PSEUDO-TRIANGULATIONS FROM SURFACES AND A NOVEL TYPE OF EDGE FLIP

OSWIN AICHHOLZER*, FRANZ AURENHAMMER, HANNES KRASSER†, AND PETER BRASS‡

Abstract. We prove that planar pseudo-triangulations have realizations as polyhedral surfaces in three-space. Two main implications are presented: The spatial embedding leads to a novel flip operation that allows for a drastical reduction of flip distances, especially between (full) triangulations. Moreover, several key results for triangulations, like flipping to optimality, (constrained) Delaunayhood, and a convex polytope representation, are extended to pseudo-triangulations in a natural way.

Key words. Pseudo-triangulation, flip distance, surface realization, locally convex function, constrained regular complex, polytope representation

AMS subject classifications. 68U05, 68W40, 52C45, 52C25, 52B55, 52B05

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 [14, 26, 11]. A *triangulation* of a finite set S of points in the plane is a maximal planar straight-line graph that uses exactly the points in S as its vertices. Each face in a triangulation is a triangle spanned by S.

In recent years, a relaxation of triangulations, called *pseudo-triangulation* (or geodesic triangulation), has received considerable attention. Here, faces bounded by three concave chains, rather than by three line segments, are allowed. Interest in these structures originates from applications to ray shooting [18, 29] and visibility [42, 43]. Meanwhile, their usefulness as an efficient data structure in many other areas has been discovered, including kinetic collision detection [1, 36, 37], rigidity [51, 47, 30], and guarding [44, 49].

Concerning the combinatorial and geometric properties of pseudo-triangulations, knowledge is comparatively sparse. Beside the structural rigidity results mentioned above, certain results on the vertex and face degree [35], on the number of possible pseudo-triangulations [45, 5], and on edge flipping operations in pseudo-triangulations [16, 47] have been obtained recently. An interesting geometric result, in [47] and generalized in [41], addresses polytope representations of pseudo-triangulations.

The present paper intends to show that pseudo-triangulations possess rich geometric properties. This provides a deeper understanding of their combinatorial properties, and of the interrelation of pseudo-triangulations of different levels of edge rank. One of our main results asserts that pseudo-triangulations have realizations as polyhedral surfaces in three-space. This insight leads to a novel edge flip operation, which has

*Institute for Software Technology, Graz University of Technology, Austria. Work done while this author was with the Institute for Theoretical Computer Science, Graz University of Technology, Austria. Research supported by APART [Austrian Programme for Advanced Research and Technology]. e-mail: oaich@ist.tugraz.at

†Institute for Theoretical Computer Science, Graz University of Technology, Austria. Research supported by the FWF [Austrian Fonds zur Förderung der Wissenschaftlichen Forschung]. e-mail: {auren,hkrasser}@igi.tugraz.at

‡Department of Computer Science, City College of New York, U.S.A. Work done while this author was with the Institut für Informatik, Freie Universität Berlin, Germany. Research supported by DFG under grant Br 1465/5-2. e-mail: peter@cs.ccny.cuny.edu

a three-dimensional interpretation in the surface, as well as a consistent geodesics interpretation in the plane. Inclusion of the new flip type allows to transform arbitrary pseudo-triangulations into each other, without changing the underlying set of vertices. A tool for rapidly adapting pseudo-triangulations (in particular, triangulations) and their realizing surfaces becomes possible, using a near-linear number of constant-size combinatorial changes. Pseudo-triangulations which are realizable as locally convex surfaces project to unique constrained regular complexes, a concept which has not been rigorously defined before. We further solve the optimization problem of finding maximal locally convex surfaces over polygonal domains, by bounding the length of (improving) flip sequences that reach the global optimum of local convexity. Finally, we provide a polytope representation of constrained regular pseudo-triangulations.

Subsection 1.1 gives a more detailed overview of our results. We expect various applications, in areas where (pseudo-)triangulations and polyhedral surfaces play a role, such as surface modeling and terrain fitting. Related concepts, that may be affected by our results, are power diagrams [8] (a generalization of Voronoi diagrams, dual to regular triangulations), weighted alpha shapes [24] and dual complexes [23] (certain subcomplexes of regular triangulations), reflex-free hulls [2] (polyhedral surfaces avoiding cavities), and roofs [3] (surface representations of straight skeletons). The focus of the present paper, though, is on providing new and fundamental properties of pseudo-triangulations, rather than on a particular application.

1.1. Overview of results. We start by revisiting edge flipping operations in pseudo-triangulations, in Section 2. A new type, the edge-removing flip, is introduced, and its relationship to existing flip types is discussed. Section 3 shows that, when using the edge-removing flip, the flip distance between any two pseudo-triangulations (and especially, triangulations) of a set of n points is reduced to $O(n \log n)$. This substantially improves over the situation without using the new type, where an $\Omega(n^2)$ lower bound holds for triangulations [31, 34]. We further derive an $O(n \log^2 n)$ bound for minimum pseudo-triangulations, *without* using edge-removing flips. This improves the previous bounds of $O(n^2)$ in [16, 47], and shows that the diameter of the high-dimensional polytope in [47] is $O(n \log^2 n)$. For special settings, even some linear flip distance bounds are achieved. The results above partially rely on new partitioning results for pseudo-triangulations in Section 4, which may be of separate interest.

Section 5 describes realizations of pseudo-triangulations as polyhedral surfaces. We introduce the concept of complete vertices, to distinguish between vertices whose heights in the surface can be predetermined or not. We characterize the class of pseudo-triangulations which are projective, in the sense that their internal edges exactly correspond to the set of non-linearity of a polyhedral surface. This generalizes the situation for triangulations, which are trivially projective. Section 6 gives a surface interpretation of all the admissible flip types from Section 2. Surface flips, though being constant-size combinatorial changes, are geometrically non-local; a single flip may change the surface heights for $\Theta(n)$ vertices. Special cases are tetrahedral surface flips, which are well known from flipping in convex hulls; see e.g. [38, 25].

Section 7 puts the results above to use, by solving the optimization problem of finding maximal locally convex functions on polygonal domains (with n vertices, some possibly internal). We prove that every triangular surface can be made locally convex, by performing surface flips which are convexifying or planarizing. The resulting projection is a pseudo-triangulation. If the underlying domain is convex, or if this domain contains no internal vertices, then $O(n^2)$ flips suffice. The former result strengthens a result in [25]. They show that incrementally inserting the vertices, and flipping

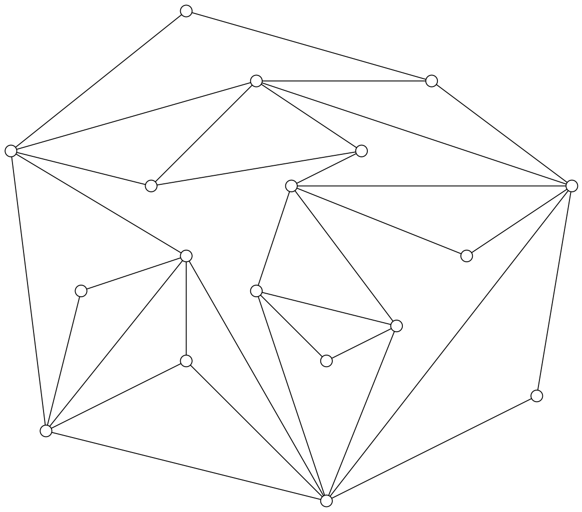

FIG. 1.1. *A pseudo-triangulation for 18 points*

away all reflex edges in the special (triangular) surface that arises from each insertion, is capable of generating the optimum. For convex domains, the optimum is a lower convex hull, and a regular triangulation in the projection. It is well known [25] that the optimum cannot be reached *within* the class of triangulations, in general. This explains the need for edge-removing (i.e., planarizing) flips, and substantiates the usefulness of pseudo-triangulations. Our results can also be viewed as a generalization of a well-known fact, namely that every triangulation of a planar point set can be transformed into the Delaunay triangulation, using so-called Delaunay flips; see e.g. [26].

Extending the scenario, by fixing a set of edges which are not to be flipped, we arrive at the concept of constrained regular complexes, in Subsection 8.2. These are projections of unique maximal surfaces that are locally convex with respect to a constraining planar straight-line graph. If this graph is connected then pseudo-triangulations are obtained. In contrast, constrained regular *triangulations* do not always exist [13, 50] (or are ambiguous structures, depending on the definition), apart from special cases like the constrained Delaunay triangulation [40, 19]. The algorithmic results from Section 7 carry over to the constrained setting. In Section 9, we show the existence of a convex polytope in n dimensions, whose vertices correspond to all the regular pseudo-triangulations constrained by a given graph. The edges of this polytope represent flips of the admissible types. This implies that the set of constrained regular pseudo-triangulations is connected under admissible flip operations – and generalizes the polytope constructions in [39, 28, 15] (the so-called associahedron, or secondary polytope) for regular triangulations. As a consequence, locally convex surfaces can be deformed into each other in a controlled way using surface flips.

1.2. Basic properties of pseudo-triangulations. We first review some basic notions and properties concerning pseudo-triangulations. For more details, we refer to e.g. [35, 16, 47]. A *(simple) polygon* is a subset of the plane, homeomorphic to a disk and with piecewise-linear boundary. We denote with $vert(P)$ the set of vertices of a polygon P. A *corner* of P is a vertex with internal angle less than π. The other vertices of P are called *non-corners*. A *side chain* of P is the chain of edges between two consecutive corners 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. Note that the corners of ∇ are the vertices of the convex hull $conv(\nabla)$ of ∇. Moreover, for each pair of corners of ∇, the geodesic between them entirely lies on ∇'s boundary, and defines a side chain of ∇.

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). 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; see Figure 1.1 for an illustration. Two faces (pseudo-triangles) may be adjacent at one or two edges. In case of such *double-adjacencies*, the union of the two 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. More precisely, $\mathcal{PT}$ 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. [51]. 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. The reader may verify that the pseudo-triangulation in Figure 1.1 has edge rank 1.

2. Flips in pseudo-triangulations revisited. So-called *flips* are operations of constant combinatorial complexity which are commonly used to modify triangulations [4, 25, 26, 31, 34, 38] and, as of late, pseudo-triangulations [43, 36, 16, 47]. The purpose of the present section is to introduce a novel flip operation, and to show its natural relationship to existing flip types.

2.1. Classical flips. The standard edge flip, also called *Lawson flip* [38], 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 [42, 51] 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 of the geodesic interior to $\nabla_1 \cup \nabla_2$ 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 2.1(a) and (b), where the edge e to be flipped is shown in bold, and the obtained geodesic is drawn dashed. 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. In a minimum pseudo-triangulation, each internal edge can be flipped to a unique different edge. 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 2.1(c).

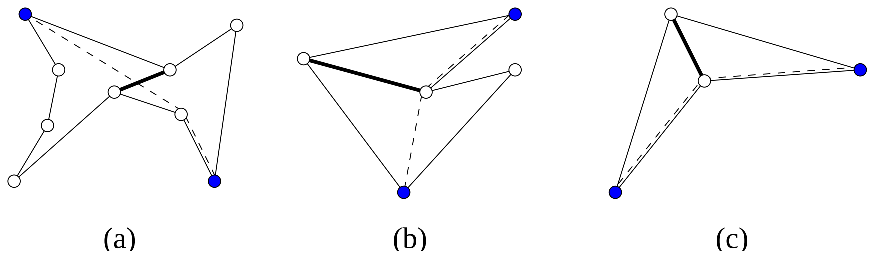

FIG. 2.1. *Exchanging flips and non-flippable edge*

Non-flippable edges have been the source for the theoretically poor behavior of certain flipping algorithms, concerning the flip distance [31, 34] as well as the non-existence of flip sequences [25].

2.2. A novel flip type. We wish to generalize the edge flip so as to cover the situation in Figure 2.1(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 does not exist. 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 each edge-removing flip decreases the edge rank by one, whereas each edge-inserting flip increases it by one. This allows for 'surfing' between pseudo-triangulations of different edge ranks. Various interesting consequences will be discussed, including a reduction of flip distances in Section 3, and a three-dimensional interpretation of the new flip type in Section 6.

2.3. Relation to other types. The new flip type can be used to simulate certain other flip types. For instance, the exchanging flip in Figure 2.1(b) can be simulated by an edge-inserting flip followed by an edge-removing flip. Of particular relevance is a type of flip that arises in the context of Delaunay and regular triangulations; see [46, 25]. 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 [51, 5]. Connect v by geodesics to (at least two) corners of the pseudo-triangle ∇ that v lies in. Each geodesic contributes one edge incident to v, and ∇ is partitioned into (two or three) pseudo-triangles.

Let us simulate the inverse operation, the removal of a degree-3 vertex v, using the new flip type. See Figure 2.2. 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. We consider the simultaneous operation as a single flip – the *vertex-removing* flip. By definition, vertex-removing flips can be applied to vertices of degree 2 only.

In summary, the repertoire of flip operations we are going to use contains the (classical) exchanging flips, the edge-removing and the vertex-removing flip, and their inverses, the inserting flips.

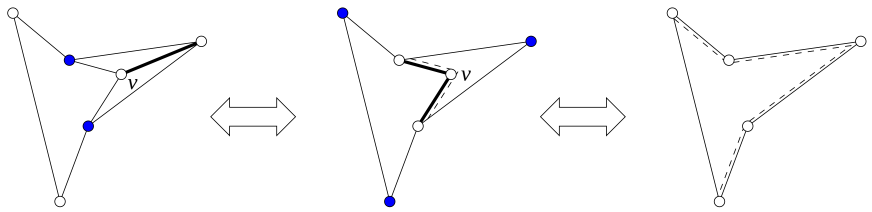

FIG. 2.2. *Edge-removing and vertex-removing flips*

Remarks Edge-removing flips arise implicitly in a greedy flip algorithm for pseudo-triangulations of *convex objects*, in [43]. Certain flips that exchange bitangents of such objects cause an edge removal (or insertion) in the corresponding pseudo-triangulation for the object centers.[1]

Instead of the vertex-removing flip, a different version – namely the exchanging flip in Figure 2.3(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 2.3(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, in Section 6. When change in edge rank is conceded, we may circumvent the flip in Figure 2.3(a) by performing two consecutive flips of the new type, namely an edge-inserting flip followed by an edge-removing flip.

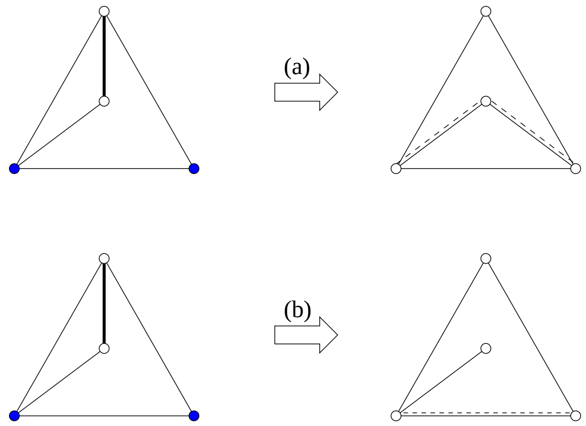

FIG. 2.3. *Ambiguous geodesics interpretation*

The vertex-inserting flip together with the exchanging flips are the constituting operations of the so-called Henneberg construction, which is capable of generating every minimum pseudo-triangulation of a point set S; see e.g. [51]. With our repertoire, arbitrary pseudo-triangulations of S can be derived from each other, without ever changing S. We will elaborate on this issue in detail in Section 3.

[1] We recently learned that edge-removing flips have been introduced independently (and at about the same time) in [41], to generalize the polytope representation in [47] to pseudo-triangulations of arbitrary edge rank. In fact, Theorem 3.4 in Section 3 implies a bound of $O(n \log n)$ on the diameter of their polytope.

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. [31, 34]. The upper bound also applies to exchanging flips in minimum pseudo-triangulations, see [16, 47], but no non-trivial lower bounds are known in this case. This section presents a series of improved upper bounds on flip distances, partially even without using the new flip type.

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 [32] that $\pi_v(P)$ is a tree that partitions P into pseudo-triangles in a unique way.

LEMMA 3.1. *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. □

COROLLARY 3.2. *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 vertex $v \in vert(P)$ 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 of the types above, by Lemma 3.1. □

Corollary 3.2 implies a flip distance of $O(n)$ between any two *pseudo-triangulations* $\mathcal{PT}_1$ and $\mathcal{PT}_2$ of a 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 3.2 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 which case $\mathcal{PT}$ is a triangulation. 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 of Lemma 3.1.

COROLLARY 3.3. *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 3.2, 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 3.4. *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 $n = |S|$. No vertex-removing flips are used.*

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 3.2. 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 4.4(1) in Section 4. Incorporate these segments into $\mathcal{T}_1$ and $\mathcal{T}_2$, respectively, in $O(n)$ flips, which is possible by Corollary 3.3. 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 3.2, another $O(n)$ flips let these two triangulations coincide. □

Remarks Theorem 3.4 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 3.4 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 [34], an instance covered by Corollary 3.2 in $O(n)$ flips.

All flips used in Theorem 3.4 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. [27]. 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.

3.3. Two linear bounds. For the sake of completeness, we include at this place two more implications that come from using our extended repertoire of flips.

Recall that vertex-removing and vertex-inserting flips have been excluded in Theorem 3.4. 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 in Theorem 3.4, 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.

LEMMA 3.5. *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.*

It suffices to prove Lemma 3.5 for full triangulations. Let $\mathcal{T}^*$ be the *star triangulation* of $conv(S)$ with respect to a fixed boundary vertex v. In $\mathcal{T}^*$, each vertex of $conv(S)$ different from v is connected to v by an edge. We show that any triangulation $\mathcal{T}$ of S can be flipped into $\mathcal{T}^*$ with $O(n)$ flips. To this end, consider a particular order for the triangles of $\mathcal{T}$.

LEMMA 3.6. *Let v be some boundary vertex of $\mathcal{T}$. There exists an order $\Delta_1, \ldots, \Delta_t$ for $\mathcal{T}$'s triangles such that $\bigcup_{i \leq j} \Delta_j$ is star-shaped as seen from v, for $j = 1, \ldots, t$.*

Proof. Put $Q_j = \bigcup_{i \leq j} \Delta_j$. To apply induction, assume that Q_j is star-shaped as seen from v, for some j; this assumption is obviously true for $Q_t = conv(S)$. Let Δ be a triangle of $\mathcal{T}$ in Q_j that is not incident to v. The part of Δ on the boundary of Q_j (if existent) defines a visibility angle $\alpha(\Delta)$ at v. Because Q_j is star-shaped and v is a boundary vertex of Q_j, there exists some triangle Δ that lies within $\alpha(\Delta)$. So $Q_j \setminus \Delta = Q_{j-1}$ is star-shaped, too. □

We now visit $\Delta_1, \ldots, \Delta_t$ in the order given by Lemma 3.6, and maintain the star triangulation, T_j^*, of Q_j with respect to v by flipping. If the current triangle Δ_j is incident to v, then no flip is performed. Else Δ_j shares either one or two edges with Q_{j-1}. Let e be such an edge. In the former case, the exchanging flip for e restores T_j^*. To obtain T_j^* in the latter case, the flip removing e is applied first, followed by a flip that removes the unique vertex w internal to Q_j. (Note that the degree of w has been lowered to 2 by the removal of e). This yields $T_t^* = \mathcal{T}^*$ after less than $2t = O(n)$ flips and proves Lemma 3.5.

The next linear bound concerns minimum pseudo-triangulations. 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, 35, 48]. In particular, a minimal pseudo-triangulation may be a full triangulation, even when its vertices are not in convex position [48]. We can show the following:

LEMMA 3.7. *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. No vertex-removing flips are used.*

Proof. Consider some spanning tree B of $\mathcal{PT}$. We construct a connected graph Z with vertex set S, that does not cross B and where each vertex is pointed. Let $v_1, \ldots, v_n$ be the vertices of S in some depth-first order for B such that, for each v_i, the edges of B incident to v_i are visited in clockwise order. To obtain Z, connect v_i and v_{i+1} for $i = 1, \ldots, n$ and $v_{n+1} = v_1$, as follows. If $v_i v_{i+1}$ is an edge of B then use this edge. Otherwise, use the geodesic that comes from 'stretching' the unique path in B between v_i and v_{i+1} to the shortest path in the plane that does not cross B. By construction, each vertex v_i either has degree 2 in Z or some geodesics run via v_i, in which case v_i is pointed in Z as well.

In general, $\mathcal{PT}$ will not contain Z as a subgraph. We flip $\mathcal{PT}$ into a pseudo-triangulation $\mathcal{PT}'$ that contains Z (and B) as follows. Imagine cutting $conv(S)$ along B and moving it apart at the cuts infinitesimally. This yields a partition of $conv(S)$ into simple polygons. Within each such polygon, $\mathcal{PT}$ can be flipped to contain the relevant part of Z, in a total of $O(n)$ flips for all polygons; see Subsection 3.1. Finally, we flip $\mathcal{PT}'$ to a minimum pseudo-triangulation that contains Z, by cutting $conv(S)$ along Z and adapting the pseudo-triangulations within these polygons, using $O(n)$ additional flips. ☐

Remarks The order in Lemma 3.6 is called a shelling order (or a translation order) for $\mathcal{T}$ with respect to v. It is well known that such an order does not exist for every triangulation if v is an *internal* vertex. Exceptions are Delaunay triangulations [21] and regular triangulations [22].

Lemma 3.7 does not imply a linear flip distance between any two pseudo-triangulations of S, because it depends on $\mathcal{PT}$ which minimum pseudo-triangulations can be reached. The pointed graph Z arises in a classical factor-2 approximation for the travelling salesman tour for S. We raise the question of whether every triangulation of a point set S contains a pointed spanning tree.

3.4. Minimum pseudo-triangulations. Our next aim is to provide a stronger version of Theorem 3.4, namely for minimum pseudo-triangulations and when using the exchanging flip type exclusively. That is, we restrict ourselves to staying within the class of minimum pseudo-triangulations. We apply arguments similar to Subsections 3.1 and 3.2.

Let P be a polygon with k corners. 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 $\pi_c(P)$ defined in Subsection 3.1. The proof of Lemma 3.1 now can be adapted, to show that every minimum pseudo-triangulation of P can be transformed into $\mu_c(P)$ by at most $k-3$ exchanging flips. The new flip type is not used here, because each internal edge is flippable in the classical sense. We obtain:

LEMMA 3.8. *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 next lemma is a variant of Corollary 3.3, for minimum pseudo-triangulations.

LEMMA 3.9. *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 modified to contain G by applying $O(nj)$ exchanging flips, if S has n vertices and $G \setminus P$ has j edges.*

Proof. Let uv be some edge of $G \setminus P$. We prove that uv can be incorporated into $\mathcal{MPT}$ using $O(n)$ exchanging flips, and without altering previously incorporated edges of G.

If vertex u is not pointed in $\mathcal{MPT} \cup \{uv\}$, we flip edges of $\mathcal{MPT}$ incident to u to restore the pointedness of u, as follows. Let $m(u)$ denote the angle formed by u's edges in the current minimum pseudo-triangulation, and let $g(u)$ be the angle formed by those edges among them which also belong to G (if there are any). We have $g(u) < m(u) < \pi$. Rotate $m(u)$, by flipping the unique edge e that bounds $m(u)$ but not $g(u)$ and such that e and $g(u)$ lie on different sides of uv, until $uv \in m(u)$. The condition is necessarily met as soon as e does not exist. This process does not alter any edge of G in $\mathcal{MPT}$, takes $O(n)$ exchanging flips, and transforms $\mathcal{MPT}$ into a minimum pseudo-triangulation $\mathcal{MPT}'$. Note that the vertex v may be non-pointed in $\mathcal{MPT}' \cup \{uv\}$.

We next flip the edges of $\mathcal{MPT}'$ which cross uv (if there are any), in the order from u to v. For the current edge e to be flipped, the following invariants hold. Edge e bounds a face ∇ incident to u. Denote with c the corner of ∇ opposite to e. ($c = u$ is possible.) The edge f obtained by flipping e lies on some geodesic to c. Thus f leaves u pointed. Moreover, as the line segment uv also lies on a geodesic to c, edge f does not cross uv. After $O(n)$ flips, a minimum pseudo-triangulation $\mathcal{MPT}''$ results, where no edge crosses uv, and such that u is still pointed in $\mathcal{MPT}'' \cup \{uv\}$.

Now either $uv \in \mathcal{MPT}''$ (and we are done) or v is not pointed in $\mathcal{MPT}'' \cup \{uv\}$. In the latter case, there is a unique face ∇'' of $\mathcal{MPT}''$ which contains uv in its interior, and v is a non-corner of ∇''. Consider the side chain of ∇'' that contains v, and let c'' be the corner of ∇'' opposite to this side chain. Flip an edge of ∇'' which is incident to v and which does not lie in the pointed angle of v in G. The obtained edge f'' lies on a geodesic to the corner c''. Therefore, f'' neither crosses uv nor destroys the pointedness of u. If $f'' = uv$ the construction is complete. Otherwise, the degree of v has been decremented in the flip, and we repeat the step above for the face that now contains uv in its interior, until uv becomes part of the current pseudo-triangulation. This takes another $O(n)$ flips. □

Lemma 3.9 directly implies an $O(n^2)$ bound on the (exchanging) flip distance in minimum pseudo-triangulations: Take for G the minimum pseudo-triangulation that $\mathcal{MPT}$ is to be transformed to. The proof of the theorem below shows that Lemma 3.9 can be used in a more sophisticated way.

THEOREM 3.10. *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 4.4(2) from Section 4. This constructs a polygonal partition Π of P whose vertex set is S, and where all vertices are pointed. Π is obtained by introducing $O(\log n)$ line segments spanned by S in each recursive step, and the number of recursive steps is $O(\log n)$.

By Lemma 3.9, $O(n \log^2 n)$ exchanging flips are sufficient to make $\mathcal{MPT}_1$ and $\mathcal{MPT}_2$ contain all the edges of Π. Finally, Lemma 3.8 allows for adapting pseudo-triangles within the polygons of Π in $O(n)$ such flips. □

Remarks Theorem 3.10 improves the recent bound of $O(n^2)$ in [16] for minimum pseudo-triangulations of point sets. Again, we conjecture that the truth is $O(n)$ flips.[2]

In [47], 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 3.10, the diameter of $\mathcal{M}(S)$ is bounded by $O(n \log^2 n)$.

As a somewhat counter-intuitive fact, the transformation between two given minimum pseudo-triangulations may be speeded up by using edge-inserting and edge-removing flips intermediately. There exist examples where $n - 3$ exchanging flips are

[2] By a very recent result in [12], 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.

necessary, a third of which is saved when performing a single edge-inserting and edge-removing flip, respectively, in the beginning and at the end; see [33]. 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.

A *constrained* pseudo-triangulation of a point set S is one that is required to contain a given planar straight-line graph G on S. It is well known that, for any choice of G, the set of constrained triangulations of S is connected by Lawson flips: Every triangulation constrained by G can be flipped into the constrained Delaunay triangulation for G and S; see [40]. The following related result holds, by arguments in the proof of Lemma 3.9.

COROLLARY 3.11. *Any two minimum pseudo-triangulations of a point set S constrained by a given (pointed) planar straight-line graph G on S can be flipped into each other using $O(n^2)$ exchanging flips.*

In case G partitions $conv(S)$ into augmented polygons (as in Figure 8.1), the $O(n \log^2 n)$ exchanging flip distance bound from Theorem 3.10 applies. Moreover, for such graphs G, the flip distance between constrained pseudo-triangulations of arbitrary edge rank is $O(n \log n)$ when edge-inserting flips and their reverses are admitted, by Theorem 3.4. The concept of constrained pseudo-triangulations plays an important role in the context of regularity, see Sections 8 and 9.

4. Partitioning results. This section presents some partitioning results concerning pseudo-triangulations, which have been referred to in Section 3. Subsections 4.1 and 4.3 give two technical lemmas. The assertions in Subsections 4.2 and 4.4 might be of separate interest.

4.1. Pseudo-convex polygons. 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.

LEMMA 4.1. *Let P be a strongly pseudo-convex polygon with k corners. There exists a nice pseudo-triangle for P with cut $O(\log k)$.*

Proof. Observe first that there always exists *some* nice pseudo-triangle ∇ for P. Select a corner w of P. Since P is strongly pseudo-convex, w 'sees' a unique side chain of P. Together with w, the two corners that define this side chain span a nice pseudo-triangle ∇. If ∇ has cut $O(\log k)$ then we are done. Else we apply induction, by constructing a pseudo-triangle whose cut is strictly smaller, as follows.

We will prove that (*) at each edge d of ∇ which is a diagonal of P, there exists an adjacent pseudo-triangle ∇', spanned by d and some corner c of P, such that ∇' has cut at most $1 + \log_2 k$. Property (*) is used as follows. Put $x = \log_2 k$, and assume that some side chain of ∇ contains at least $2x + 5$ diagonals of P. Choose a middle-most diagonal, d, and generate an adjacent pseudo-triangle ∇' as above. Flip edge d, which creates two pseudo-triangles. Each of them avoids at least $x + 2$ diagonals that bound ∇, and contains at most $x + 1$ diagonals that bound ∇'. So for both, the cut is strictly smaller than the cut of ∇. Moreover, because P is strongly pseudo-convex, exactly one created pseudo-triangle is nice.

It remains to prove (*). We construct the desired pseudo-triangle ∇' as follows; see Figure 4.1. Let $d = uv$, and denote with Q the part of P cut off by d and not containing ∇. Note that u and v are corners of Q. Consider the minimum shortest-path tree of Q with root u, as defined in Subsection 3.4. Let ∇_u be the unique face of this tree incident to uv. (Here and below, define things for v analogously.)

Case (1) Q contains only one corner $c \neq u, v$. Then $\nabla_u = \nabla_v$, and we take this face for the pseudo-triangle ∇' in question. No diagonals of P lie on the side chains uc and vc of ∇'.

Case (2) Q contains at least two corners $\neq u, v$. Then ∇_u and ∇_v differ in their third corner, call it c_u and c_v respectively, and thus have diagonals of P on their boundaries. Let d_u be the first diagonal on the side chain uc_u of ∇_u, and let d_u end at the vertex n_u. P is strongly pseudo-convex, so n_u is a non-corner of Q, because d_u would be a corner tangent, otherwise. Also, n_u does not lie on the side chain of Q that starts at v and is different from uv, because non-convex geodesics inside P and via u and n_u would exist, otherwise. Therefore, the geodesic from n_u to v starts with a diagonal D of P. Observe next that, by $\nabla_u \neq \nabla_v$, their side chains uc_u and vc_v cross. Hence the two diagonals d_u and d_v cross unless $n_u = n_v$. In both cases at least one diagonal, say d_u, cuts off from Q at least half of its corners. Therefore the diagonal D above, that is adjacent to d_u at vertex n_u, cuts off $k' \leq \frac{k}{2}$ corners. Clearly, $k' \geq 1$, and we recur on the polygon cut off by D (shaded in Figure 4.1). The diagonal d_u lies on the boundary of the pseudo-triangle ∇' being constructed.

Each occurrance of Case (2) increments the cut of ∇', which initially is 1 because of the diagonal $d = uv$. By the choice of D, Case (2) occurs at most $\log_2 k$ times. This proves (*) and completes the argumentation. □

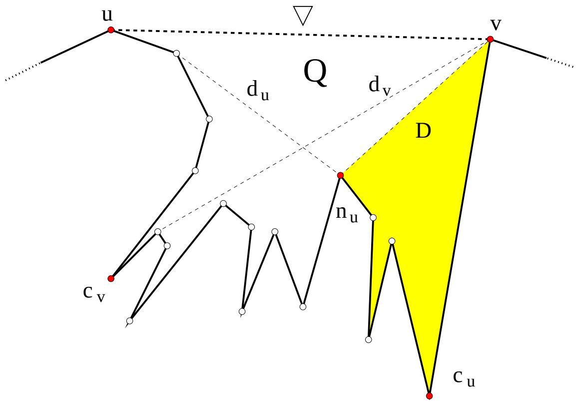

FIG. 4.1. *Constructing* ∇'

4.2. Pseudo-triangulations with small cut. We are now ready to prove the following structural result for minimum pseudo-triangulations of simple polygons.

THEOREM 4.2. *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 4.1. 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 4.1 or is among the at most 6 edges incident to some corner of f. □

Remarks Theorem 4.2 is asymptotically optimal. There exist polygons with n vertices where every possible minimum pseudo-triangulation contains some face with cut $\Omega(\log n)$; see [6]. The theorem is related to a result in [35] which shows, for every point set S, the existence of a minimum pseudo-triangulation with constant face complexity. Another related result, in [18], shows that every n-gon P admits a minimum pseudo-triangulation $\mathcal{MPT}$ such that each line segment internal to P crosses only $O(\log n)$ edges of $\mathcal{MPT}$.

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

LEMMA 4.3. *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 $\frac{2i}{3}$ points of M. We claim that, for each pair c, c' of corners of ∇, there exist at least $\frac{2i}{3}+1$ 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 $2i+3$ 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 $\frac{2i}{3}$ points each. □

Remarks The fraction $\frac{2}{3}$ in Lemma 4.3 is optimal, even when ∇ 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.

4.4. Partition theorem. The results in the preceeding subsections combine to the following partition theorem for augmented polygons.

THEOREM 4.4. *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 4.3 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 4.2. 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 4.3 that may have to be used to split ∇. □

Remarks Theorem 4.4 is similar in flavor to a result in [17], which asserts that any simple n-gon can be split by a diagonal into two subpolygons with at most $\frac{2}{3}n$ vertices.

5. Surfaces for pseudo-triangulations. A *polyhedral surface* is the graph of a continuous and piecewise-linear function whose domain is the plane, or a polygonal subset thereof. Polyhedral surfaces are also known as polyhedral terrains. A two-dimensional cell complex $\mathcal{C}$ is called *projective* if there exists some polyhedral surface whose set of non-linearity projects exactly to the set of all (internal) edges of $\mathcal{C}$. A subclass are the *regular* complexes, which are the projections of *convex* surfaces, i.e., whose defining function is convex. Examples of regular two-dimensional complexes are line arrangements, Voronoi diagrams and power diagrams (even their generalizations to higher order), and Delaunay triangulations; see e.g. [9] and references therein. Not every triangulation, however, is regular (as Figure 7.3 shows) but it is trivial that triangulations are projective.

We intend to show, in Subsection 5.2, that pseudo-triangulations have natural realizations as polyhedral surfaces. We characterize the class of projective pseudo-triangulations in Subsection 5.3, and discuss regular pseudo-triangulations and their constrained variants in later sections.

The underlying domain is – as in previous sections – an augmented polygon, not only to gain generality compared to the point set case, but also because certain questions concerning a three-dimensional realization become more natural and interesting in this setting.

5.1. Complete vertices. We start with introducing a *status* for vertices, which will allow us to determine which vertices can be fixed in height when lifting a given pseudo-triangulation to three dimensions.

Consider an augmented polygon (P, S), and let $\mathcal{PT}$ be a pseudo-triangulation of (P, S). The status of a vertex v of $\mathcal{PT}$ is *complete* if v is a corner in each of its incident pseudo-triangles. The status of v is *incomplete*, otherwise. Note that each incomplete vertex uniquely corresponds to a pseudo-triangle of $\mathcal{PT}$ where this vertex is a non-corner. Incomplete vertices have to be pointed, whereas complete vertices are pointed if and only if they are corners of P. In a (full) triangulation, all vertices are complete. On the other end of the spectrum, every minimum pseudo-triangulation of (P, S) realizes the minimum number of complete vertices.

Observe that no vertex status is affected by an exchanging flip, whereas an edge-removing flip alters the status of exactly one vertex from complete to incomplete. A vertex-removing flip does not alter the status of any remaining vertex. Making all vertices complete by edge-inserting flips always results in a triangulation.

Remarks We believe that the *completeness* of a vertex is a concept being more natural than *pointedness*. For example, the vertices of $conv(S)$ – though being pointed – behave like non-pointed vertices in the flipping process, because no geodesic ever runs via them. This effect becomes even more apparent when the underlying domain (P, S) is not convex. Corners as well as non-corners of the boundary polygon P may be pointed in a pseudo-triangulation $\mathcal{PT}$ of (P, S), but the (forced) pointedness of P's corners seems artificial. What actually counts is '(non-)pointedness with boundary effects eliminated', that is, the completeness status of a vertex.

The number of incomplete vertices is $i = p - k$, if P has k corners and $\mathcal{PT}$ has p pointed vertices. This relation saves one of the four parameters in the formulas for the number of edges and faces in a pseudo-triangulation of (P, S), see Subsection 3.2 – another fact in favor of the new concept.

Complete vertices can be distinguished from incomplete ones in any polygonal partition, as they are defined via corners of polygonal faces. In fact, all the results in Sections 2, 3, and 4 could have been derived using this terminology. We refrained from doing so only to keep to existing notation in related work. In the three-dimensional setting below, the notion of complete vertices is indispensible, however.

5.2. Surface theorem. In this subsection, we state and prove a main geometric result of this paper.

THEOREM 5.1. *(Surface Theorem) Let (P, S) be an augmented polygon, and let $\mathcal{PT}$ be any pseudo-triangulation thereof. Let $\boldsymbol{h}$ be a vector assigning a height to each complete vertex of $\mathcal{PT}$. For each choice of $\boldsymbol{h}$, there exists a unique polyhedral surface $\mathcal{F}$ above the domain P, that respects $\boldsymbol{h}$ and whose edges project vertically to (a subset of) the edges of $\mathcal{PT}$.*

Proof. Let $v_1, \ldots, v_r$ denote the complete vertices of $\mathcal{PT}$, and let $v_{r+1}, \ldots, v_n$ be the incomplete ones. We set up a system of linear equations for the vertex heights $z_1, \ldots, z_n$ that determines a polygonal surface $\mathcal{F}$ and has a unique solution.

If v_i is a complete vertex, the system contains the equation $z_i = h_i$. Otherwise, its height z_i is constrained as follows. There is a unique pseudo-triangle ∇ in $\mathcal{PT}$ where v_i is no corner; an internal angle larger that π occurs in ∇ at v_i. Let v_j, v_k, v_ℓ be the corners of ∇. We have $v_i \in conv(v_j, v_k, v_\ell) = conv(\nabla)$ by definition of a pseudo-triangle. Requiring coplanarity of the lifted vertices v_i, v_j, v_k, v_ℓ leads to an equation of the form

$$\alpha(v_j v_k v_\ell) \cdot z_i - \alpha(v_i v_k v_\ell) \cdot z_j - \alpha(v_i v_j v_\ell) \cdot z_k - \alpha(v_i v_j v_k) \cdot z_\ell = 0$$

where $\alpha(abc)$ denotes the (unsigned) area of a triangle abc. This expresses z_i as a convex combination of z_j, z_k, and z_ℓ.

A linear system $A \cdot \boldsymbol{z} = \boldsymbol{b}$ is obtained, where A is an $n \times n$ matrix whose i-th row corresponds to the equation for z_i, and where $\boldsymbol{b}$ is a vector that coincides with $\boldsymbol{h}$ in the first r entries and contains zero entries, otherwise. Any solution $\boldsymbol{z}$ defines a valid surface $\mathcal{F}$, because triangular facets of $\mathcal{F}$ are trivially fixed by three heights, and for each pseudo-triangular facet of $\mathcal{F}$ its vertices are forced to be coplanar. Thus $\mathcal{F}$ is piecewise-linear, continuous, and each edge of $\mathcal{F}$ projects to an edge of $\mathcal{PT}$.

We now prove that a unique solution $\boldsymbol{z}$ exists, by showing $det(A) \neq 0$. Call an $n \times n$ matrix M *affirmative* if (1) M is positive exactly at its diagonal entries, (2) no row sum in M is negative, and (3) at least one row sum is positive. M is called *strictly affirmative* if each allowable submatrix of M (i.e., which comes from deleting the i-th row and the i-th column of M, for $t < n$ indices i) is affirmative. It is easy to prove $det(M) > 0$ for any strictly affirmative matrix M; see [10].

It remains to show that A is strictly affirmative. Observe first that A is affirmative: the first $r \geq 1$ rows of A are unit rows, and the remaining $n - r$ rows represent convex combinations for the diagonal elements. Let A' be an allowable submatrix of A. Then Properties (1) and (2) clearly hold for A'. Property (3) is obvious while unit rows (that correspond to complete vertices) are present in A', and is retained because heights for complete vertices participate with negative coefficients in remaining rows. We conclude that A is strictly affirmative. ☐

Remarks The set of edges of the surface $\mathcal{F}$ in Theorem 5.1 may project to a proper subset of edges of the pseudo-triangulation $\mathcal{PT}$ of (P, S). We will discuss this phenomenon in detail in Subsection 5.3.

The existence of *some* projection surface is obvious when (P, S) contains no internal vertices, that is, if $S = vert(P)$. Then all internal edges of $\mathcal{PT}$ are diagonals of P. For each pseudo-triangle of $\mathcal{PT}$, its three corners define a plane, which can be placed appropriately in space. Note that these planes can be chosen such that the resulting surface is convex at each internal edge; this observation leads to Corollary 9.2 in Section 9.

Theorem 5.1 holds for arbitrary (augmented) polygonal domains, including polygons with holes. We will use this generalization in Subsections 6.3 and 8.2. Note that the non-corners of holes are corners of the domain, and thus are complete vertices.

The vector $\boldsymbol{b}$ is an eigenvector of the transpose A^T with eigenvalue 1, that is, $A^T \cdot \boldsymbol{b} = \boldsymbol{b}$. In fact, A^T (and thus A) has r eigenvalues of 1.

Interestingly, a matrix of the form similar to A arises in the context of power diagrams and Gale transforms [10]. Moreover, the polytope representation of pseudo-triangulations we are going to describe in Section 9, could alternatively be obtained using Gale transform techniques; cf. [15]. A systematic use of Gale transforms may shed additional light into the geometric properties of pseudo-triangulations.

By applying allowable row operations to A, it can be shown that its inverse A^{-1} contains only non-negative entries. This implies a monotonicity property which will be important later.

LEMMA 5.2. *The solution $\boldsymbol{z}$ of the system $A \cdot \boldsymbol{z} = \boldsymbol{b}$ is a monotone (linear) function of $\boldsymbol{b}$, and thus of the height vector $\boldsymbol{h}$.*

The Maxwell-Cremona theorem [20] asserts that, for a given planar straight-line graph G, there exists a (certain) projection surface if and only if G admits a so-called stress, i.e., an assignment of a real-valued tension to each edge of G such that all vertices of G are in equilibrium state. In this projection surface, the outer facet is required to be planar (unlike in the surface in Theorem 5.1). For *convex* augmented polygons (P, S), the Maxwell-Cremona theorem combines with Theorem 5.1, namely when height zero is chosen for each (necessarily complete) boundary vertex of the convex polygon P. This gives the corollary below, which also follows from recent results in [41].

COROLLARY 5.3. *Let $\mathcal{PT}$ be a pseudo-triangulation of a planar point set S. There exists a non-trivial stress for the edges of $\mathcal{PT}$, provided that $\mathcal{PT}$ is not minimum. More generally, the dimension of the space of stresses for $\mathcal{PT}$ equals the number of its complete internal vertices.*

We mention that Theorem 5.1 also implies a variant of Corollary 5.3 for arbitrary augmented polygons (P, S), asserting that every pseudo-triangulation of (P, S) admits a non-trivial stress for the set $S \setminus vert(P)$ of *internal* vertices, unless P is a pseudo-triangle.

5.3. Projective pseudo-triangulations. The surface $\mathcal{F}$ in Theorem 5.1 may fail to be strictly convex or reflex for some edges of $\mathcal{PT}$. In this case, $\mathcal{F}$ yields a proper subset of edges of $\mathcal{PT}$ in the projection. One reason is that the height vector $\boldsymbol{h}$ is degenerate. For instance, if $\boldsymbol{h} = 0$ then $\mathcal{F}$ contains a single facet that projects to the polygon P in the plane, regardless of the shape of $\mathcal{PT}$. On the other hand, pseudo-triangles in double-adjacency give rise to a single pseudo-triangular facet of $\mathcal{F}$, for every choice of $\boldsymbol{h}$.

Let us call a pseudo-triangulation $\mathcal{PT}$ of (P, S) *stable* if no subset of incomplete

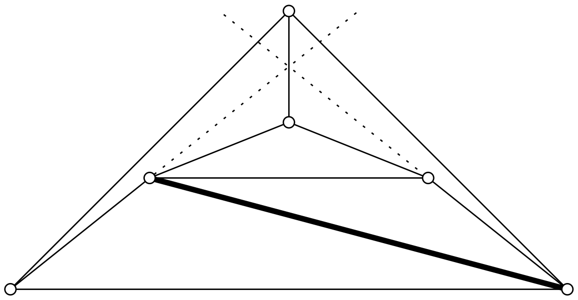

FIG. 5.1. *A stable though non-projective pseudo-triangulation*

vertices in $S \setminus vert(P)$ can be eliminated (along with their incident edges) such that (1) a valid pseudo-triangulation $\mathcal{PT}'$ remains, and (2) the status of each vertex of $\mathcal{PT}'$ is the same as in $\mathcal{PT}$. In particular, double-adjacencies are outruled in a stable pseudo-triangulation. Trivial stable instances are full triangulations, and pseudo-triangulations of (P, S) for $S = vert(P)$. To see an example of a pseudo-triangulation which is not stable, omit the three bold edges in Figure 7.3 which do not belong to the internal triangle. Interestingly, no double-adjacencies occur in this example.

THEOREM 5.4. *A pseudo-triangulation $\mathcal{PT}$ of (P, S) is projective only if $\mathcal{PT}$ is stable. If $\mathcal{PT}$ is stable then the point set S can be perturbed (by some arbitrarily small ε) such that $\mathcal{PT}$ becomes projective.*

Proof. Assume first that $\mathcal{PT}$ is not stable. Eliminate some set of incomplete vertices from $S \setminus vert(P)$ such that a pseudo-triangulation $\mathcal{PT}'$ of (P, S) is obtained where all vertices agree in status with their counterparts in $\mathcal{PT}$. Fix an arbitrary height vector $\boldsymbol{h}$ for S. By Theorem 5.1, there exist unique surfaces $\mathcal{F}$ for $\mathcal{PT}$ and $\mathcal{F}'$ for $\mathcal{PT}'$. Since only incomplete vertices have been eliminated, and the status of no remaining vertex has been changed, we have $\mathcal{F}' = \mathcal{F}$. Now let ∇ be a pseudo-triangle of $\mathcal{PT}'$ but not of $\mathcal{PT}$. The restriction of $\mathcal{F}'$ to ∇ is a pseudo-triangular facet. Because $\mathcal{F}' = \mathcal{F}$, the restriction of $\mathcal{F}$ to ∇ is planar, too. This implies that $\mathcal{PT}$ is not projective.

Now let $\mathcal{PT}$ be a stable pseudo-triangulation. Let $\boldsymbol{h}$ be a height vector for S with paiwise different entries. Assume that the surface $\mathcal{F}$ for $\mathcal{PT}$ and $\boldsymbol{h}$ is planar at some edge uv. Let c and c' be the two corners opposite to uv in the two pseudo-triangular facets of $\mathcal{F}$ adjacent at uv. By assumption, the vertex v lies in the plane through u, c, and c'. Suppose first that v is a complete vertex. As the heights of u, c, and c' also depend on vertices different from v, we can displace v by some arbitrarily small ε in an appropriate direction, such that $\mathcal{F}$ becomes non-planar at uv. Now suppose that v is incomplete. Because $\mathcal{PT}$ is stable, the height of v does not depend on the same set of (three) vertices as do the heights of u, c, and c'. Otherwise, v could be eliminated (possibly jointly with other vertices) without changing $\mathcal{F}$, a contradiction to the assumed stability of $\mathcal{PT}$. So, again, v can be displaced to make $\mathcal{F}$ non-planar at uv. This implies that S can be perturbed such that $\mathcal{PT}$ becomes projective. ◻

Figure 5.1 illustrates a stable pseudo-trianglation which fails to be projective. Neither of the two incomplete vertices can be eliminated without changing the status of the upmost internal vertex from complete to incomplete. On the other hand, as the prolongations (shown dotted) of two edges meet at a third edge, the edge drawn in bold flattens out in all possible surface realizations. Note that a slight movement of any single vertex restores the projectivity in this example.

By Theorem 5.4, we restrict our attention – in the remainder of this paper – to vertex sets where every stable pseudo-trianglation is also projective. For such a vertex set S and a pseudo-triangulation $\mathcal{PT}$ thereof, a height vector $\boldsymbol{h}$ is called *generic* if $\boldsymbol{h}$ witnesses the projectivity of $\mathcal{PT}$.

Remarks Note the difference between the concepts of *stable* and *minimum* pseudo-triangulation. In the latter, no single edge (but possibly some vertices) can be eliminated such that a valid pseudo-triangulation remains. There is no obvious relation between stable and minimum pseudo-triangulations, but the following can be observed. If $conv(S)$ is a triangle and S contains vertices internal to P, then no minimum pseudo-triangulation $\mathcal{MPT}$ of (P, S) is stable. All internal vertices are incomplete, and their elimination leaves a valid pseudo-triangle, P, without changing the status of any vertex of P. By Theorem 5.4, this implies that $\mathcal{MPT}$ is not projective.

6. Surface interpretation of flips. We next provide a surface interpretation for all the admissible flip operations in Section 2. To this end, consider some projective pseudo-triangulation $\mathcal{PT}$ of an augmented polygon (P, S). (Recall from Subsection 5.3 that 'projective' is equivalent to 'stable', by our convention on the vertex set S.) Let $\mathcal{PT}'$ be the unique pseudo-triangulation obtained from $\mathcal{PT}$ by applying a single admissible flip. Assign height vectors $\boldsymbol{h}$ and $\boldsymbol{h}'$ to $\mathcal{PT}$ and $\mathcal{PT}'$, such that $\boldsymbol{h}$ is generic for $\mathcal{PT}$ and coincides with $\boldsymbol{h}'$ for all vertices being complete in both structures. By Theorem 5.1, there exist two unique projection surfaces $\mathcal{F}$ and $\mathcal{F}'$ for $\mathcal{PT}$ and $\mathcal{PT}'$. We call the operation that transforms $\mathcal{F}$ into $\mathcal{F}'$ a *surface flip.*

In a combinatorial sense, surface flips cause local and constant-size changes. Geometrically, however, these operations are by no means local. A single flip may change the surface heights for many vertices, even 'far away' from the region where the combinatorial change takes place. The change in height depends on the structure of the entire surface. Thus general surface flips are *context-sensitive* operations.

For expository reasons, we will first study surface flips in their *context-free* version, in the next two subsections. The following notation will be used. Consider an exchanging or an edge-removing flip. ∇_1 and ∇_2 are the two adjacent pseudo-triangles involved in the flip, and R denotes their union. We restrict the corresponding surface flip to R, such that $\mathcal{PT}$ and $\mathcal{PT}'$ are the two (projective) pseudo-triangulations of R before and after the flip, and $\mathcal{F}$ and $\mathcal{F}'$ are the corresponding unique surfaces. The operation that takes $\mathcal{F}$ into $\mathcal{F}'$ constitutes the context-free surface flip. We come back to its context-sensitive version in Subsection 6.3.

6.1. Convexifying flips. Suppose the flip in question is an exchanging flip. The following combinatorial change is caused. The flip replaces the edge $e = \nabla_1 \cap \nabla_2$ by the edge e' that is the intersection of the two pseudo-triangles $\mathcal{PT}'$ consists of. Note that ∇_1 and ∇_2 are in no double-adjacency, by our assumption that $\mathcal{PT}$ is projective. R has exactly four corners which are the vertices v_i being complete in both of $\mathcal{PT}$ and $\mathcal{PT}'$. Their heights have been fixed as $z_i = h_i$, in a non-coplanar manner. The non-corners of R are incomplete in both of $\mathcal{PT}$ and $\mathcal{PT}'$. If they change incidence to pseudo-triangles during the flip, then the planarity conditions for their heights are affected. For such a vertex v_j the change is from $z_j = \gamma_j$ to $z_j = \gamma'_j$, where γ_j and γ'_j are convex combinations of the heights of the corners in the respective pseudo-triangles.

To view the change geometrically, see Figure 6.1. Consider the two facets f_1 and f_2 of the surface $\mathcal{F}$ before the flip, and assume that $\mathcal{F}$ is reflex at edge e. Then exactly one corner of f_1 lies below the plane through f_2, and vice versa. As corners do not

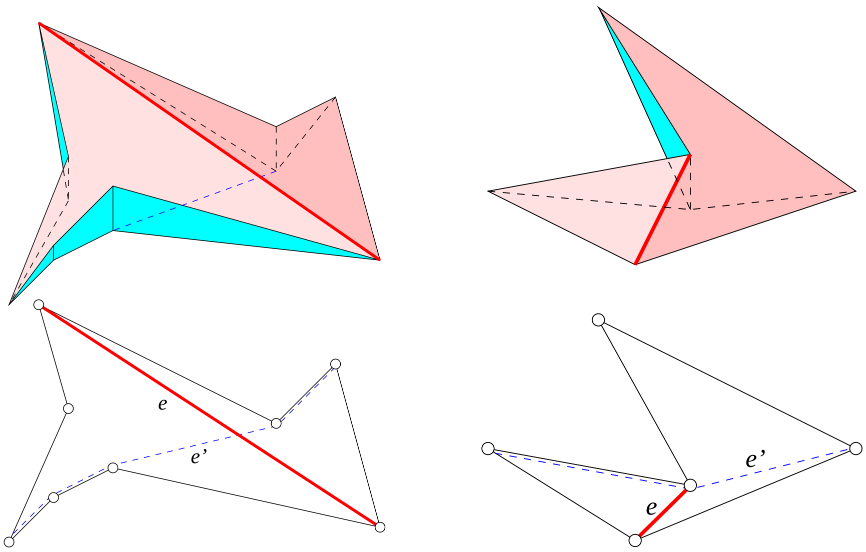

FIG. 6.1. *Two convexifying flips*

change height in the flip, this implies that the obtained surface $\mathcal{F}'$ is convex at edge e'. We will therefore use the term *convexifying* flip. $\mathcal{F}$ and $\mathcal{F}'$ bound a polyhedron from above and below, respectively. This polyhedron is a tetrahedron if f_1 and f_2 are triangles, and the flip degenerates to the Lawson flip for triangulations. In this case, we will also talk of a *tetrahedral* flip – a special case well known from flipping in triangular surfaces; see e.g. [38].

6.2. Planarizing flips. We discuss the edge-removing flip next. Combinatorially, the following change takes place. $\mathcal{PT}$ consists of two pseudo-triangles ∇_1 and ∇_2 in single adjacency. The flip removes their common edge e and joins ∇_1 and ∇_2 to the pseudo-triangle R which thus represents $\mathcal{PT}'$. The three corners of R are complete in both of $\mathcal{PT}$ and $\mathcal{PT}'$. So their heights are given explicitly. One vertex v_e incident to e changes its status from complete in $\mathcal{PT}$ to incomplete in $\mathcal{PT}'$. The condition for its height changes from $z_e = h_e$ to $z_e = \gamma_e'$, where γ_e' expresses coplanarity with the lifted corners of R. All other vertices v_j of R are incomplete in both of $\mathcal{PT}$ and $\mathcal{PT}'$. Their height constraints change from $z_j = \gamma_j$ to $z_j = \gamma_j'$ where γ_j expresses coplanarity with the lifted corners of ∇_1 or ∇_2.

Geometrically, the surface $\mathcal{F}$ (which we assume to be reflex at edge e again) is flipped to $\mathcal{F}'$ as is depicted in Figure 6.2. $\mathcal{F}'$ is a single pseudo-triangular facet and thus is planar. We call such a flip a *planarizing* flip. Again, $\mathcal{F}$ and $\mathcal{F}'$ bound a polyhedron from above and below, respectively.

6.3. Context-sensitive view. The context-free flips described in the previous two subsections leave the heights of the complete vertices of $\mathcal{PT}'$ unchanged. The incomplete vertices of $\mathcal{PT}'$ are lowered. Complete vertices may, however, be incomplete when being embedded in a larger pseudo-triangulation. A surface flip – in its general context-sensitive form – may alter their heights as well. This alteration is uniquely described by the corresponding two linear systems. No alteration occurs in a convexifying flip which is tetrahedral, because the system remains unchanged.

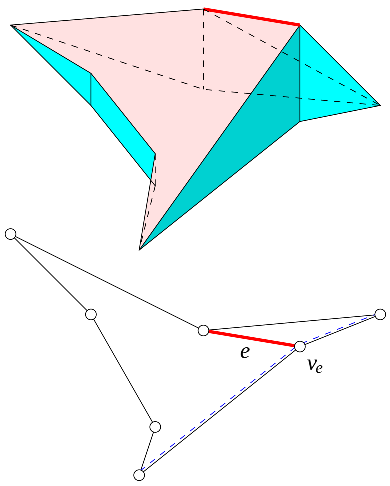

FIG. 6.2. *Planarizing flip*

LEMMA 6.1. *Let $\mathcal{F}'$ be obtained from the surface $\mathcal{F}$ by applying a (context-sensitive) convexifying or planarizing surface flip φ. Then $\mathcal{F}'$ nowhere lies strictly above $\mathcal{F}$.*

Proof. We show that φ can be simulated by a sequence of flips, each lowering certain parts of the surface. Let (P, S) and $\mathcal{PT}$, respectively, be the augmented polygon and the pseudo-triangulation in question. Thus $\mathcal{F}$ is the surface for $\mathcal{PT}$ before flipping. Consider the polygon $R \subseteq P$ the context-free version of φ is restricted to, and denote this flip by φ_R.

Introduce new edges to $\mathcal{PT}$ such that a pseudo-triangulation, $\mathcal{PT}_1$, is obtained in which all corners of R are complete vertices. Assign heights as determined by $\mathcal{F}$ to the complete vertices of $\mathcal{PT}_1$. The surface that corresponds to $\mathcal{PT}_1$ and these heights is still $\mathcal{F}$. Now perform the flip φ_R and denote the obtained pseudo-triangulation by $\mathcal{PT}_2$ and its surface by $\mathcal{F}_2$. This lowers the non-corners of R (which are incomplete in $\mathcal{PT}_2$) but leaves the heights of R's corners fixed. To see that no other incomplete vertex of $\mathcal{PT}_2$ can gain height, let us cut out the region R from P and consider it as a hole. The restrictions of $\mathcal{PT}_1$ and $\mathcal{PT}_2$, respectively, to $P \setminus R$ have the same set of complete vertices, which now includes all vertices of R. Let $\boldsymbol{h}_1$ and $\boldsymbol{h}_2$ be their height vectors in the linear systems for $\mathcal{F}$ and $\mathcal{F}_2$. We know that $\boldsymbol{h}_2 \leq \boldsymbol{h}_1$ (element-wise) because a subset of R's vertices has been lowered, which implies $\boldsymbol{z}_2 \leq \boldsymbol{z}_1$ by Lemma 5.2. We conclude that $\mathcal{F}_2$ nowhere lies strictly above $\mathcal{F}$.

Observe that pseudo-triangles in $\mathcal{PT}_1$ may be incident to R only with their corners. Thus $\mathcal{F}_2$ may be reflex at edges of $\mathcal{PT}_1$ which are not in $\mathcal{PT}$. However, no reflex edge is changed to being convex. If all reflex edges of $\mathcal{F}_2$ have been present as reflex edges in $\mathcal{F}$ already, then the surface flip is complete. That is, $\varphi_R = \varphi$. Otherwise, appropriate context-free flips (convexifying or planarizing) are applied to the new reflex edges in $\mathcal{F}_2$, in the way described above and until all such edges disappear. Note that reflex edges flipped away earlier can reappear, so that the process may oscillate. It converges, though, because each reflex edge is flippable at all times, and there is an obvious lower bound on the heights, as certain vertices (the corners of P, for instance) never change height. This implies that the limit is the unique surface $\mathcal{F}'$ that $\mathcal{F}$ is flipped into by the context-sensitive surface flip φ. $\square$

Figure 6.3 illustrates two oscillating flips. The edges e_1 and e_2 are both reflex at the beginning. They alternately disappear when context-free planarizing flips are applied. Their endpoints v_1 and v_2 are lowered alternately until the limit is reached – the pseudo-triangulation that contains the shaded and the blank pseudo-triangle.

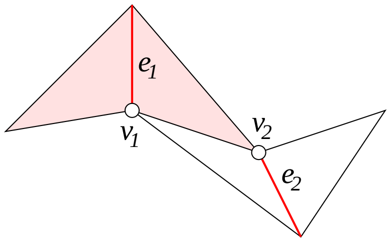

FIG. 6.3. *Oscillating flips*

6.4. Trivial flips. In a convexifying or a planarizing surface flip, a pseudo-triangulation of (P, S) may be created which is not stable. This happens if a subset $B \in S \setminus vert(P)$ of incomplete vertices is generated, such that the removal of B leaves a valid pseudo-triangle ∇ without changing the status of vertices. As a special case, a double-adjacency occurs. The obtained pseudo-triangulation is not projective, by Theorem 5.4. In fact, the surface flip changes the part of the surface above ∇ to planar. That is, a *trivial* flip takes place implicitly, which eliminates from the surface the edges and vertices that stem from B.

In the pseudo-triangulation of (P, S), the removal of B can be realized by a finite sequence of admissible flips (including vertex-removing flips, of course). This restores a projective pseudo-triangulation. Surface flips that cause trivial flips are a natural generalization of the degree-3 vertex removal in triangular surfaces, where a single tetrahedron is split off; see e.g. [38].

Remarks Each (non-tetrahedral) surface flip requires a system of linear equations to be solved; see Subsection 5.2. The corresponding matrix has a special structure, though. The first $n - i$ rows are unit rows, if i of the n vertices in the current pseudo-triangulation are incomplete. So the system can be solved in $O(n + i^3)$ time, which covers the geometrical part of the surface flip. The combinatorial part obviously can be done within this complexity.

7. Locally convex functions. Let P be a polygon in the plane. A real-valued function f with domain P is called *locally convex* if f is convex on each line segment internal to P. For convex polygons P, every locally convex function is convex continuable on the exterior of P, but this is not true for arbitrary polygons.

Let now (P, S) be an augmented polygon, and let $\boldsymbol{h}$ be a vector assigning some real value to each of the n vertices in S. We are interested in the following optimization problem:

Given (P, S) and $\boldsymbol{h}$, find the maximal locally convex function f^* on the domain P which fulfills $f^*(v_i) \leq h_i$ for each $v_i \in S$.

The function f^* is unique, because it is the pointwise maximum of all locally convex functions on P that satisfy the constraints given by $\boldsymbol{h}$. Moreover, f^* is piecewise linear, and non-linearity occurs only at line segments spanned by the set S. We prove that the graph $\mathcal{F}^*$ of f^* projects to a pseudo-triangulation of (P, S'), for some $S' \subseteq S$.

We further show that $\mathcal{F}^*$ can be constructed by a finite sequence of convexifying or planarizing surface flips.

If (P,S) is convex then a well-known special case is obtained: $\mathcal{F}^*$ is the lower convex hull of the point set S when being lifted to the heights in $\boldsymbol{h}$. In this case, we succeed in generating the (convex) surface $\mathcal{F}^*$ with $O(n^2)$ surface flips of the kind above. The bound $O(n^2)$ further can be achieved when (P,S) is an arbitrary polygon without internal vertices. Both flip distance results are asymptotically optimal, by known lower bounds; see e.g. [26].

7.1. Flipping to optimality. Consider some full triangulation $\mathcal{T}$ of the augmented polygon (P,S). All the n vertices of $\mathcal{T}$ are complete vertices, and we assign to $\mathcal{T}$ the height vector $\boldsymbol{h}$ used above to specify the maximal locally convex surface $\mathcal{F}^*$. This defines a unique triangular surface $\mathcal{F}(\mathcal{T},\boldsymbol{h})$. We state another core result of this paper.

THEOREM 7.1. *(Optimality Theorem) Let $\mathcal{F}^*(\mathcal{T},\boldsymbol{h})$ be a surface obtained from $\mathcal{F}(\mathcal{T},\boldsymbol{h})$ by applying convexifying and planarizing surface flips (in any order) as long as reflex edges do exist. Then $\mathcal{F}^*(\mathcal{T},\boldsymbol{h}) = \mathcal{F}^*$, for any choice of the initial triangulation $\mathcal{T}$. The optimum $\mathcal{F}^*$ is reached after a finite number of surface flips.*

Proof. By Lemma 6.1, convexifying and planarizing surface flips always generate a sequence of surfaces that is (strictly) totally ordered by decreasing height. As, by the uniqueness result in Theorem 5.1, these surfaces are pairwise different, so are the corresponding pseudo-triangulations. An exponential upper bound on the sequence length results, because the number of possible planar straight-line graphs on the vertex set S is exponentially bounded in n; see [7]. The obtained surface $\mathcal{F}^*(\mathcal{T},\boldsymbol{h})$ contains no reflex edge, as each such edge can be flipped at any time. This implies that $\mathcal{F}^*(\mathcal{T},\boldsymbol{h})$ is locally convex. Moreover, $\mathcal{F}^*(\mathcal{T},\boldsymbol{h})$ is maximal, because it fulfills $\boldsymbol{h}$ with equality for all complete vertices. Thus $\mathcal{F}^*(\mathcal{T},\boldsymbol{h}) = \mathcal{F}^*$ follows from the uniqueness of $\mathcal{F}^*$. □

If the vector $\boldsymbol{h}$ constraining the vertex heights is generic then all facets in the surfaces generated above are pseudo-triangles. That is, only (projective) pseudo-triangulations are generated. In particular, $\mathcal{F}^*$ projects to a pseudo-triangulation.

$\mathcal{F}^*$ can be alternatively computed by solving a linear program with at most $\binom{n}{4}$ constraints. Each constraint stems from a quadrilateral $Q \subset P$ with vertices from S, and postulates convex position for the vertex heights.

Theorem 3.4 from Section 3 implies the existence of a sequence of $O(n \log n)$ surface flips that transforms any given triangular surface $\mathcal{F}(\mathcal{T},\boldsymbol{h})$ for (P,S) into $\mathcal{F}^*$. This sequence includes convexifying and planarizing flips and, in general, their inverses. Trivial flips may be implicated, if vertices of $\mathcal{F}(\mathcal{T},\boldsymbol{h})$ are missing in $\mathcal{F}^*$. The sequence contains no inverse trivial flips, however. This outrules the re-insertion of vertices, and makes the bound substantial. Note that the sequence need not be *improving*, i.e., exclusively consisting of convexifying and planarizing (and trivial) surface flips such as in Theorem 7.1. We summarize:

THEOREM 7.2. *Let (P,S) be an augmented polygon with $|S| = n$, and let $\boldsymbol{h}$ be a height vector for S. Let f^* be the unique maximal locally convex function on P that is bounded from above by $\boldsymbol{h}$.*

(1) The graph $\mathcal{F}^*$ of f^* projects to a pseudo-triangulation in the generic case.

(2) $\mathcal{F}^*$ can be constructed from any triangular surface $\mathcal{F}$ for (P,S) and $\boldsymbol{h}$, by applying any improving flip sequence of sufficient length.

(3) There exists a sequence of $O(n \log n)$ flips (excluding the inverse trivial type) which transforms $\mathcal{F}$ into $\mathcal{F}^*$ and can be computed in polynomial time if $\mathcal{F}^*$ is known.

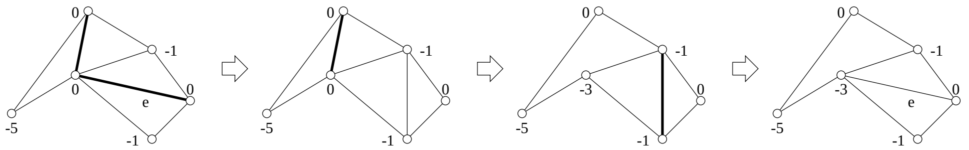

FIG. 7.1. *Edge e reappears*

A principal problem, which complicates the construction of improving sequences of polynomial length (as in [26]), is that edges flipped away may reappear. Figure 7.1 illustrates this phenomenon. Numbers denote vertex heights, and reflex edges are shown in bold. Edge e is flipped first, and a planarizing flip for the other reflex edge follows. This makes the edge between the two remaining triangles reflex. Flipping this edge next lets e reappear, though as a convex edge. But an appropriate flip in the neighborhood of e (not shown in the figure) makes e reflex again. Interestingly, a different improving sequence – that does not start with e – avoids this undesirable effect. We show in Subsections 7.2 and 7.3 that the reappearance of edges always can be avoided for less general domains (P, S).

Remarks Flipping to local convexity is not always possible *within* the class of triangulations. Different flip sequences may terminate at different triangular surfaces that contain reflex edges, as Figure 7.2 shows. Numbers at vertices denote surface heights. The initial surface contains two reflex edges, connecting vertices of height 1. Flipping either edge yields one of the two triangular surfaces shown in the upper part. Each contains only one reflex edge (shown in bold), which cannot be flipped away without destroying the triangulation. The global optimum $\mathcal{F}^*$ is a pseudo-triangular surface, shown in the lower part.

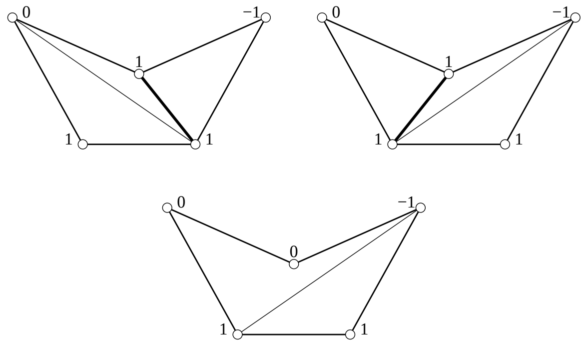

FIG. 7.2. *Local and global optima*

7.2. Convex domains. Suppose that the underlying domain (P, S) is convex. Then $\mathcal{F}^*$ bounds from below the convex hull of a three-dimensional point set. Let us show that reappearence of edges does not take place, if improving surface flips are performed in a carefully chosen order. In particular, the following holds.

LEMMA 7.3. *Let (P, S) be an augmented* convex *polygon, and let $\boldsymbol{h}$ be a height vector for S. Then $\mathcal{F}^*$ can be generated, from any given triangular surface for (P, S) and $\boldsymbol{h}$, by $O(n^2)$ improving surface flips, for $n = |S|$.*

Proof. To avoid the discussion of special cases, assume that $\boldsymbol{h}$ is generic. Let $\mathcal{F}$ be the current surface for (P, S), which is triangular at the beginning. For a vertex v

of $\mathcal{F}$, denote with $I(v)$ the set of all vertices of $\mathcal{F}$ that share a facet with v. Note that $v \in I(v)$. Let $low(I(v))$ be the lower part of the boundary of $conv(I(v))$. We use the following algorithm to flip $\mathcal{F}$ into $\mathcal{F}^*$.

Step (1) While there exists a vertex v of $\mathcal{F}$ that lies strictly above $low(I(v))$ do:
Case (1.1) Vertex v is complete. As v lies strictly above $low(I(v))$, there exist reflex edges incident to v. While v is complete, we flip away such edges, using convexifying and planarizing surface flips. We know, from Lemma 6.1, that $low(I(v))$ is lowered in this process, whereas v keeps its height. So there always exists some reflex edge incident to v which can be flipped. Each flip decrements the degree of v. Note that pseudo-triangular facets with many vertices may be created, but only if v is a corner there. Obviously, the last flip is planarizing. It either changes the status of v to incomplete, and we continue with Case (1.2), or it implicates a trivial flip that removes v (possibly together with some other vertices in v's neighborhood). In particular, v gets removed if its degree before the flip has been three.
Case (1.2) Vertex v is incomplete. Let f be the unique pseudo-triangular facet of $\mathcal{F}$ where v is a non-corner. While v has more than two edges, we flip away reflex edges, according to the following rule. As v lies strictly above $low(I(v))$, either both edges of f incident to v are reflex, or another reflex edge is incident to v. In the latter case, we flip this edge. The former case needs more care. Consider the projections, e_1 and e_2, of these two reflex edges of f in the pseudo-triangulation. Between them, there is another edge incident to the projection v' of v. Therefore, for at least one of e_1 and e_2, say e_1, the geodesic between the two corners opposite to e_1 (in the two pseudo-triangles adjacent at e_1) does not run via v'. We select e_1 for the edge to be flipped. This flip creates no edge incident to v. So either the degree of v is decremented, or v is removed in a trivial flip.

Step (2) Now all vertices v of $\mathcal{F}$ lie on $low(I(v))$. The surface $\mathcal{F}$ therefore is triangular, and all vertices in the corresponding triangulation are complete. While $\mathcal{F}$ is not convex, we apply convexifying (in fact, tetrahedral) flips. No vertex heights change. This yields the convex and triangular surface $\mathcal{F}^*$.

The total number of flips in Cases (1.1) and (1.2) is bounded by $n-1$, the maximum possible degree of the vertex v. Step (1) runs through these cases at most $n-3$ times, because each time one internal vertex is removed. The number of flips needed in Step (2) is at most $\binom{n}{2}$, by well-known arguments [26]. □

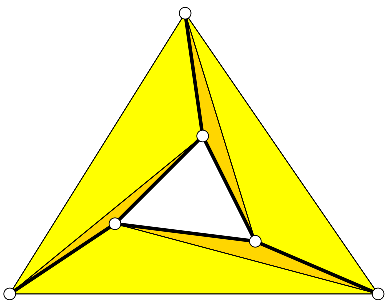

FIG. 7.3. *Local optimum for convex P*

Remarks Even though P is convex such that $\mathcal{F}^*$ is triangular, only a local optimum may be reached when tetrahedral flips are applied exclusively. See Figure 7.3 which is taken from [25]. P has height zero, and the three vertices in $S \setminus vert(P)$ lie in some plane above P. No reflex edge (bold) can be flipped, and no single vertex can be removed, to obtain a valid triangulation. Planarizing (and trivial) flips are needed to reach the optimum $\mathcal{F}^*$, which is a single triangle, P.

It remains unclear whether the algorithm above can be modified to run efficiently for *non-convex* augmented polygons (P, S). Step (1) may leave non-corners of P incomplete, such that Step (2) is not guaranteed to take $O(n^2)$ flips.

7.3. Vertex-empty domains. Improving flip sequences of (at most) quadratic length also exist for any polygonal domain (P, S) without internal vertices, that is, in the case $vert(P) = S$. Let Δ be an *ear* of the polygon P, that is, a triangle that can be cut off from P by some diagonal. It is well known that every triangulation of P contains at least two ears.

LEMMA 7.4. *Let Δ be an ear of P, and let $\boldsymbol{h}$ be a height vector for $vert(P)$. Denote with $\mathcal{F}^*$ and $\mathcal{F}^*(P \setminus \Delta)$, respectively, the optimal surfaces for P and $P \setminus \Delta$ with respect to $\boldsymbol{h}$. Then $\mathcal{F}^*(P \setminus \Delta)$ can be transformed into $\mathcal{F}^*$ by attaching the triangular facet for Δ and applying $O(n)$ improving flips.*

Proof. Consider the following sequence of flips. Imagine that a triangular facet f with the following properties is attached to $\mathcal{F}^*(P \setminus \Delta)$: The projection of f is Δ, and the corner c of f that does not belong to $\mathcal{F}^*(P \setminus \Delta)$ has some height h_c large enough to make the resulting surface convex at the respective diagonal d of P. This surface is optimum for P and the vector $\boldsymbol{h}$ with c's height replaced by h_c. Now, continuously lower h_c, and maintain the optimum surface by flips as below, until h_c attains the original value as given by $\boldsymbol{h}$, and the surface becomes $\mathcal{F}^*$.

If the diagonal d stays convex throughout then no flips are performed. Otherwise, right after the surface becomes reflex at d, we flip d. This flip is either planarizing or convexifying. It does not change the convexity of any other edge, though, and thus restores optimality. To describe the generic step, let F_c be the part of the surface modified by flips so far. F_c is separated by diagonals $d_1, \ldots, d_k$ from its complement in the surface. Therefore, when decreasing h_c, all edges of the surface except possibly $d_1, \ldots, d_k$ stay convex. As soon as one of $d_1, \ldots, d_k$ becomes reflex, the next flip is performed. This restores optimality and extends F_c by one facet.

The number of flips is $O(n)$ because each flip extends F_c. It remains to be observed that exactly the same sequence of flips can be applied if the facet f is attached to $\mathcal{F}^*(P \setminus \Delta)$ with c's original height: Each diagonal d flipped above now becomes reflex as soon as d bounds the current locally convex part incident to c. So flipping d indeed constitutes an improving flip. □

We remark that each edge constructed above corresponds to an edge of the shortest-path tree $\pi_c(P)$ for P and c, defined in Subsection 3.1.

LEMMA 7.5. *Let P be a simple polygon with n vertices, and let $\boldsymbol{h}$ be some height vector for $vert(P)$. The optimum $\mathcal{F}^*$ can be constructed by an improving flip sequence of length $O(n^2)$, from any given triangular surface for P and $\boldsymbol{h}$.*

Proof. We construct $\mathcal{F}^*$ by simulating an incremental algorithm. Let $\mathcal{T}$ be the triangulation the initial surface projects to. Traverse the triangles of $\mathcal{T}$ in some adjacency order. Let Q be the union of the triangles visited so far. We maintain the following invariant for the current surface $\mathcal{F}$. The restriction of $\mathcal{F}$ to the polygon Q is optimum for Q (i.e., maximal and locally convex), and the restriction of $\mathcal{F}$ to $P \setminus Q$ coincides with the initial triangular surface.

The generic step processed a triangle Δ of $\mathcal{T}$ by constructing the optimum surface for the polygon $Q \cup \Delta$. As Δ is an ear of $Q \cup \Delta$, this takes $O(n)$ improving flips by Lemma 7.4. These flips leave the triangular surface for $P \setminus (Q \cup \Delta)$ unaffected, because all its vertices are complete. The number of steps (triangles) is $n - 3$. □

8. (Constrained) regular complexes. The concept of local convexity of polyhedral surfaces is closely related to the concept of constrained regularity of polygonal partitions. Let us elaborate the consequences of the results in Section 7 in terms of constrained regularity.

8.1. Regular triangulations. A triangulation $\mathcal{T}$ of a point set S is called *regular* if there exists some convex function on $conv(S)$ whose set of non-linearity coincides with the set of internal edges of $\mathcal{T}$. Regular triangulations arise as duals of power diagrams [8] (a generalization of Voronoi diagrams) and have several applications [24, 23]. Figure 7.3 shows a well-known example of a non-regular triangulation.

We discuss the implications of Section 7 for regular triangulations, that is, for the case $P = conv(S)$. For each height vector $\boldsymbol{h}$, the surface $\mathcal{F}^*$ is convex in this case. It is the lower part, $low(S_{\boldsymbol{h}})$, of the boundary of the convex hull of the lifted point set $S_{\boldsymbol{h}}$. A unique regular triangulation $\mathcal{T}(S)_{\boldsymbol{h}}$ for S and $\boldsymbol{h}$ is obtained by projecting $low(S_{\boldsymbol{h}})$. Lemma 7.3 implies:

THEOREM 8.1. *Let $\mathcal{T}$ and $\boldsymbol{h}$ be a given triangulation and a height vector, respectively, for a planar set S of n points. Then $\mathcal{T}$ can be flipped to the regular triangulation $\mathcal{T}(S)_{\boldsymbol{h}}$ using $O(n^2)$ flips, which are improving in the corresponding surfaces.*

Note that the *upper* part of the boundary of the convex hull of $S_{\boldsymbol{h}}$ can be obtained as well, by reflecting $S_{\boldsymbol{h}}$ at the plane $z = 0$ first. This yields another interesting result.

COROLLARY 8.2. *The convex hull of a three-dimensional n-point set M can be constructed by taking an arbitrary triangular surface for M, and performing $O(n^2)$ improving surface flips.*

Observe that vertices of S may be absent in $\mathcal{T}(S)_{\boldsymbol{h}}$, namely if some points of $S_{\boldsymbol{h}}$ lie strictly above $low(S_{\boldsymbol{h}})$. We call $\boldsymbol{h}$ a *convex* height vector (for S) if no point of $S_{\boldsymbol{h}}$ lies strictly above $low(S_{\boldsymbol{h}})$. If $\boldsymbol{h}$ is convex then no (proper) pseudo-triangle is ever created when flipping to regularity: Pseudo-triangles stem from planarizing flips, which are never performed because the respective edges are convex already. Therefore all convexifying flips are tetrahedral flips, and all vertices show up in $\mathcal{T}(S)_{\boldsymbol{h}}$. In particular, if $S_{\boldsymbol{h}}$ happens to lie on a paraboloid with vertical axis of rotation (say, $z = \boldsymbol{x}^2$) then $\mathcal{T}(S)_{\boldsymbol{h}}$ is the Delaunay triangulation $DT(S)$ of S; see e.g. [26]

Remarks It is well known that each triangulation of S can be transformed into $DT(S)$ by flips that are convexifying and tetrahedral in the respective surface (so-called Delaunay flips), but not every regular triangulation can be constructed in this way. Theorem 8.1 tells us that pseudo-triangulations – rather than triangulations – are the right framework for flipping to regularity.

$O(n^2)$ convexifying flips always suffice for generating $DT(S)$. On the other hand, $\Omega(n^2)$ is a lower bound, even if S is in convex position; see e.g. [26]. This implies that Theorem 8.1 is optimal in the worst case. Theorem 8.1 is a stronger version of a result in [25], which asserts that incremental insertion of points, and application of tetrahedral flips to the special triangular surface that arises from each insertion, is capable of generating the regular triangulation, in $O(n^2)$ flips.

Recently, the quest for a *Delaunay pseudo-triangulation* of a point set S has been formulated [47]. If 'Delaunay' is interpreted as 'convex at each edge of the projection surface', and if the *entire* class of pseudo-triangulations for S is taken as the underlying setting (rather than some class of fixed edge rank) then our results imply that $\mathcal{T}(S)_{\boldsymbol{h}}$ (or $DT(S)$, respectively) is the solution. The Delaunay *minimum* pseudo-triangulation does not exist in this setting, unless S is in convex position: No minimum pseudo-triangulation of S is regular if S contains some internal vertex v, because the incompleteness of v forces every surface to be non-convex in the neighborhood of v.

8.2. Constrained regular pseudo-triangulations. Let us return to general augmented polygons (P, S). Here the maximal locally convex surface $\mathcal{F}^*$ is not convex continuable on the exterior of P, in general. Thus $\mathcal{F}^*$ disagrees with the lower convex hull of the lifted point set $S_{\boldsymbol{h}}$ and contains pseudo-triangular facets. The pseudo-triangulation that $\mathcal{F}^*$ projects to always contains the edges of P. It can be interpreted as the regular pseudo-triangulation for S *constrained* by P. Constrained triangulations, and especially, constrained *Delaunay* triangulations are well-known and versatile concepts; see e.g. [40, 19, 4]. Very recently, constrained *regular* triangulations have been considered [13, 50]. This structure is not guaranteed to exist for general height vectors $\boldsymbol{h}$, however. We intend to show that our framework allows for a definition of constrained regular pseudo-triangulations, that always exist and are unique.

To exploit our concepts in more generality, let G be a planar straight-line graph with vertex set S. Local convexity of a function f generalizes naturally, by requiring f to be convex on each line segment in $conv(S)$ that does not cross G. We call G *admissible* if, apart from the isolated vertices of G, each component of G is connected to the boundary of $conv(S)$. An admissible graph G defines a partition of $conv(S)$ as follows; see Figure 8.1. Cut $conv(S)$ along G, and move it apart at the cuts infinitesimally. This gives a collection of augmented polygons (Q_i, S_i) with $\bigcup_i S_i = S$.

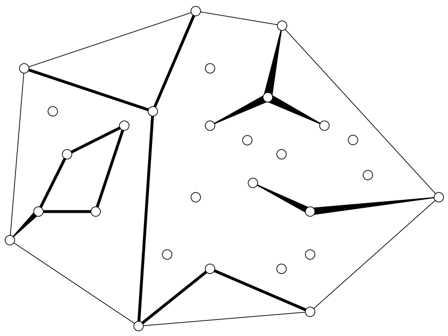

FIG. 8.1. *Cutting $conv(S)$ along G*

Theorems 7.1 and 7.2 directly apply to this setting. Let $\mathcal{T}$ be a triangulation of S that includes all the edges of G. For a given height vector $\boldsymbol{h}$ for S, consider the unique triangular surface for $\mathcal{T}$ and $\boldsymbol{h}$. Let $\mathcal{F}_i$ denote the restriction of this surface to the polygon Q_i. Then $\mathcal{F}_i$ can be made locally convex, by flipping away all reflex edges that do not belong to G. This gives a unique pseudo-triangulation $\mathcal{PT}_i$ for each augmented polygon (Q_i, S_i). Altogether, a unique constrained regular pseudo-triangulation $\mathcal{PT}(G)_{\boldsymbol{h}}$ for G and $\boldsymbol{h}$ is obtained.

The corresponding maximal locally convex function $f_{G,\boldsymbol{h}}$ on the domain $conv(S)$ is piecewise linear but, in general, not continuous. Its set of discontinuity is a subset of G, and occurs at certain pseudo-triangular faces of $\mathcal{PT}(G)_{\boldsymbol{h}}$, namely where local convexity (with respect to G) and maximality (with respect to $\boldsymbol{h}$) contradict. The graph of $f_{G,\boldsymbol{h}}$ thus is no surface in our sense. We summarize:

THEOREM 8.3. *Let S be a set of n points in the plane, and let G and $\boldsymbol{h}$, respectively, be an admissible straight-line graph and a height vector for S. There exists a unique constrained regular pseudo-triangulation $\mathcal{PT}(G)_{\boldsymbol{h}}$ for G and $\boldsymbol{h}$. $\mathcal{PT}(G)_{\boldsymbol{h}}$ can be constructed by improving surface flips, starting from any triangulation of S that conforms with G. Moreover, the $O(n \log n)$ bound in Theorem 7.2 applies.*

Remarks If G partitions *conv*(S) into augmented *convex* polygons, then no incomplete vertices occur in $\mathcal{PT}(G)_{\boldsymbol{h}}$, which thus is a triangulation for every $\boldsymbol{h}$. The graph $F_{G,\boldsymbol{h}}$ of $f_{G,\boldsymbol{h}}$ then is a surface, namely the continuous concatenation of lower convex hulls on top of these convex polygons.

There are other interesting instances of admissible constraining graphs G. For example, G may define a minimum pseudo-triangulation $\mathcal{M}$ of S. In this case, $\mathcal{PT}(G)_{\boldsymbol{h}}$ gives ways of canonically refining $\mathcal{M}$. Or, G may define a spanning tree of S, a choice used in [4] to transform crossing-free spanning trees into each other.

If the height vector $\boldsymbol{h}$ is convex then all vertices of S show up in $\mathcal{PT}(G)_{\boldsymbol{h}}$. Moreover, only triangles are generated in the flipping process. $\mathcal{PT}(G)_{\boldsymbol{h}}$ is a triangulation and $f_{G,\boldsymbol{h}}$ is continuous. The surface $F_{G,\boldsymbol{h}}$ is reflex only at edges of G, namely at those being absent in the regular triangulation $\mathcal{T}(S)_{\boldsymbol{h}}$. In the special case where the lifted point set $S_{\boldsymbol{h}}$ lies on the paraboloid $z = \boldsymbol{x}^2$, the constrained Delaunay triangulation [40] of S and G is obtained.

If the graph G is not admissible, then G splits *conv*(S) into augmented polygons with possible holes. $\mathcal{PT}(G)_{\boldsymbol{h}}$ still is uniquely defined, because Theorem 5.1 generalizes to augmented polygons with holes in a straightforward way. However, the constrained regular complex $\mathcal{PT}(G)_{\boldsymbol{h}}$ fails to be a pseudo-triangulation, in general, because pseudo-triangular facets with holes may arise. It is worth mentioning that the polytope constructed in Section 9 generalizes to representing arbitrary constrained regular complexes.

An attempt to define constrained regular complexes as full triangulations results in one of the two phenomena: If convexity on all edges different from G is required, then the complex does not always exist. On the other hand, ambiguity occurs if reflex edges not belonging to G are admitted, even if they are all non-flippable (in the classical sense). Their structure depends on the flip sequence rather than on G and $\boldsymbol{h}$ alone. That is, several solutions (local optima) do exist in general. See Figure 7.2, where G delineates a non-convex polygon.

9. Polytope representation. Let S be a set of n points in the plane. Let G denote an admissible straight-line graph on S, as defined in Subsection 8.2. We establish the existence of a convex polytope that represents all the regular pseudo-triangulations on S and constrained by G. This generalizes the polytope constructions in [39] (the associahedron) and in [28] (the secondary polytope, see also [15]), which concern the (unconstrained) regular triangulations of a point set S.

Let S' be a subset of S, which includes all the vertices of the (unique) connected component formed by G and the boundary of *conv*(S). We consider the set $\Pi(G)$ of all pseudo-triangulations $\mathcal{PT}$ that live on such subsets S' and conform with G. Clearly, for each height vector $\boldsymbol{h}$ for S, the constrained regular pseudo-triangulation $\mathcal{PT}(G)_{\boldsymbol{h}}$ is a member of $\Pi(G)$.

We construct, for each $\mathcal{PT} \in \Pi(G)$ and each $\boldsymbol{h}$, a unique piecewise linear function $f_{\mathcal{PT},\boldsymbol{h}}$ on the domain *conv*(S). Recall from Subsection 8.2 that the admissible graph G defines a partition of *conv*(S) into polygons Q_i. Let $\mathcal{PT}_i$ denote the part of $\mathcal{PT}$ inside Q_i. Then $\boldsymbol{h}$ assigns a height to each complete vertex of the pseudo-triangulation $\mathcal{PT}_i$. So, by the surface theorem (Theorem 5.1), there exists a unique polyhedral surface $\mathcal{F}_i$ for $\mathcal{PT}_i$ and $\boldsymbol{h}$. The concatenation of the surfaces $\mathcal{F}_i$ is the graph of the function $f_{\mathcal{PT},\boldsymbol{h}}$ to be constructed. Note that $f_{\mathcal{PT},\boldsymbol{h}}$ is discontinuous, in general.

Recall further that $f_{\mathcal{PT},\boldsymbol{h}}$ is locally convex with respect to G exactly if $\mathcal{PT} = \mathcal{PT}(G)_{\boldsymbol{h}}$. The optimality theorem (Theorem 7.1) thus implies $f_{\mathcal{PT},\boldsymbol{h}} \leq f_{\mathcal{PT}',\boldsymbol{h}}$ for all

$\mathcal{PT}' \in \Pi(G)$. Integrating both sides yields

$$\int_{x \in conv(S)} f_{\mathcal{PT},\boldsymbol{h}}(x)\ dx \leq \int_{x \in conv(S)} f_{\mathcal{PT}',\boldsymbol{h}}(x)\ dx. \tag{9.1}$$

Consider the lefthand-side integral in (9.1). The expressed volume is a linear (and homogenous) function of $\boldsymbol{h}$, because $f_{\mathcal{PT},\boldsymbol{h}}$ is piecewise linear and, by Lemma 5.2, the value of $f_{\mathcal{PT},\boldsymbol{h}}$ at each vertex of $\mathcal{PT}$ linearly depends on $\boldsymbol{h}$. Let $\boldsymbol{q}(\mathcal{PT})$ be the coefficient vector of this linear function. As analogous observations hold for the righthand side in (9.1), we obtain

$$\boldsymbol{q}(\mathcal{PT}) \cdot \boldsymbol{h} \leq \boldsymbol{q}(\mathcal{PT}') \cdot \boldsymbol{h}, \text{ for all } \mathcal{PT}' \in \Pi(G). \tag{9.2}$$

We now interpret $\boldsymbol{q}(\mathcal{PT})$ as a point in R^n, and consider the convex polytope

$$\mathcal{Q}(G) = conv\{\boldsymbol{q}(\mathcal{PT}) \mid \mathcal{PT} \in \Pi(G)\}.$$

Further, we define the following set of height vectors:

$$\mathsf{V}(G, \mathcal{PT}) = \left\{\boldsymbol{h} \mid \mathcal{PT} = \mathcal{PT}(G)_{\boldsymbol{h}}\right\}.$$

Inequality (9.2) implies that, for each $\boldsymbol{h} \in \mathsf{V}(G, \mathcal{PT})$, the set $\{\boldsymbol{q}(\mathcal{PT}) \mid \mathcal{PT} \in \Pi(G)\}$ lies in a halfspace whose boundary contains $\boldsymbol{q}(\mathcal{PT})$. So $\boldsymbol{q}(\mathcal{PT})$ is a vertex of the polytope $\mathcal{Q}(G)$ if and only if $\mathcal{PT} = \mathcal{PT}(G)_{\boldsymbol{h}}$, for some height vector $\boldsymbol{h}$. Note that $\boldsymbol{q}(\mathcal{PT})$ lies in the interior of $\mathcal{Q}(G)$ if $\mathcal{PT}$ is not constrained regular.

The collection of all sets $\mathsf{V}(G, \mathcal{PT})$, for $\mathcal{PT} \in \Pi(G)$, defines a cell complex $\mathcal{C}(G)$ in R^n, because $\mathcal{PT}(G)_{\boldsymbol{h}}$ exists and is unique for each $\boldsymbol{h}$, by Theorem 8.3. $\mathsf{V}(G, \mathcal{PT})$ contains, for each member $\boldsymbol{h}$, the vector $\lambda \cdot \boldsymbol{h}$ for $\lambda \geq 0$. Moreover, $\boldsymbol{h}_1, \boldsymbol{h}_2 \in \mathsf{V}(G, \mathcal{PT})$ implies $\lambda \cdot \boldsymbol{h}_1 + (1-\lambda) \cdot \boldsymbol{h}_2 \in \mathsf{V}(G, \mathcal{PT})$, for $0 < \lambda < 1$. That is, $\mathsf{V}(G, \mathcal{PT})$ is a convex polyhedral cone with apex 0.

Consider a facet $g = \mathsf{V}(G, \mathcal{PT}_1) \cap \mathsf{V}(G, \mathcal{PT}_2)$ of the complex $\mathcal{C}(G)$. For $\boldsymbol{h} \in g$ we have $\boldsymbol{h} \cdot (q(\mathcal{PT}_1) - q(\mathcal{PT}_2)) = 0$, by inequality (9.2). So the line segment connecting $q(\mathcal{PT}_1)$ and $q(\mathcal{PT}_2)$ is normal to g. This implies that the edges of $\mathcal{Q}(G)$ are dual to the facets of $\mathcal{C}(G)$. For $\boldsymbol{h}$ lying in the relative interior of such a facet, there are exactly two constrained regular pseudo-triangulations $\mathcal{PT}_1$ and $\mathcal{PT}_2$ that yield the same function $f_{\mathcal{PT}_1,\boldsymbol{h}} = f_{\mathcal{PT}_2,\boldsymbol{h}}$. Therefore, the graphs of $f_{\mathcal{PT}_1,\boldsymbol{h}_1}$ and $f_{\mathcal{PT}_2,\boldsymbol{h}_2}$, for $\boldsymbol{h}_1 \in \mathsf{V}(G, \mathcal{PT}_1)$ and $\boldsymbol{h}_2 \in \mathsf{V}(G, \mathcal{PT}_2)$, differ by a minimum combinatorial change – which corresponds to a flip that transforms $\mathcal{PT}_1$ and $\mathcal{PT}_2$ into each other. We conclude:

THEOREM 9.1. *Let S be a finite set of points in the plane, and let G be an admissible straight-line graph spanned by S. There exists a convex polytope for G, whose vertices are in one-to-one correspondence with the regular pseudo-triangulations of S constrained by G. Moreover, the edges of this polytope correspond to flips as in Section 2.*

Remarks The existence of $\mathcal{Q}(G)$ implies that any two constrained regular pseudo-triangulations $\mathcal{PT}(G)_{\boldsymbol{h}}, \mathcal{PT}(G)_{\boldsymbol{h}'} \in \Pi(G)$ are connected by a sequence of flips that retain constrained regularity. That is, there exists a sequence $\boldsymbol{h} = \boldsymbol{h}_1, \boldsymbol{h}_2, \ldots, \boldsymbol{h}_k = \boldsymbol{h}'$ of height vectors, such that the functions $f_{G,\boldsymbol{h}_i}$ are all locally convex with respect to G. If G is chosen such that $f_{G,\boldsymbol{h}_i}$ is continuous (which, for instance, is the case if G defines a convex partition of $conv(S)$), then a tool for morphing locally convex surfaces is obtained. We raise the question for the maximal length k of the sequence above, that is, for the diameter of $\mathcal{Q}(G)$. The diameter obviously depends on both $n = |S|$ and G.

It is not hard to see that every pseudo-triangulation of a simple polygon P (without internal vertices) is constrained regular. The corresponding polytope $\mathcal{Q}(P)$ therefore represents *all* possible pseudo-triangulations of P. Moreover, any two pseudo-triangulations of P can be tranformed into each other with $O(n)$ flips; see Subsection 3.1. We obtain:

COROLLARY 9.2. *Let P be a simple polygon with n vertices. There exists a polytope $\mathcal{Q}(P)$ whose vertices and edges, respectively, represent all possible pseudo-triangulations of P, and the flips between them. The diameter of $\mathcal{Q}(P)$ is $O(n)$.*

Corollary 9.2 shows that, when generalizing from convex to non-convex polygons, the diameter of the associahedron [39] remains linear.

Does there exist a polytope representation of all possible pseudo-triangulations of an augmented polygon? It does for the subclass of all minimum pseudo-triangulations of a point set S, by results in [47]. In view of our extended repertoire of flips, this question arises naturally. We recently learned that an affirmative answer for the point set case has been given in [41], by generalizing the construction in [47]. Our results on flip distances in Section 3 imply that the diameters of the polytopes in [47] and [41] are $O(n \log^2 n)$ and $O(n \log n)$, respectively. What is the relation of these polytopes to the polytopes in Theorem 9.1 and Corollary 9.2?

10. Conclusion. We have demonstrated that pseudo-triangulations enjoy rich geometric and combinatorial properties, and thus have contributed to the theory of pseudo-triangulations. Several key concepts for triangulations have been generalized to pseudo-triangulations. A series of algorithmic implications have been described, which we believe to be of value in practical applications. The three main ingredients have been (1) the derivation a new type of flip, (2) a proof of the realizability of pseudo-triangulations as polyhedral surfaces, and (3) the establishment of a link to locally convex functions. For a better understanding of pseudo-triangulations, several notions have been revitalized or introduced the first time, like the edge rank, the completeness status of vertices, and the stability of pseudo-triangulations. Let us conclude this paper with an extension of our results which seems most interesting to us.

The relationship between pseudo-triangulations and locally convex functions, expressed in Theorem 7.2, opens a way to define pseudo-complexes in higher dimensions, and flips therein. Also, the techniques used to prove the surface theorem (Theorem 5.1) do not depend on the underlying dimension. These tools may finally resolve the '3D pseudo-triangulation mystery'. We are exploring the consequences of this approach in a separate paper.

Acknowledgements The authors gratefully acknowledge insightful discussions on the presented topic with 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] H.Ahn, S.-W.Cheng, O.Cheong, J.Snoeyink. The reflex-free hull. 13th Canadian Conf. Computational Geometry 2001, 9-12.

[3] O.Aichholzer, F.Aurenhammer. Straight skeletons for general polygonal figures in the plane. In: Voronoi's Impact on Modern Sciences II, A.M.Samoilenko (ed.), Proc. Institute of Mathematics of the National Academy of Sciences of Ukraine 21, 1998, 7-21.

[4] O.Aichholzer, F.Aurenhammer, F.Hurtado. Sequences of spanning trees and a fixed tree theorem. Computational Geometry: Theory and Applications 21 (2002), 3-20.

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

[6] O.Aichholzer, M.Hoffmann, B.Speckmann, C.D.Tóth. Degree bounds for constrained pseudo-triangulations. Manuscript, Inst. Theor. Comput. Sci., Graz University of Technology, Austria, 2003.

[7] M.Ajtai, V.Chvátal, M.M.Newborn, E.Szemerédi. Crossing-free subgraphs. Ann. Discrete Math. 12 (1982), 9-12.

[8] F.Aurenhammer. Power diagrams: properties, algorithms, and applications. SIAM J. Computing 16 (1987), 78-95.

[9] F.Aurenhammer. A criterion for the affine equivalence of cell complexes in R^d and convex polyhedra in R^{d+1}. Discrete & Computational Geometry, 2 (1987), 49-64.

[10] F.Aurenhammer. A relationship between Gale transforms and Voronoi diagrams. Discrete Applied Mathematics, 28 (1990), 83-91.

[11] 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.

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

[13] S.Balaven, C.Bennis, J.-D.Boissonnat, M.Yvinec. Conforming orthogonal meshes. 11th International Meshing Roundtable, Ithaca, New York, USA, 2002.

[14] 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.

[15] L.J.Billera, P.Filliman, B.Sturmfels. Constructions and complexity of secondary polytopes. Advances in Mathematics 83 (1990), 155-179.

[16] H.Brönnimann, L.Kettner, M.Pocchiola, J.Snoeyink. Counting and enumerating pseudo-triangulations with the greedy flip algorithm. Manuscript, Polytechnic University, Brooklyn, U.S.A, 2001.

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

[18] 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.

[19] L.P.Chew. Constrained Delaunay triangulations. Algorithmica 4 (1989), 97-108.

[20] H.Crapo, W.Whiteley. Spaces of stresses, projections, and parallel drawings of spherical polyhedra. Beiträge Algebra Geom. 35 (1994), 259-281.

[21] L.De Floriani, B.Falcidieno, C.Pienovi, G.Nagy. On sorting triangles in a Delaunay tesselation. Rep. Inst. Mat. Appl., Consiglio Nazionale della Richerche, Genova, Italy, 1988.

[22] H.Edelsbrunner. An acyclicity theorem for cell complexes in d dimensions. Combinatorica 10 (1990), 251-260.

[23] H.Edelsbrunner. The union of balls and its dual shape. Discrete & Computational Geometry 13 (1995), 415-440.

[24] H.Edelsbrunner, E.P.Mücke. Three-dimensional alpha shapes. ACM Transactions Graph. 13 (1994), 43-72.

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

[26] 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.

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

[28] I.M.Gel'fand, M.M.Kapranov, A.V.Zelevinsky. Newton polyhedra of principal A-determinants. Soviet Math. Dokl. 308 (1989), 20-23.

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

[30] 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. 19th Ann. ACM Sympos. Computational Geometry, to appear.

[31] S.Hanke, T.Ottmann, S.Schuierer. The edge-flipping distance of triangulations. Journal of Uni-

versal Computer Science 2 (1996), 570-579.

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

[33] C.Huemer. Master Thesis, Inst. Theor. Comput. Sci., Graz University of Technology, Austria, 2003.

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

[35] 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.

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

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

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

[39] C.W.Lee. The associahedron and triangulations of the n-gon. European J. Combinatorics 10 (1989), 173-181.

[40] D.T.Lee, A.K.Lin. Generalized Delaunay triangulation for planar graphs. Discrete & Computational Geometry 1 (1986), 201-217.

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

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

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

[44] M.Pocchiola, G.Vegter. On polygon covers. Contemporary Mathematics 223 (1999), 257-258.

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

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

[47] 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.

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

[49] B.Speckmann, C.D.Toth. Allocating vertex π-guards in simple polygons via pseudo-triangulations. Proc. 14th ACM-SIAM Symp. Discrete Algorithms, 2003, 109-118.

[50] J.R.Shewchuk, Updating and constructing constrained Delaunay and constrained regular triangulations by flips. Proc. 19th Ann. ACM Sympos. Computational Geometry, to appear.

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