result for minimum pseudo-triangulations of simple polygons.

Theorem 3 *For every polygon P with n vertices, there exists a minimum pseudo-triangulation of P where each face has cut $O(\log n)$.*

PROOF. We first partition P into strongly pseudo-convex polygons. Diagonals on non-convex geodesics and corner tangents are used, such that each introduced diagonal is incident to some corner in both polygons it bounds. (These diagonals will contribute to the cut of the final faces in the minimum pseudo-triangulation to be constructed, but their number is at most 6 per face.) Each strongly pseudo-convex polygon Q with more than 3 corners is partitioned further as follows.

Integrate a nice pseudo-triangle ∇ with small cut for Q, whose existence is guaranteed by Lemma 5. Because ∇ is nice, it does not violate the pointedness of any vertex. Moreover, each diagonal of Q on ∇'s boundary is incident to two corners of the polygon it cuts off from Q. These polygons are partitioned recursively.

A minimum pseudo-triangulation $\mathcal{MPT}$ of P results. Each face f of $\mathcal{MPT}$ has cut $O(\log n)$: A diagonal on f's boundary comes from Lemma 5 or is among the at most 6 edges incident to some corner of f. ◻

Remarks Theorem 3 is asymptotically optimal. There exist polygons with n vertices where every minimum pseudo-triangulation contains some face with cut $\Omega(\log n)$; see [4]. The theorem is related to a result in [20] which shows, for every point set S, the existence of a minimum pseudo-triangulation with constant face complexity. Another related result, in [10], shows that every n-gon P admits a minimum pseudo-triangulation $\mathcal{MPT}$ such that each line segment interior to P crosses only $O(\log n)$ edges of $\mathcal{MPT}$.

5.2 Partition theorem

We continue with a ham-sandwich type result for pseudo-triangles.

Lemma 6 *Let ∇ be a pseudo-triangle that contains a set M of i points in its interior. There exists a point $p \in M$ whose geodesics to two corners of ∇ divide M into two subsets of cardinality $\frac{2i}{3}$ or less.*

PROOF. For each point $p \in M$, the geodesics from p to the three corners of ∇ partition ∇ into three pseudo-triangles (faces). Such a face f is called *sparse* if f encloses at most $\frac{2\,i}{3}$ points of M. We claim that, for each pair c, c' of corners of ∇, there exist at least $\frac{2\,i}{3}+1$ sparse faces: Consider the sorted order of M, as given by the shortest-path tree with root c for M. The j-th point of M in this order spans a face that contains strictly less than j points.

We conclude that there are at least $2\,i+3$ sparse faces in total. So the mean number of sparse faces per point in M exceeds two, which implies that there exists a point $p \in M$ incident to three sparse faces. Among them, let f be the face that contains the most points, which are at least $\frac{i}{3}$. We take the two geodesics that span f to partition ∇. This yields two parts with at most $\frac{2\,i}{3}$ points each. ◻

Lemma 6 combines with Theorem 3 to the following partition theorem for augmented polygons.

Theorem 4 *Let (P, S) be an augmented polygon, and let $I = S \setminus vert(P)$. There exist polygonal partitions Π_1 and Π_2 of (P, S) such that*

(1) Π_1 uses $O(1)$ line segments spanned by S, and assigns at most $\frac{2}{3} \cdot |I|$ points of I to each polygon

(2) Π_2 uses $O(\log n)$ line segments spanned by S, assigns at most $\frac{2}{3} \cdot |I|$ points of I to each polygon, and guarantees the pointedness of each vertex of S.

Proof. To construct Π_1, let $\mathcal{T}$ be some triangulation of the polygon P. Call a polygon $Q \subset P$ *sparse* if Q contains at most $\frac{2}{3} \cdot |I|$ points of I. Let ∇ be any face of $\mathcal{T}$. If each part of $P \setminus \nabla$ is sparse then we are done, because we can partition P with ∇, and ∇ with two line segments as in Lemma 6 if ∇ is non-sparse. Otherwise, we continue with the face of $\mathcal{T}$ adjacent to ∇ in the (unique) non-sparse part of $P \setminus \nabla$, until the first condition is met.

To construct Π_2 we proceed analogously, but use a minimum pseudo-triangulation $\mathcal{MPT}$ of P with face cuts bounded by $O(\log n)$. The existence of $\mathcal{MPT}$ is given by Theorem 3. The $O(\log n)$ edges of ∇ that are used to partition (P, S) retain the pointedness of all vertices, as do the two segments from Lemma 6 that may have to be used to split ∇. ◻

Remarks The fraction $\frac{2}{3}$ in Lemma 6 is optimal, even if ∇ is a triangle. The set M may consist of three groups of $\frac{i}{3}$ points such that, for each choice of $p \in M$, the two groups not containing p end up in the same subset.

Theorem 4 is similar in flavor to a result in [9], which asserts that any simple n-gon can be split by a diagonal into two subpolygons with at most $\frac{2}{3}n$ vertices.

6 Relation between flip types

The new flip type introduced in Section 2 can be used to simulate certain other flip types. Let us briefly comment on this fact.

As an example, the exchanging flip in Figure 1(b) can be simulated by an edge-inserting flip followed by an edge-removing flip. Interestingly, this is not possible for the exchanging flip in Figure 1(a). A more important example concerns a flip that arises in the context of Delaunay and regular triangulations; see [28, 11]. This flip inserts a new vertex v in the interior of a triangle Δ, and connects v by edges to the three vertices of Δ. Vertex insertion is meaningful for pseudo-triangulations as well [32, 3]. Connect v by geodesics to (at least two) corners of the pseudo-triangle ∇

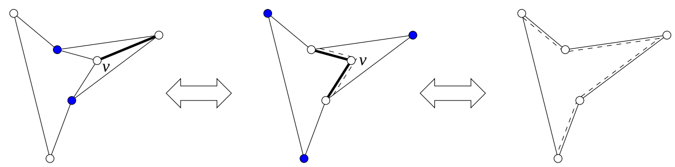

Figure 2: Edge-removing and vertex-removing flips

that v lies in. Each geodesic contributes one edge incident to v, and ∇ is partitioned into (two or three) pseudo-triangles.

The inverse operation, namely the removal of a degree-3 vertex v, can be simulated using edge-removing flips. See Figure 2. Apply an edge-removing flip to one of v's edges first, which leaves a partition of ∇ into two pseudo-triangles in double-adjacency. Then, carry out two edge-removing flips *simultaneously*. This deletes v and leaves ∇ empty, because no edges are created by the geodesics rule. This simultaneous operation can be considered a single flip – the *vertex-removing* flip. By definition, such a flip is applicable to vertices of degree 2 only. Vertex-removing flips (as well as edge-removing flips) play an important role for surface realizations of pseudo-triangulations in three-space [2].

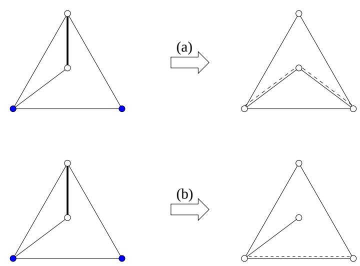

Figure 3: Ambiguous geodesics interpretation

Remarks Instead of the vertex-removing flip, a different version – namely the exchanging flip in Figure 3(a) – has been commonly used. It also leads to a valid pseudo-triangulation (which now does contain the vertex v). However, care has to be taken not to misinterpret this version as in Figure 3(b), where the geodesic still lies inside the union of the two pseudo-triangles involved. Also, this version conflicts with a three-dimensional interpretation of flips in surfaces [2]. When change in edge rank is conceded, we may circumvent the flip in Figure 3(a) by performing two consecutive flips of the new type, namely an edge-inserting flip followed by an edge-removing flip.

Vertex-removing and vertex-inserting flips are not used in Theorems 1 and 2. Dropping this restriction makes things easy, because every pseudo-triangulation contains some vertex of constant degree, which can be removed with $O(1)$ flips. A flip distance of $O(n)$ is obvious in this setting. However, removing a vertex does not only change the vertex set S, but rather changes the underlying domain (the polygon P) if removal was for a boundary vertex. In contrast, in the setting for the theorems above, both S and P remain unchanged. The situation where S but not P is allowed to change is of some interest, because no internal vertex of constant degree might exist in a pseudo-triangulation. We are able to show the following:

Lemma 7 *Let S be a planar n-point set. Any two pseudo-triangulations of S can be flipped into each other in $O(n)$ flips, and without changing the underlying domain $conv(S)$, if the entire repertoire of flips from Section 2 is used.*

Acknowledgements We gratefully acknowledge discussions on the presented topic with Michel Pocchiola, Günter Rote, and Francisco Santos.

References

[1] P.K.Agarwal, J.Basch, L.J.Guibas, J.Hershberger, L.Zhang. Deformable free space tilings for kinetic collision detection. In B.R.Donald, K.Lynch, D.Rus (eds.), Algorithmic and Computational Robotics: New Directions (Proc. 5th Workshop Algorithmic Found. Robotics), 2001, 83-96.

[2] O.Aichholzer, F.Aurenhammer, P.Brass, H.Krasser. Spatial embedding of pseudo-triangulations. Proc. 19^{th} Ann. ACM Sympos. Computational Geometry 2003, to appear.

[3] O.Aichholzer, F.Aurenhammer, H.Krasser, B.Speckmann. Convexity minimizes pseudo-triangulations. Proc. 14^{th} Canadian Conf. Computational Geometry 2002, 158-161.

[4] O.Aichholzer, M.Hoffmann, B.Speckmann, C.D.Tóth. Degree bounds for constrained pseudo-triangulations. Manuscript, Institute for Theoretical Computer Science, Graz University of Technology, Austria, 2003.

[5] F.Aurenhammer, Y.-F.Xu. Optimal triangulations. In: P.M.Pardalos, C.A.Floudas (eds), Encyclopedia of Optimization 4, Kluwer Academic Publishing, 2000, 160-166.

[6] M.Bern, D.Eppstein. Mesh generation and optimal triangulation. In: D.-Z.Du, F.Hwang (eds), Computing in Euclidean Geometry, Lecture Notes Series on Computing 4, World Scientific, 1995, 47-123.

[7] S.Bespamyatnikh. Transforming pseudo-triangulations. Manuscript, Dept. Comput. Sci., University of Texas at Dallas, 2003.

[8] H.Brönnimann, L.Kettner, M.Pocchiola, J.Snoeyink. Counting and enumerating pseudo-triangulations with the greedy flip algorithm. Manuscript, 2001.

[9] B.Chazelle. A theorem on polygon cutting with applications. Proc. 23^{rd} IEEE Symp. FOCS, 1982, 339-349.

[10] B.Chazelle, H.Edelsbrunner, M.Grigni, L.J.Guibas, J.Hershberger, M.Sharir, J.Snoeyink. Ray shooting in polygons using geodesic triangulations. Algorithmica 12 (1994), 54-68.

[11] H.Edelsbrunner, N.R.Shah. Incremental topological flipping works for regular triangulations. Algorithmica 15 (1996), 223-241.

[12] S.Fortune. Voronoi diagrams and Delaunay triangulations. In: D.-Z.Du, F.Hwang (eds), Computing in Euclidean Geometry, Lecture Notes Series on Computing 4, World Scientific, 1995, 225-265.

[13] J.Friedman, J.Hershberger, J.Snoeyink. Efficiently planning compliant motion in the plane. SIAM J. Computing 25 (1996), 562-599.

[14] M.T.Goodrich, R.Tamassia. Dynamic ray shooting and shortest paths in planar subdivisions via balanced geodesic triangulations. J. Algorithms 23 (1997), 51-73.

[15] R.Haas, D.Orden, G.Rote, F.Santos, B.Servatius, H.Servatius, D.Souvaine, I.Streinu, W.Whiteley. Planar minimally rigid graphs and pseudo-triangulations. Proc. 19^{th} Ann. ACM Sympos. Computational Geometry, to appear.

[16] S.Hanke, T.Ottmann, S.Schuierer. The edge-flipping distance of triangulations. Journal of Universal Computer Science 2 (1996), 570-579.

[17] J.Hershberger. An optimal visibility graph algorithm for triangulated simple polygons. Algorithmica 4 (1989), 141-155.

[18] C.Huemer. Master Thesis, Institute for Theoretical Computer Science, Graz University of Technology, Austria, 2003.

[19] F.Hurtado, M.Noy, J.Urrutia. Flipping edges in triangulations. Discrete & Computational Geometry 22 (1999), 333-346.

[20] L.Kettner, D.Kirkpatrick, A.Mantler, J.Snoeyink, B.Speckmann, F.Takeuchi. Tight degree bounds for pseudo-triangulations of points. Computational Geometry: Theory and Applications 25 (2003), 3-12.

[21] D.Kirkpatrick, J.Snoeyink, B.Speckmann. Kinetic collision detection for simple polygons. Intern. J. Computational Geometry & Applications 12 (2002), 3-27.

[22] D.Kirkpatrick, B.Speckmann. Kinetic maintenance of context-sensitive hierarchical representations for disjoint simple polygons. Proc. 18^{th} Ann. ACM Sympos. Computational Geometry 2002, 179-188.

[23] C.L.Lawson. Properties of n-dimensional triangulations. Computer Aided Geometric Design 3 (1986), 231-246.

[24] D.Orden, F.Santos. The polyhedron of non-crossing graphs on a planar point set. Manuscript, Universidad de Cantabria, Santander, Spain, 2002.

[25] M.Pocchiola, G.Vegter. Minimal tangent visibility graphs. Computational Geometry: Theory and Applications 6 (1996), 303–314.

[26] M.Pocchiola, G.Vegter. Topologically sweeping visibility complexes via pseudo-triangulations. Discrete & Computational Geometry 16 (1996), 419-453.

[27] D.Randall, G.Rote, F.Santos, J.Snoeyink. Counting triangulations and pseudo-triangulations of wheels. Proc. 13^{th} Canadian Conf. Computational Geometry 2001, 117-120.

[28] V.T.Rajan, Optimality of the Delaunay triangulation in R^d. Discrete & Computational Geometry 12 (1994), 189-202.

[29] G.Rote, F.Santos, I.Streinu. Expansive motions and the polytope of pointed pseudo-triangulations. In: Discrete & Computational Geometry - The Goodman-Pollack Festschrift, B.Aronov, S.Basu, J.Pach, M.Sharir (eds.), Algorithms and Combinatorics, Springer, Berlin, 2003, 699-736.

[30] G.Rote, C.A.Wang, L.Wang, Y.Xu. On constrained minimum pseudo-triangulations. Manuscript, Inst. f. Informatik, FU-Berlin, 2002.

[31] B.Speckmann, C.D.Toth. Allocating vertex π-guards in simple polygons via pseudo-triangulations. Proc. 14^{th} ACM-SIAM Symposium on Discrete Algorithms, 2003, 109-118.

[32] I.Streinu. A combinatorial approach to planar non-colliding robot arm motion planning. Proc. 41^{st} IEEE Symp. FOCS, 2000, 443-453.