Pre-Triangulations and Liftable Complexes

Oswin Aichholzer
IST, Graz University of Technology
oaich@ist.tugraz.at

Franz Aurenhammer
IGI, Graz University of Technology
auren@igi.tugraz.at

Thomas Hackl [1]
IST, Graz University of Technology
thackl@ist.tugraz.at

1. INTRODUCTION

Polygonal complexes in the plane have been objects of interest in combinatorial geometry from various points of view. With the advent of computational geometry, it soon became apparent that combinatorial and geometric properties of certain polygonal complexes prove useful for structuring geometric data and designing efficient algorithms. Classical examples are line arrangements that arise as duals of finite point sets [8], Voronoi diagrams that capture proximity information among geometric objects [2, 10], and triangulations that connect or partition geometric objects in predefined domains [5, 10].

Whereas generalizations of line arrangements and Voronoi diagrams meanwhile have been studied extensively, the discovery of a structure that generalizes triangulations but still retains their basic properties (e.g., planarity, simple face shape, and flippability) happened more recently. In a so-called pseudo-triangulation, faces bounded by three concave chains, rather than by three line segments, are allowed. Pseudo-triangulations enjoy a variety of combinatorial and geometric properties, and lead to efficient data structures and algorithms in several areas. See e.g. [18, 22, 1, 17], and references therein, respectively.

The aim of this paper is to generalize triangulations further – in a natural way and to the utmost in a certain sense. In a pseudo-triangulation, a face (pseudo-triangle) is a simply connected polygonal region where exactly three vertices have no reflex angle. Dropping simplicity, we arrive at a concept we will call a *pre-triangle*, and following suit, a *pre-triangulation* of a given domain. We show that pre-triangulations arise in three different contexts: In the characterization of complexes that are liftable to three-space in a strong sense, in flip sequences for general polygonal complexes, and as graphs of maximal locally convex functions. Below we give some background on these topics and briefly outline our results.

Supported by the FWF Joint Research Project 'Industrial Geometry' S9205-N12.

[1] This work was done while Thomas Hackl was with the Institute for Theoretical Computer Science, University of Technology, Graz, Austria. Parts of the results are presented in the master thesis [12].

.

1.1 Liftable complexes

The issue of characterizing complexes that can be 'lifted' to space has been frequently investigated in the mathematical literature. A classical theorem of Steinitz [21] implies that, for every polygonal complex in the plane, there exists a convex 3-polyhedron with isomorphic boundary. The Maxwell-Cremona theorem, see e.g. [7], characterizes polygonal complexes whose edges are (exactly) the vertical projections of the edges of a polyhedral surface. Complexes with this property are sometimes called *projective* complexes. When the projection surface is required to be convex, the well-studied class of *regular* complexes is obtained. Several criteria for characterizing regular complexes (in general dimensions) exist; see e.g. [3] for a short bibliography.

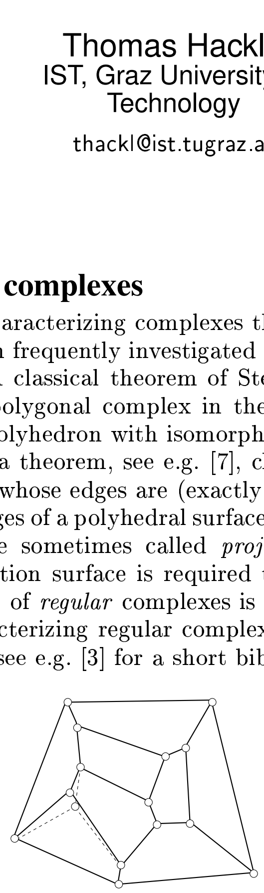

Figure 1: A Schlegel diagram

We introduce the notion of *combinatorial projectivity*, which informally means projectivity with probability one after a random ϵ-perturbation of the vertices of the complex in question. Triangulations are projective complexes that are also combinatorial projective (though they are not necessarily regular). Schlegel diagrams [11], and thus Voronoi diagrams and power diagrams [3], are regular and therefore projective. However, these complexes are not combinatorial projective. It is known that, apart from special cases, any movement of their vertices destroys their projectivity; see Figure 1.

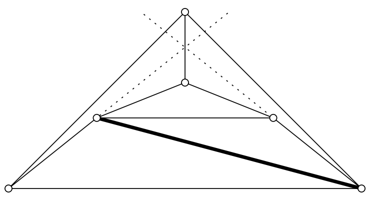

Figure 2: A pseudo-triangulation

On the other hand, the complex shown in Figure 2 (taken from [1]) is not projective, because the edge drawn in bold

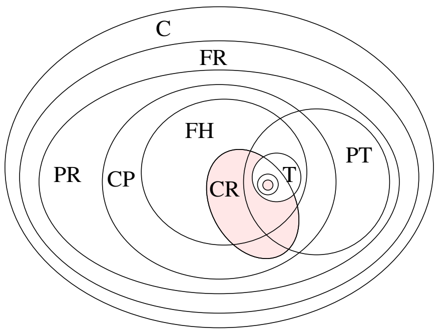

C	... polygonal complexes
FR	... face-reducible complexes
PR	... pre-triangulations
CP	... combinatorial projective complexes
FH	... face-honest complexes
CR	... combinatorial regular complexes
PT	... pseudo-triangulations
T	... triangulations
◯	... Delaunay triangulations constrained by R
(shaded)	... location for Delaunay minimum complex of R

Figure 3: Hierarchy of polygonal complexes in a given region R

can be shown to flatten out in all possible projection surfaces. But ϵ-perturbing almost surely restores projectivity in this case, that is, this complex is combinatorial projective.

In Sections 3 through 6 we derive a criterion that completely characterizes when a given polygonal complex is combinatorial projective. This characterization hinges on the concept of pre-triangulations, introduced in Section 3. More specifically, a complex is combinatorial projective if and only if it coincides with its *M-skeleton*, a certain pre-triangulation defined in Section 4. In that section we also show that the surface theorem for pseudo-triangulations in [1] holds in a more general setting. Loosely speaking, this theorem asserts that three (non-trivial) vertex heights per face can be chosen in any given pseudo-triangulation such that each face lifts to planarity. We characterize the class of complexes where the surface theorem applies. This leads to the class of so-called *face-reducible* complexes which contains the pre-triangulations as a proper subclass.

In Section 7 we turn to convex projection surfaces and define *combinatorial regularity*, a property stronger than combinatorial projectivity. We introduce the class of *face-honest* complexes, which are basically those where each single face can be lifted to a different plane. Both face-honest complexes and combinatorial regular complexes are certain pre-triangulations. Combinatorial regular complexes in simply connected regions are shown to be face-honest pseudo-triangulations. The various considered complex classes and their containment relations are illustrated in Figure 3. The picture is nonredundant in the sense that subclasses are proper and class overlaps are nonempty.

1.2 Flips and convexity

One of the most basic properties of convex sets is their facial structure [6]. In this sense, every convex function generates a cell complex in its domain of definition. Requiring maximality with respect to finitely many predefined function values leads to a complex whose faces are polygonal. This insight dates back to the classical observation that Delaunay triangulations are projected lower convex hulls; see e.g. [10]. A similar relation exists in a more general setting [1]: Maximal *locally convex* functions generate constrained regular pseudo-triangulations if the domain of definition is a nonconvex polygon. We extend the situation to full generality. In Section 8 we prove that graphs of locally convex functions on arbitrary polygonal domains generate combinatorial regular complexes, and thus pre-triangulations, in the generic case. This result turns out useful in the design of a new and powerful flipping operation for polygonal complexes, whose properties are studied in Sections 9 to 12.

So-called *flips* are a frequently used means to modify polygonal complexes. Various types of flip operations have been considered in the literature, mainly for triangulations and pseudo-triangulations. To give a few relevant citations, we refer to [15, 9, 14] and [18, 22, 1, 17], respectively. All these flips are defined via geodesic lines in the domain where the flip takes place.

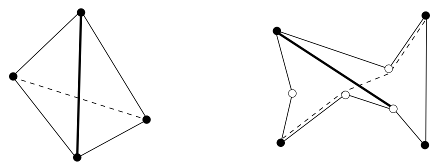

Figure 4: Exchanging flips

Figure 4 shows the standard Lawson flip for triangulations (left) and the exchanging flip for pseudo-triangulations (right). An edge-removing flip and a vertex-removing flip are illustrated in Figure 5 left and right, respectively. In both figures, the edges to be flipped are drawn in bold, and the geodesic lines that define the edges created in the flip (if any) are shown dashed. The class of pseudo-triangulations is closed under flips of these four types.

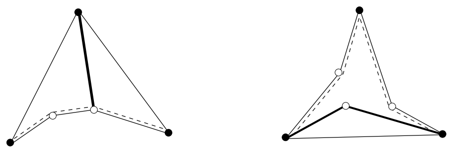

Figure 5: Removing flips

Using a radically different approach that is based on the machinery of locally convex functions, we derive a general flip operation in Sections 9 and 10. This operation works for the entire class of face-reducible complexes and covers all the classical flip types. The smallest class that is closed under this operation are the pre-triangulations.

In Section 11 we introduce the *Delaunay minimum complex*, which is a variant of the well-known Delaunay triangulation [2, 10] for general polygonal domains. This unique structure is the complex of smallest combinatorial size that still retains the desired 'Delaunay properties': an analog of local Delaunayhood, and reachability by improving flips. The Delaunay minimum complex is strongly related to the concept of constrained regular pseudo-triangulations introduced in [1], and generalizes the concept of pointed Delaunay pseudo-triangulation in [20]. We prove that any given face-reducible complex in a polygonal region R can be flipped to the Delaunay minimum complex of R by means of improving flips. This connectivity result is a generalization of the optimality theorem in [1]. Section 12 shows that, within simply connected regions, every triangulation can be flipped to a predefined constrained regular pseudo-triangulation, in a way such that all intermediate complexes are face-honest pseudo-triangulations (except immediately before a vertex is removed). Face-honest flip sequences are desirable because they change the complex at hands in a local way.

2. POLYGONAL COMPLEXES

This section provides the definitions and notions we will use to work on polygonal complexes.

Let R be a bounded subset of the plane. We call R a *polygonal region* if the boundary of R is piecewise linear and coincides with the boundary of the interior of R. Neither connectedness nor simple connectedness of a polygonal region R is required. The boundary components of R are called *edges* and *vertices* of R. A *corner* of R is a vertex of R with no internal angle larger than π. All other vertices of R are termed *noncorners* of R. See Figure 6. The depicted region consists of three connected components, two being simply connected. Corners and noncorners are distinguished as black dots and white dots. Note that more than one internal angle may arise at a single vertex. For corners, all these angles have to be convex.

A *polygon* is a polygonal region that is homeomorphic to a disk. Observe that the convex hull of a polygonal region R is a convex polygon whose vertices are corners of R. This implies that every polygonal region has at least 3 corners.

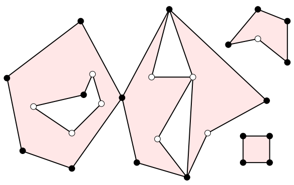

Figure 6: General polygonal region

A *polygonal partition*, $\mathcal{C}$, of R is a partition of R into (finitely many) polygonal regions. Such a region f is called a *face* of $\mathcal{C}$ if the interior of f is connected. Faces need not be simply connected; they may contain holes and thus are no polygons, in general. The *edges* and *vertices* of $\mathcal{C}$ are the edges and vertices of its faces. An edge (vertex) of $\mathcal{C}$ is called *internal* (to R) if it does not lie on the boundary of R. $\mathcal{C}$ is termed a *polygonal complex* in R if each internal edge is an edge of two different faces.

We will restrict attention to polygonal complexes in this paper. The vertices of any polygonal complex $\mathcal{C}$ are assumed to be in general position[1] in the plane. As the vertices of the underlying polygonal region R arise as vertices of $\mathcal{C}$, the vertices of R are required to be in general position as well.

Consider an arbitrary subset B of faces of a polygonal complex $\mathcal{C}$. Let v be a vertex of some face in B. Adopting notation from [1], vertex v is called *complete in* B if v is a corner of each face in B that is incident to v. Otherwise, v is called *incomplete in* B. For vertices of $\mathcal{C}$ which are not incident to some face in B, the completeness status with respect to B is left undefined.

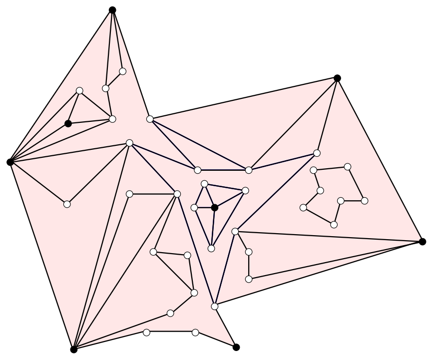

Figure 7: A polygonal complex

For convenience, we will mostly view a polygonal complex $\mathcal{C}$ as the set of its faces. If a vertex v is complete in $\mathcal{C}$, then there is no subset B of faces of $\mathcal{C}$ where v is incomplete. Conversely, if v is incomplete in $B \subset \mathcal{C}$, then v is also incomplete in $\mathcal{C}$. We define the *degree* of B as the number of vertices of $\mathcal{C}$ that are complete in B. The notion of degree of a face set is crucial for the developments in the present paper. Notice that the union of the faces in B is a polygonal region, and that each corner of this region has to be complete in B. This implies that the degree of B, and in particular the degree of $\mathcal{C}$, is at least 3.

Figure 7 shows a polygonal complex of degree 8. Complete and incomplete vertices are drawn as black dots and white dots, respectively. We will keep this convention (which is also compatible with Figure 6) throughout this paper. Observe that the (in)completeness of a vertex is a notion based on the underlying complex, whereas the (non)corner property of a vertex is a notion based on the underlying polygonal region. Note finally that adjacent faces of a complex may touch at many edges, but each internal edge has to belong to exactly two faces.

3. PRE-TRIANGULATIONS

This section introduces the concept of pre-triangulation. We start by recalling the (related) definition of a pseudo-triangulation [18, 22, 1, 17]. We then define minimum complexes in polygonal regions, and finally give a characterization of face sets of degree 3. Such face sets will turn out important in most of our subsequent investigations.

A *pseudo-triangle* is a polygon with exactly 3 corners. A *pseudo-triangulation* is a polygonal complex all whose faces are pseudo-triangles.

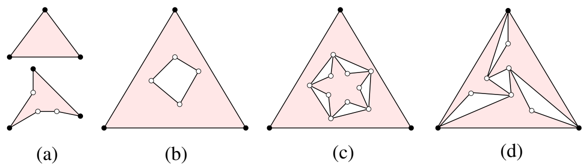

Figure 8: Five valid pre-triangles

[1]A set S of points in the plane is in *general position* if no 3 points of S are collinear and no 4 points of S are cocircular.

Define a *pre-triangle* as an arbitrary polygonal region with exactly 3 corners. Clearly, any pseudo-triangle is a pre-triangle, but the latter may contain convex holes, because no vertex of such a hole is an additional corner. In fact, the interior of a pre-triangle may be disconnected, as it may consist of many edge-disjoint faces. See Figure 8. Part (a) illustrates two pseudo-triangles, whereas parts (b), (c), and (d) show pre-triangles which are not pseudo-triangles. The pre-triangle in (c) consists of two faces, and the pre-triangle in (d) consists of three faces.

A polygonal complex $\mathcal{C}$ is called a *pre-triangulation* if $\mathcal{C}$ can be partitioned into subsets $B_1, \ldots, B_t$ of faces such that, for each subset B_i, the union of the faces in B_i is a pre-triangle, and the faces of B_i are pairwise edge-disjoint. Clearly, every pseudo-triangulation is also a pre-triangulation. However, a pre-triangulation may contain faces with holes (Figure 9) and even faces with a large number of corners (Figure 8c).

Let $\mathcal{C}$ be some polygonal complex. Recall that each corner of the underlying polygonal region R is complete in $\mathcal{C}$. This leads us to term $\mathcal{C}$ a *minimum* polygonal complex if the corners of R are the only vertices that are complete in $\mathcal{C}$. In other words, the degree of $\mathcal{C}$ equals the number of corners of R, the minimum that can be achieved. To give examples, the complex in Figure 13 is minimum whereas the complexes in Figures 7 and 9 are not.

For pseudo-triangulations, the definition above is consistent with the original definition that uses vertex pointedness: In a minimum pseudo-triangulation [22] each vertex is pointed, that is, its incident edges span a convex angle. Observe that a polygonal complex $\mathcal{C}$ in R is minimum if and only if $\mathcal{C}$ can be 'filled up' to a minimum pseudo-triangulation in R by adding edges between vertices of $\mathcal{C}$. In particular, the number of edges of a minimum complex in R is not determined by its vertex set. The natural counterpart to minimum complexes are complexes where all vertices are complete. Triangulations are an example.

The following lemma characterizes face sets of degree 3 in a polygonal complex $\mathcal{C}$. Here and in later sections, we denote with $U(B)$ the union of a subset B of faces of $\mathcal{C}$.

LEMMA 1. *Let B be any subset of faces of $\mathcal{C}$. Then B is of degree 3 if and only if $U(B)$ is a pre-triangle and B forms a minimum polygonal complex in $U(B)$.*

PROOF. Assume that B is of degree 3. As $U(B)$ has at least 3 corners, and each such corner is complete in B, we conclude that $U(B)$ has exactly 3 corners. That is, $U(B)$ is a pre-triangle, and the polygonal complex formed by B in $U(B)$ is minimum.

Conversely, assume that $U(B)$ is a pre-triangle, and that the polygonal complex formed by B in $U(B)$ is minimum. Then the degree of this complex equals the number of corners of $U(B)$, which is 3 because $U(B)$ is a pre-triangle. That is, B is of degree 3. ☐

4. THE M-SKELETON

Utilizing pre-triangulations, we now define a substructure for polygonal complexes, the so-called M-skeleton, which is the key to combinatorial projectivity. We introduce the class of face-reducible complexes and demonstrate that the surface theorem for pseudo-triangulations in [1] can be extended to this more general class. Face-reducible complexes are also the largest class so, as Section 6 shows.

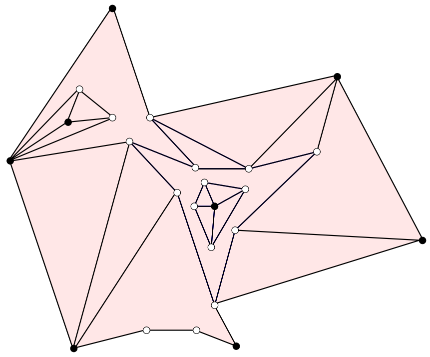

Figure 9: M-skeleton of the complex in Figure 7

Let $\mathcal{C}$ be some polygonal complex. A subset B of faces of $\mathcal{C}$ is called 3*-reducible* if there exists a superset $B' \supseteq B$ of faces of $\mathcal{C}$ such that B' is of degree 3. For example, if $\mathcal{C}$ is a pre-triangulation (in particular, any (pseudo-)triangulation), then for each face $f \in \mathcal{C}$, the singleton set $\{f\}$ is 3-reducible. A polygonal complex all whose faces have this property will be called a *face-reducible* complex. Such complexes may contain faces of any shape. See Figure 7. The reader is encouraged to check that this complex is face-reducible (using Figure 9 as an aid).

Let $\mathcal{C}$ be a face-reducible complex for the rest of this section. For a face $f \in \mathcal{C}$, we will denote with M_f a maximal subset of faces of $\mathcal{C}$ that contains f and that is of degree 3. Recall from Lemma 1 that $U(M_f)$ is a pre-triangle. We state and prove two basic properties of 'M_f-sets'.

LEMMA 2. *Let $f, g \in \mathcal{C}$. If $M_f \cap M_g \neq \emptyset$ then $M_f = M_g$.*

PROOF. Let $M_f \cap M_g \neq \emptyset$. We claim that $M_f \cup M_g$ is of degree 3. This implies $M_f = M_g$ by the maximality of these sets.

Put $I = U(M_f \cap M_g)$ and $Q = U(M_f \cup M_g)$. As, by Lemma 1, M_f and M_g form minimum complexes in $U(M_f)$ and $U(M_g)$, respectively, each vertex v of I is incomplete in the face set $M_f \cup M_g$ unless v is a corner of Q. Thus $M_f \cup M_g$ forms a minimum complex in Q. Let c be a corner of I. Then, by the minimality of the complex in Q, c is a corner of at least one of $U(M_f)$ and $U(M_g)$, and c is a corner of Q only if c is a corner of both $U(M_f)$ and $U(M_g)$. But each of the (at most 6) possible corners of Q is a corner of $U(M_f)$ or $U(M_g)$, and I has at least 3 corners. This implies that Q cannot have more than 3 corners. By Lemma 1, $M_f \cup M_g$ is of degree 3. ☐

COROLLARY 1. *The set M_f is unique for each face $f \in \mathcal{C}$. Moreover, the collection of these sets defines a partition of $\mathcal{C}$.*

PROOF. As we have $f \in M_f$ for each set M_f, the uniqueness of M_f follows from Lemma 2. To prove the partition property, let us write $f \sim g$ if $M_f = M_g$ holds for $f, g \in \mathcal{C}$. Then $\sim$ is an equivalence relation on $\mathcal{C}$. We show that $[f]_\sim = M_f$ are its equivalence classes. Clearly, $g \in [f]_\sim$ implies $M_g = M_f$ and thus $g \in M_f$. Conversely, $g \in M_f$ implies $g \in M_g \cap M_f$ and thus $M_g = M_f$ by Lemma 2, that is, $g \in [f]_\sim$ holds. ☐

From Corollary 1 and Lemma 1 we know that the regions $U(M_f)$, for all $f \in \mathcal{C}$, partition the underlying region $U(\mathcal{C})$ of $\mathcal{C}$ into pre-triangles in a unique way. We will term the resulting pre-triangulation the *M-skeleton* of $\mathcal{C}$. Figure 9 illustrates an example. Note that not all faces of this pre-triangulation are simply connected, although there are no 'holes' in the underlying region.

The following lemma is similar in spirit to Lemma 2. It will be needed in the design of a general flipping operation, in Section 9.

LEMMA 3. *Let $f, g \in \mathcal{C}$ be two faces that share some edge. Then either $M_f = M_g$ or the degree of $M_f \cup M_g$ is 4.*

PROOF. Assume $M_f \neq M_g$. Then M_f and M_g are disjoint, by Lemma 2. $U(M_f)$ and $U(M_g)$ have an edge e in common, because f and g do so, and $e = U(M_f) \cap U(M_g)$ holds because we have $M_f = M_g$, otherwise. Like in the proof of Lemma 2 (put $I = e$), each endpoint c of e leads to a loss of at least one of the 6 possible corners of $Q = U(M_f \cup M_g)$. Thus Q has at most 4 corners. Moreover, if an endpoint c of e is complete in $M_f \cup M_g$ but is a noncorner of Q, then c causes a loss of two possible corners of Q. We conclude that $M_f \cup M_g$ is of degree 4. $\square$

Face-reducible complexes enjoy a strong lifting property which stems from their M-skeletons and that we are going to describe next. Define a *polyhedral surface* as (the graph of) a continuous and piecewise-linear function φ whose domain is a polygonal region R. $\varphi(L)$ is called a *facet* of φ if L is a maximal interior-connected subset of R where φ is linear.

THEOREM 1. *Let $\mathcal{C}$ be a face-reducible polygonal complex. Further, let h be a vector assigning a height h_i to each complete vertex v_i of $\mathcal{C}$. For each choice of h, there exists a unique polyhedral surface φ with $\varphi(v_i) = h_i$ for all i, and such that $\varphi(f)$ is a subset of a facet of φ, for all faces $f \in \mathcal{C}$.*

PROOF. As $\mathcal{C}$ is face-reducible, the M-skeleton of $\mathcal{C}$ exists (and is unique). Let $\mathcal{PR}$ denote this pre-triangulation. The conditions required for the surface theorem in [1] to hold can be formulated as follows: For each incomplete vertex v of $\mathcal{C}$, there is a unique pre-triangle $\nabla = U(M_f)$ of $\mathcal{PR}$ such that (1) v is incomplete in M_f and (2) v lies in the convex hull of the 3 corners of ∇.

Consider condition (1). If v is not a vertex of $\mathcal{PR}$ then there is a unique ∇ such that v is internal to ∇. Otherwise, there is a unique ∇ where v is a noncorner, because an internal angle larger than π occurs in ∇ at v. In both cases, v is incomplete in the face set M_f with $\nabla = U(M_f)$, by Lemma 1. Condition (2) holds because the vertices of the convex hull of ∇ are corners of ∇. $\square$

5. PLANAR FACE SETS

Throughout this section, let $\mathcal{C}$ be a face-reducible complex. Theorem 1 makes explicit that $\mathcal{C}$ can be lifted to a polyhedral surface in space. However, not all edges of $\mathcal{C}$ might have their counterparts in this surface, for several reasons. For instance, the choice of the height vector h may force more than three complete vertices to be coplanar in the lifting. Also, geometric degeneracies of $\mathcal{C}$ may be the reason; consult Figure 2. Using our concept of 3-reducible face sets, we are able to characterize those edges of $\mathcal{C}$ which, even under 'generic' conditions, cannot be made to show up in the surface.

Let S be the set of all vertices of $\mathcal{C}$. Let S_ϵ be some replacement within distance ϵ of each vertex in S, for arbitrarily small $\epsilon > 0$. (We assumed S to be in general position, so the order type[2] of S_ϵ equals the order type of S.) An *ϵ-perturbation*, $\mathcal{C}_\epsilon$, of $\mathcal{C}$ is the polygonal complex with vertex set S_ϵ and with the same combinatorial structure as $\mathcal{C}$. As order types are preserved, a vertex is complete in $\mathcal{C}_\epsilon$ if and only if its original is complete in $\mathcal{C}$.

A subset B of faces of $\mathcal{C}$ is called *combinatorial planar* if, for any ϵ-perturbation $\mathcal{C}_\epsilon$ of $\mathcal{C}$, and for any height vector h for the complete vertices of $\mathcal{C}_\epsilon$, the subset $\varphi(B)$ of the surface φ for $\mathcal{C}_\epsilon$ and h lies in a single plane. Recall that φ uniquely exists by Theorem 1.

LEMMA 4. *Let B be any subset of faces of $\mathcal{C}$. Then B is combinatorial planar if and only if B is 3-reducible.*

PROOF. Let B be 3-reducible, and let $B' \supseteq B$ be a corresponding face set of degree 3. Then exactly 3 vertices are complete in B'. For every height vector h for $\mathcal{C}$, the plane through the corresponding 3 points in space determines a possible surface for B', because all other vertices of the complex formed by B' are incomplete in B', and thus are incomplete in $\mathcal{C}$. By the uniqueness of this surface, all surface vertices for B' have to lie in this plane. This reasoning remains true for any ϵ-perturbation $\mathcal{C}_\epsilon$ of $\mathcal{C}$. Therefore B' is combinatorial planar. As $B \subseteq B'$, the set B is combinatorial planar as well.

Now assume that B is combinatorial planar. Then the complex formed by B has to be minimum: Any vertex v that is complete in B but is not a corner of $U(B)$ destroys the coplanarity of $\varphi(B)$ in the surface φ for $\mathcal{C}_\epsilon$ and h, provided $\mathcal{C}_\epsilon$ and v's height in h are chosen appropriately. Let $B_L \supseteq B$ be a maximal set of faces of $\mathcal{C}$ such that B_L defines a minimum complex in $U(B_L)$. Then the vertices that are complete in B_L are just the corners of $U(B_L)$. All these corners are also complete in $\mathcal{C}$, by the maximality of B_L, and their heights can be chosen individually. On the other hand, as B is combinatorial planar, $\varphi(B)$ has to lie in a plane, in the surface φ for any $\mathcal{C}_\epsilon$ and h. This can only be achieved if there exists some set B', with $B \subseteq B' \subseteq B_L$, such that $U(B')$ has exactly 3 corners whose corresponding 3 points in space span the afore-mentioned plane. By Lemma 1, B' is of degree 3. Thus B' witnesses that B is 3-reducible. $\square$

6. COMBINATORIAL PROJECTIVITY

We now define a notion of projectivity for arbitrary polygonal complexes which outrules complexes that are projective (or non-projective) only because of geometric artifacts.

Consider some polygonal complex $\mathcal{C}$. Let $\mathcal{C}_r$ be a random[3] ϵ-perturbation of $\mathcal{C}$. We call the complex $\mathcal{C}$ *combinatorial projective* if, with probability 1, a polygonal surface exists for $\mathcal{C}_r$ whose set of edges projects *exactly* to the set of edges of $\mathcal{C}_r$.

Our aim is to characterize the class of combinatorial projective complexes. The following assertion can be proven by arguments similar to the proof of Lemma 4. Let $\mathcal{C}$ be a face-reducible complex, and let k be the degree of $\mathcal{C}$. Let $h_r \in [0,1]^k$ be a random height vector for the k complete vertices of $\mathcal{C}_r$.

[2]The *order type* of a set S of points in the plane assigns to each ordered triple $\{p, q, r\} \subset S$ its orientation, either clockwise or counterclockwise.

[3]By *random* we mean drawn from uniform distribution.

LEMMA 5. *Let B be a subset of faces of $\mathcal{C}$. If B is not combinatorial planar then, in the (by Theorem 1 unique) surface φ for $\mathcal{C}_r$ and h_r, the set $\varphi(B)$ is not coplanar with probability* 1.

The next lemma formulates the main observation for a characterization of combinatorial projective complexes.

LEMMA 6. *Let $\mathcal{C}$ be a face-reducible complex. Then, with probability* 1*, the edges in the surface for $\mathcal{C}_r$ and h_r bijectively correspond to the edges of the M-skeleton of $\mathcal{C}$.*

PROOF. Let e be any internal edge of $\mathcal{C}$. To prove the lemma, it obviously suffices to consider such edges. Let f_1 and f_2 be the two faces of $\mathcal{C}$ that share e.

Assume that e is not an edge of the M-skeleton, that is, e is internal to $U(M_f)$ for some face f of $\mathcal{C}$. Then the set $\{f_1, f_2\}$ is 3-reducible, because $\{f_1, f_2\} \subseteq M_f$ and M_f is of degree 3. So, by Lemma 4, $\{f_1, f_2\}$ is combinatorial planar. This implies that e does not correspond to an edge of the surface for $\mathcal{C}_r$ and h_r.

Conversely, assume that e is an edge of the M-skeleton. Then $M_{f_1} \neq M_{f_2}$. The existence of some face set B of degree 3 and with $\{f_1, f_2\} \subseteq B$ implies $M_{f_1} = M_{f_2}$ which is a contradiction. Therefore $\{f_1, f_2\}$ is not 3-reducible. By Lemma 4, $\{f_1, f_2\}$ is not combinatorial planar. Moreover, by Lemma 5, edge e corresponds to an edge of the surface for $\mathcal{C}_r$ and h_r with probability 1. $\square$

THEOREM 2. *A polygonal complex $\mathcal{C}$ is combinatorial projective if and only if $\mathcal{C}$ coincides with its M-skeleton.*

PROOF. For the class of face-reducible complexes, the assertion is true by Lemma 6. To complete the proof, we assume that $\mathcal{C}$ is not face-reducible and argue that $\mathcal{C}$ then cannot be combinatorial projective.

Faces that are not 3-reducible (of which $\mathcal{C}$ now is supposed to contain at last one) have more than 3 corners. So at least one internal edge can be added to each such face f without changing any vertex from incomplete to complete. We keep adding such edges until a face-reducible complex $\mathcal{C}'$ is obtained. Let $B(f)$ denote the set of faces of $\mathcal{C}'$ that a face f of $\mathcal{C}$ splits into. Then not all sets $B(f)$ are 3-reducible in $\mathcal{C}'$, because all singletons $\{f\}$ would be 3-reducible in $\mathcal{C}$, otherwise. So, by Lemma 4, there is a set $B(f)$ that is not combinatorial planar. In the surface φ' for $\mathcal{C}'_r$ and h_r, the set $\varphi'(B(f))$ is not part of a single facet of φ' with probability 1, by Lemma 5. By the uniqueness of the surface theorem, there is no other way of constructing the required surface for $\mathcal{C}_r$ and h_r. This implies that $\mathcal{C}$ is not combinatorial projective. $\square$

COROLLARY 2. *Every combinatorial projective complex is a pre-triangulation.*

The second part of the proof of Theorem 2 also shows the following:

COROLLARY 3. *Theorem 1 cannot be extended beyond the class of face-reducible polygonal complexes.*

Figure 7 shows an example for a face-reducible complex which is not combinatorial projective; the complex does not coincide with its M-skeleton (see Figure 9). A pseudo-triangulation which is not combinatorial projective is shown in Figure 11a (upper drawing); the four pseudo-triangles form a face set which is combinatorial planar. However, combinatorial planarity of a set of two or more faces does not necessarily destroy combinatorial projectivity. For example, the pre-triangle in Figure 8d constitutes both a combinatorial projective complex and a combinatorial planar face set of size three. Note also that, although each combinatorial projective complex is a pre-triangulation, not all of its *faces* need to be pre-triangles; see the pentagonal star in Figure 8c.

7. COMBINATORIAL REGULARITY

In this section we focus on convex projection surfaces. We introduce the classes of combinatorial regular complexes and face-honest complexes and elaborate on their interrelation. Face-honest complexes are of interest for two more reasons. Loosely speaking, each face of such a complex can be lifted to an individual plane in space. Moreover, these complexes allow for easy modification with flipping operations; see Section 12.

Let φ be some polyhedral surface. An edge e of φ is called *convex* if there exists a line segment ℓ that nowhere lies above φ and that intersects e at exactly one point interior to both e and ℓ. Edges of φ which are not convex are called *reflex*. Note that the image $\varphi(e')$ of each boundary edge e' of the domain of φ is a convex edge.

A polygonal complex $\mathcal{C}$ is called *combinatorial regular* if $\mathcal{C}$ is combinatorial projective and there exists some ϵ-perturbation $\mathcal{C}_\epsilon$ and some height vector h for $\mathcal{C}$ such that all edges in the surface for $\mathcal{C}_\epsilon$ and h are convex. By Corollary 2, combinatorial regular complexes are pre-triangulations. Known examples are (constrained) Delaunay triangulations [16], and more generally, the constrained regular pseudo-triangulations [1]. Schlegel diagrams [11] and thus Voronoi diagrams are not combinatorial regular (though they are well known to be regular in the classical sense) because these complexes are not combinatorial projective.

LEMMA 7. *Let $\mathcal{C}$ be a combinatorial regular complex. Then each internal vertex of $\mathcal{C}$ is complete.*

PROOF. Let $\mathcal{C}_\epsilon$ and h witness the combinatorial regularity of $\mathcal{C}$. Consider an incomplete vertex v of $\mathcal{C}$, and let f be the unique face of $\mathcal{C}$ where v is a noncorner. To get a contradiction, suppose v is an internal vertex of $\mathcal{C}$. As there is an internal angle greater than π in f at v, there exists some line segment $\ell \subset U(\mathcal{C})$ such that ℓ crosses all the edges of $\mathcal{C}$ incident to v. But $\mathcal{C}$ is combinatorial regular, so the corresponding edges in the surface φ for $\mathcal{C}_\epsilon$ and h are all convex. Therefore $\varphi(f)$ cannot be part of a single facet of φ – a contradiction to the definition of φ. $\square$

A polygonal complex is termed *face-honest* if $M_f = \{f\}$ holds for each of its faces f. Every face-honest complex $\mathcal{C}$ coincides with its M-skeleton; see Figure 9 for an illustration. Therefore, $\mathcal{C}$ is a pre-triangulation that is combinatorial projective, by Theorem 2. Moreover, all faces of $\mathcal{C}$ are indeed pre-triangles. On the other hand, even when a pseudo-triangulation is combinatorial regular, it is not necessarily face-honest. Figure 8d reveals this fact. Triangulations are always face-honest: All their vertices are complete, so their individual triangles constitute the only face sets of degree 3. See Figure 3 for the interaction of face-honest complexes and combinatorial regular complexes with other classes.

Below we derive some results for the case where the underlying region $U(\mathcal{C})$ of the complex $\mathcal{C}$ in question is a polygon.

LEMMA 8. *Let $\mathcal{C}$ be a face-reducible complex, and suppose each internal vertex of $\mathcal{C}$ is complete. If $U(\mathcal{C})$ is a polygon then $\mathcal{C}$ is a face-honest pseudo-triangulation.*

PROOF. Let f be a face of $\mathcal{C}$. As $\mathcal{C}$ is face-reducible, the set M_f exists. Each vertex v that is incomplete in M_f is also incomplete in $\mathcal{C}$, and thus v is a vertex of $U(\mathcal{C})$, by assumption. Therefore, if $U(\mathcal{C})$ is required to be a polygon, then the pre-triangle $U(M_f)$ is just a pseudo-triangle, and M_f contains a single face, $U(M_f) = f$. $\square$

Lemmas 7 and 8 combine to the following statement.

THEOREM 3. *Let $\mathcal{C}$ be a combinatorial regular complex whose underlying region is a polygon. Then $\mathcal{C}$ is a face-honest pseudo-triangulation.*

COROLLARY 4. *Let $\mathcal{PT}$ be a pseudo-triangulation without internal vertices in a polygon R. Then $\mathcal{PT}$ is face-honest.*

PROOF. As $\mathcal{PT}$ contains no internal vertices, each internal edge of $\mathcal{PT}$ is a diagonal of R. This implies that $\mathcal{PT}$ is combinatorial regular. The assertion now follows from Theorem 3. $\square$

8. LOCALLY CONVEX SURFACES

We show next that combinatorial regular complexes – and thus certain pre-triangulations – arise from graphs of locally convex functions on polygonal domains. This generalizes results in [1] where locally convex functions are introduced and utilized in the context of pseudo-triangulations.

Let a polygonal domain R be given. A surface φ on R is called *locally convex* if each edge of φ is convex.[4] Let now $S \subset R$ be a finite set of points that includes all the vertices of R. Further, let h be a vector that assigns an upper height bound $h(v)$ to each point $v \in S$. We define F_h as the maximal (i.e., highest) locally convex surface on R that satisfies $F_h(v) \leq h(v)$ for all $v \in S$.

The surface F_h exists and is unique: Any surface φ that consists of a single facet, $\varphi(R)$, is locally convex and can be lowered to satisfy h, and F_h is the pointwise maximum of all surfaces that are locally convex and satisfy h. The facets of F_h project to the faces of a polygonal complex in R which we denote by $\mathcal{C}(F_h)$.

LEMMA 9. *All vertices of $\mathcal{C}(F_h)$ are included in S. Moreover, $F_h(v) = h(v)$ holds for each complete vertex v of $\mathcal{C}(F_h)$.*

PROOF. Let w be a vertex of $\mathcal{C}(F_h)$. Assume $w \notin S$. Then w has to be complete in $\mathcal{C}(F_h)$: If w lies on the boundary of R then clearly no face of $\mathcal{C}(F_h)$ has an internal angle larger than π at w. If w is an internal vertex of $\mathcal{C}(F_h)$ then the same is true, because the image $F_h(e)$ of each edge e incident to w is a convex edge. But for any complete vertex v of $\mathcal{C}(F_h)$ we must have $v \in S$ and $F_h(v) = h(v)$, because otherwise the height of v can be increased without violating the local convexity of the surface or the restrictions in h, in contradiction to the maximality of F_h. $\square$

[4] Equivalently, $\varphi(\ell)$ defines a convex function, for every line segment $\ell \subset R$ that does not cross the boundary of R. This definition of locally convexity is used in [1].

Let h_ϵ be a vector obtained from h by perturbing each coordinate of h by at most ϵ. We call h *generic* (for R and S) if there exists some $\epsilon > 0$ such that $\mathcal{C}(F_{h_\epsilon}) = \mathcal{C}(F_h)$ holds for all possible vectors h_ϵ. Observe that if R is a convex polygon then F_h is the lower convex hull [19] of the spatial point set $\{\binom{v}{h(v)} \mid v \in S\}$. If, in addition, h is generic then $\mathcal{C}(F_h)$ is a regular triangulation [9]. This fact generalizes for arbitrary polygonal regions R as follows.

THEOREM 4. *Let h be a generic height vector. Then $\mathcal{C}(F_h)$ is a pre-triangulation that is combinatorial regular.*

PROOF. By Lemma 9, $F_h(v) = h(v)$ holds for each complete vertex v of $\mathcal{C}(F_h)$. Thus F_h is the unique surface for $\mathcal{C}(F_h)$ and h predicted in Theorem 1. We show that $\mathcal{C}(F_h)$ coincides with its M-skeleton. This implies the assertion because then, by Theorem 2, $\mathcal{C}(F_h)$ is combinatorial projective, and h is a height vector for $\mathcal{C}(F_h)$ such that all edges of F_h are convex.

Let g be an arbitrary face of $\mathcal{C}(F_h)$. As h is generic, there exists an $\epsilon > 0$ and a maximal subset $B \ni g$ of faces of $\mathcal{C}(F_h)$ such that the set $F_{h_\epsilon}(B)$ lies in a plane for all h_ϵ. (Note that $B = \{g\}$ is possible.) This implies that exactly 3 vertices in B are complete. That is, B is of degree 3, which leads us to conclude $B = M_g$ by the maximality of B. But each facet of F_h projects to a single face of $\mathcal{C}(F_h)$, which shows that $\mathcal{C}(F_h)$ coincides with its M-skeleton. $\square$

COROLLARY 5. *If the domain R of F_h is a polygon, and h is generic, then $\mathcal{C}(F_h)$ is a face-honest pseudo-triangulation.*

PROOF. Combine Theorem 3 and Theorem 4. $\square$

9. A GENERAL FLIPPING SCHEME

Flip operations are a common tool for locally modifying polygonal complexes. By now, the most general class where flip operations have been defined in the literature are the pseudo-triangulations. (An exception are diagonal partitions of convex polygons studied in [13].) The repertoire includes Lawson flips [15], exchanging flips [18, 22], and removing flips and their inverses [1, 17]. Still, several interesting classes of complexes, in particular the combinatorial regular complexes whose relevance as optimal surfaces is documented in Section 8, cannot be reached and modified with existing flip operations. The combinatorial regular complex in Figure 8c serves as an example.

Let $\mathcal{C}$ be a face-reducible complex. Below we define a flip operation that applies to any internal edge e of $\mathcal{C}$, in a way such that face-reducibility is retained in the resulting complex. All the flips known for triangulations and pseudo-triangulations arise as special cases of this operation.

OPERATION FLIP(e). *Let f and g be the two faces of $\mathcal{C}$ incident to e. Consider the subcomplex $\mathcal{C}_e = M_f \cup M_g$ of $\mathcal{C}$. Choose a height for each complete vertex of $\mathcal{C}_e$ such that, if $M_f \neq M_g$, the edge $\varphi(e)$ is reflex in the surface φ for $\mathcal{C}_e$. Let the vector h contain these heights and, in addition, the entry ∞ for each incomplete vertex of $\mathcal{C}_e$. Replace $\mathcal{C}_e$ by $\mathcal{C}(F_h)$.*

FLIP(e) deletes a subset E of edges of $\mathcal{C}$ and creates a subset E' of edges disjoint from E, in a unique way. Note that $E \neq \emptyset$ because FLIP(e) always deletes the edge e. That is, each internal edge of $\mathcal{C}$ is flippable with this operation. We might have $E' = \emptyset$ if the flip is removing. The cardinalities of E and E' may be large, though. We offer a detailed discussion of the flip types covered by the operation FLIP in Section 10.

LEMMA 10. *The subcomplex $\mathcal{C}(F_h)$ constructed by FLIP(e) is combinatorial regular.*

PROOF. By Theorem 4 it suffices to show that h was chosen in a generic way. If $M_f = M_g$ then $M_f \cup M_g$ is of degree 3. For any entries in h for these 3 complete vertices, we obviously have $\mathcal{C}(F_h)$ unchanged. If $M_f \neq M_g$ then $M_f \cup M_g$ is of degree 4, by Lemma 3. The corresponding 4 entries in h were chosen such that e yields a reflex edge in the surface for $M_f \cup M_g$. This implies that the respective 4 surface vertices are not coplanar. So $\mathcal{C}(F_h)$ remains unchanged when h is ϵ-perturbed. We conclude that h is generic. $\square$

THEOREM 5. *The class of face-reducible complexes is closed under the operation FLIP. Moreover, pre-triangulations are the smallest class having this property.*

PROOF. Let $\mathcal{C}$ be some face-reducible complex in the polygonal region $R = U(\mathcal{C})$. Let e be an internal edge of $\mathcal{C}$, and denote with $\mathcal{C}'$ the complex that is obtained from $\mathcal{C}$ by applying FLIP(e).

The subcomplex $\mathcal{C}(F_h)$ of $\mathcal{C}'$ constructed by FLIP(e) is combinatorial regular, by Lemma 10. Thus $\mathcal{C}(F_h)$ is a pre-triangulation in $Q = U(M_f \cup M_g)$, for f and g being the faces incident to e. In the complement $R \setminus Q$, the complex $\mathcal{C}'$ coincides with $\mathcal{C}$. So each face $b \in \mathcal{C}'$ external to Q is 3-reducible by the assumption that $\mathcal{C}$ is face-reducible. That is, the part of $\mathcal{C}'$ in $R \setminus Q$ is a face-reducible complex.

If, in addition, $\mathcal{C}$ is a pre-triangulation then there exists a partition of $\mathcal{C}$ into face sets B_i such that $U(B_i)$ is a pre-triangle and no edge of $\mathcal{C}$ is internal to $U(B_i)$, for each set B_i. So B_i is of degree 3, which implies $B_i \subseteq M_b$ if $b \in B_i$ holds for a face b external to Q. This implies that B_i is not affected by FLIP(e) in this case. Consequently, the part of $\mathcal{C}'$ in $R \setminus Q$ is a pre-triangulation.

But the concatenation of two pre-triangulations (or of two face-reducible complexes) is a complex of the same type, as is easily seen from the definition of such complexes.

To complete the proof, we observe that there exist Delaunay triangulations constrained by R that can be flipped to a pre-triangulation which is neither combinatorial projective nor a pseudo-triangulation. Figure 11b exemplifies this fact. In view of the containment relations in Figure 3, this shows that any class closed under FLIP has to include the class of pre-triangulations. $\square$

10. INSTANCES OF FLIP

Let us discuss the effect of the operation FLIP in different scenarios. First of all, the interested reader may convince himself that FLIP is able to simulate the exchanging flips in Figure 4 and the removing flips in Figure 5; see Subsection 1.2. In the figures below, the edge e to be flipped is shown in bold, and only the region $U(M_f \cup M_g)$ that gets restructured is displayed. Here f and g are the two faces incident to e. Edges created by the respective flip are drawn as dashed lines.

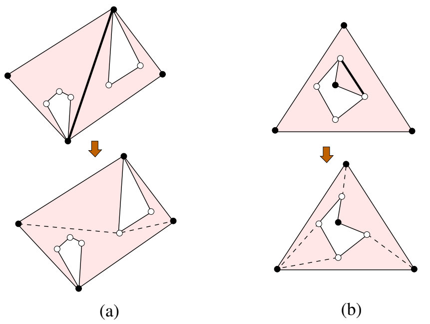

Figure 10: Two exchanging flips

Figure 10 illustrates how the exchanging flip generalizes for pre-triangles. The degree of the face set $M_f \cup M_g$ is 4, and the status of all vertices (complete or incomplete) is preserved, as in a usual exchanging flip. However, more than one edge is created in flip (a) as well as in flip (b). Moreover, even more than one edge may be deleted, as Figure 12a shows. Observe that FLIP is not necessarily a symmetric operation: Flipping one edge (or all edges in any order) that just have been created need not give the initial complex.

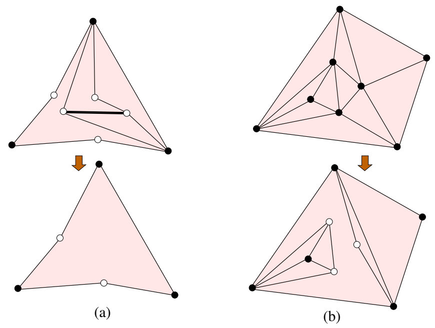

Figure 11: (a) Removing flip (b) Many flips

Figure 11a shows a generalization of the removing flip. The degree of $M_f \cup M_g$ is 3 for this type of flip. The edges and (incomplete) vertices internal to $U(M_f \cup M_g)$ form a minimum complex. They all get removed, and a single pre-triangular face remains. New edges are never created. It should be noted that incomplete vertices and their incident edges may get removed in a generalized exchanging flip as well.

A natural question is why FLIP(e) is based on maximal face sets of degree 3 rather than on single faces of $\mathcal{C}$. Figure 12 provides an answer. The flip (a) is in accordance with the definition of FLIP and retains face-reducibility. In flip (b) that is solely based on the two faces incident to the flipped edge, face-reducibility is lost, and with it, the applicability of the surface theorem (Theorem 1) and its advantages, which will be exploited next.

11. DELAUNAY MINIMUM COMPLEX

By the well-known paraboloid lifting function $h(x) = x^2$, Delaunay triangulations in a convex region R correspond to maximal locally convex surfaces generated by the vertex heights $h(x)$. (In this easy case, such surfaces are just lower convex hulls of the lifted vertices.) This correspon-

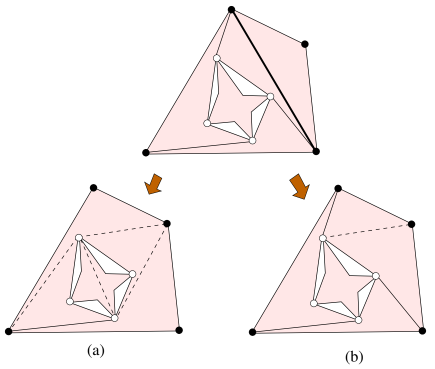

Figure 12: (a) Correct flip (b) 'Wrong' flip

dence is extended to nonconvex polygons R in [1] and leads to Delaunay triangulations constrained by (the edges of) R. We define below a unique complex of *smallest* combinatorial size in an *arbitrary* polygonal region R that still shows the desired 'Delaunay properties': an analog of local Delaunayhood, and the reachability by improving flip operations.

Let p be a height vector for the vertices of R such that $p(c) = c^2$ for each corner c of R, and $p(v) = \infty$ for each noncorner v of R. Consider the complex $\mathcal{D}_R = \mathcal{C}(F_p)$. By our general position assumption for R, no four vertices of R are cocircular. This implies that the vector p is generic. So, by Theorem 4, the complex $\mathcal{D}_R$ is a pre-triangulation that is combinatorial regular. Moreover, $\mathcal{D}_R$ is a minimum complex in R because, by Lemma 9, its only vertices are those of R and each noncorner of R is incomplete in $\mathcal{D}_R$. We term $\mathcal{D}_R$ the *Delaunay minimum complex* of R. Figure 13 gives an example.

Observe that $\mathcal{D}_R$ differs from the constrained Delaunay triangulation [16] of R, unless R is a convex polygon and $\mathcal{D}_R$ is the classical Delaunay triangulation [2, 10] of the vertices of R. In case R is a polygon, $\mathcal{D}_R$ is a minimum pseudo-triangulation which is face-honest, by Corollary 5. The latter structure has been also considered in [20] and is a constrained regular pseudo-triangulation as defined in [1].

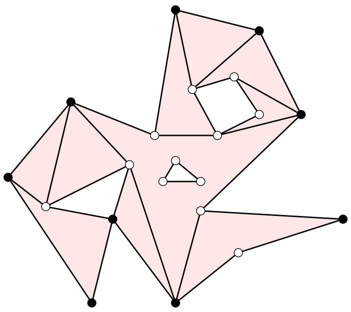

Figure 13: Delaunay minimum complex

The class of all possible face-reducible complexes in a given region R is connected under the operation FLIP in the following sense.

THEOREM 6. *Let a polygonal region R be given. Each face-reducible complex in R can be transformed to the Delaunay minimum complex $\mathcal{D}_R$ by using finitely many operations FLIP.*

PROOF. Let $\mathcal{C}$ be a face-reducible complex in R. We construct a finite sequence of flips that transforms $\mathcal{C}$ into $\mathcal{D}_R$. Let h be a height vector for the vertices of $\mathcal{C}$, with $h(c) = c^2$ for each corner c of R, and with heights for the remaining vertices of $\mathcal{C}$ being sufficently large such that $\mathcal{C}(F_h) = \mathcal{D}_R$.

Apply h to the complete vertices of $\mathcal{C}$. This defines a unique surface φ on R by Theorem 1. If φ contains reflex edges then choose such an edge $\varphi(e)$, transform $\mathcal{C}$ to a new complex $\mathcal{C}'$ by performing FLIP(e) for the edge e of $\mathcal{C}$, and put $\mathcal{C} = \mathcal{C}'$ and repeat.

The complexes $\mathcal{C}_i$ created above are face-reducible, by Theorem 5. So the procedure is well-defined. We first prove that the created surfaces φ_i are pairwise different[5]. Let FLIP(e) transform $\mathcal{C}_i$ to $\mathcal{C}_{i+1}$. Denote by Q the region restructured by FLIP(e), and let $\mathcal{S}$ and $\mathcal{S}'$, respectively, be the subcomplex of $\mathcal{C}_i$ and $\mathcal{C}_{i+1}$ in Q. By adding all internal edges of $\mathcal{S}'$ to $\mathcal{C}_i$ (and splitting faces of $\mathcal{C}_i$ accordingly) we obtain a polygonal complex A in R. Extend h to a height vector h_A for the complete vertices of A such that the surface for A and h_A is φ_i. (This is possible although A is not face-reducible, in general.) Similarly, extend h to h'_A such that the surface for A and h'_A is φ_{i+1}. Then $h'_A \leq h_A$ holds (element-wise, with strict inequality for some elements) because e is the only internal edge for $\mathcal{S}$ and $\varphi_i(e)$ is reflex, and all internal edges e' of $\mathcal{S}'$ have a convex image $\varphi_{i+1}(e')$. But surface heights can be shown to strictly decrease with the height vector (see [1], Lemma 5.2). This implies that the surfaces φ_i are pairwise different.

In fact, the complexes $\mathcal{C}_i$ are pairwise different as well: If $\mathcal{C}_i$ and $\mathcal{C}_j$ have different sets of complete vertices then clearly $\mathcal{C}_i \neq \mathcal{C}_j$. Otherwise, by the uniqueness of φ_i for fixed h, $\varphi_i \neq \varphi_j$ implies $\mathcal{C}_i \neq \mathcal{C}_j$. Observe next that no new vertices are created by the operation FLIP. But the number of polygonal complexes with vertices from a fixed set of n points is finite (in fact, exponential in n). We conclude that the procedure above terminates with a surface φ^* where all edges are convex.

In the corresponding complex $\mathcal{C}^*$, the only vertices that are complete are the corners of R: Any other complete vertex v of $\mathcal{C}^*$ would be incident to edges whose images are reflex in φ^*, by the choice of $h(v)$. This implies $\varphi^* = F_h$ by the uniqueness of F_h. Observe that $\mathcal{C}(F_h) = \mathcal{D}_R$ is the M-skeleton of $\mathcal{C}^*$. If $\mathcal{C}^* \neq \mathcal{D}_R$ then $\mathcal{C}^*$ contains internal edges (and possibly, internal and incomplete vertices) that have no counterparts in F_h. As long as such an edge e exists, we apply FLIP(e). For the two faces f and g incident to e, we have $M_f = M_g$. Therefore, FLIP(e) removes all edges internal to $U(M_f)$ (in particular, the edge e) and creates no edges. The number of such flips is bounded by the number of edges of $\mathcal{C}^*$. The surface φ^* remains unchanged, and the complex $\mathcal{D}_R$ is the result. $\square$

Theorem 6 is rather general. Flips can be applied in any order to edges with reflex images, and we even can drop the requirement that such edges have to be flipped before edges without surface counterparts. Moreover, when starting with a pre-triangulation this class is never left; see Theorem 5. We remark that Theorem 6 is a generalization of the optimality theorem for pseudo-triangulations in [1].

As a special case, let R be a polygon, and let $\mathcal{PT}$ be some minimum pseudo-triangulation in R without internal ver-

[5] Thanks go to Paco Santos for pointing out, in the context of pseudo-triangulations, the following elegant argument.

tices. Then $\mathcal{PT}$ is face-honest by Corollary 4. Let e be an edge of $\mathcal{PT}$ which is incident to two pseudo-triangles ∇ and ∇'. We have $M_\nabla = \nabla \neq \nabla' = M_{\nabla'}$ and the region affected by FLIP(e) is $Q = \nabla \cup \nabla'$. Moreover, Q has exactly 4 corners by Lemma 3. So FLIP(e) is an exchanging flip; see Section 10. Let h be the height vector used above to define $D_R = \mathcal{C}(F_h)$. We may call FLIP(e) a *Delaunay flip* if the edge $\varphi(e)$ in the surface φ for $\mathcal{PT}$ and h is reflex. As FLIP(e) transforms $\mathcal{PT}$ to a minimum pseudo-triangulation of R, and D_R is a complex of this type, we obtain:

COROLLARY 6. *Any minimum pseudo-triangulation of a polygon R can be flipped to D_R by repeated application of Delaunay flips and without leaving the class of minimum pseudo-triangulations.*

If Delaunay flips are performed in a well-chosen order then $O(n^2)$ flips are sufficient, as can be seen by adapting a result in [1]. Note finally that the edges of D_R do not form a subset of the edges of the constrained Delaunay triangulation [16] of R, in general.

12. FACE-HONEST FLIPPING

Throughout this section, we restrict attention to polygonal complexes whose underlying region R is a polygon. We show that any triangulation $\mathcal{T}$ in R can be flipped to the complex $\mathcal{C}(F_h)$, for any given (generic) height vector h for the vertices of $\mathcal{T}$, in a way such that each intermediate complex is a pseudo-triangulation. (We allow $\mathcal{T}$ to contain internal vertices.) In fact, all obtained complexes are face-honest and thus are combinatorial projective, except in cases where an internal vertex has to be removed in the subsequent flip and face-honesty is impossible to keep. This guarantees that only standard flips for pseudo-triangulations are performed.

OBSERVATION 1. *Let $\mathcal{PT}$ be a pseudo-triangulation in R, and let v be an internal and complete vertex of $\mathcal{PT}$ with at least 4 incident edges. Then at least one such edge can be flipped by exchange such that no edge incident to v is created.*

PROOF. Let Q be the union of the faces of $\mathcal{PT}$ incident to v, and let $\mathcal{C}$ be the pseudo-triangulation in Q defined by these faces. For an edge vw of $\mathcal{C}$, define $\alpha(vw)$ as the internal angle at v of the union of the two faces of $\mathcal{C}$ adjacent in vw. Imagine that removing edge flips are applied to all edges vw where w is a non-corner of Q that is complete in $\mathcal{C}$. Each such flip makes a unique vertex incomplete. This vertex is different from v (and thus is w) because $\alpha(vw) < \pi$ holds by our assumption on w. So v stays complete in the obtained complex, and therefore at least 3 edges of $\mathcal{C}$ are still incident to v. Such an edge vw can be flipped in $\mathcal{C}$ by exchange, and without creating an edge incident to v, if $\alpha(vw) < \pi$. But this is guaranteed for at least one such edge, because we assumed v to be incident to at least 4 edges of $\mathcal{C}$. $\square$

THEOREM 7. *Let R be a polygon, let $\mathcal{T}$ be a triangulation in R, and let h be a generic height vector for the vertices of $\mathcal{T}$. Then $\mathcal{T}$ can be transformed by FLIP to the complex $\mathcal{C}(F_h)$, in a way such that all intermediate complexes are face-honest pseudo-triangulations, except immediately before a vertex is removed.*

PROOF. Denote with $\mathcal{PT}$ the current complex obtained from flipping; initially, $\mathcal{PT} = \mathcal{T}$. Recall that all vertices of $\mathcal{T}$ are complete. We construct a desired sequence of flips that terminates with $\mathcal{PT} = \mathcal{C}(F_h)$. Let φ denote the unique surface that results from $\mathcal{PT}$ when h is applied to all complete vertices. We say that an internal vertex v of $\mathcal{PT}$ fulfills the *hull condition* if $\varphi(v)$ lies strictly below the lower convex hull of its neighbored vertices in φ.

Step (1) As long as there exists an internal vertex v of $\mathcal{PT}$ that violates the hull condition, we do the following. Apply exchanging flips to the edges incident to v until only 3 such edges remain. This is possible by Observation 1. After each flip all the internal vertices of $\mathcal{PT}$ are complete, because this was true before the flip, and the flip was exchanging. Two more applications of FLIP to v's remaining edges make v incomplete (and $\mathcal{PT}$ temporary non-projective) and remove v, respectively, leaving all internal vertices complete again.

Step (2) While there exists an edge e in $\mathcal{PT}$ such that $\varphi(e)$ is reflex, do the following: Apply FLIP(e). If e is exchanged then no vertex alters its completeness status. If e gets removed then an endpoint of e which is a vertex of R becomes incomplete: Each internal vertex fulfills the hull condition before the flip, and $\varphi(e)$ was reflex. So all internal vertices stay complete after FLIP(e). If some internal vertices now violate the hull condition then repeat from Step (1).

The total number of flips performed in Step (1) is clearly $O(n^2)$, if $\mathcal{T}$ contains n vertices. In Step (2) only edges with a reflex image are flipped. The number of these flips is finite by the arguments in the proof of Theorem 6. In each created complex all internal vertices are complete, except immediately before a vertex is removed. So, with these exceptions, each such complex is a face-honest pseudo-triangulation, by Lemma 8.

In particular the final complex, $\mathcal{PT}^*$, is face-honest and therefore combinatorial projective. The surface φ^* for $\mathcal{PT}^*$ contains no reflex edge. This implies that, for each edge e of $\mathcal{PT}^*$, its image $\varphi^*(e)$ is a convex edge. As only vertices that violate the hull condition are removed, and such vertices cannot belong to $\mathcal{C}(F_h)$, we conclude $\mathcal{PT}^* = C(F_h)$. $\square$

We conjecture that only quadratically many flips are performed in Step (2). This would yield a flip sequence of total length $O(n^2)$. A result related to Theorem 7 has been proved in [12]. Namely, any triangulation $\mathcal{T}$ can be flipped to the complex $\mathcal{C}(F_h)$ such that no edges with convex images in the corresponding surface are flipped, and such that all intermediate complexes are combinatorial projective (with the above exceptions before vertex removal).

13. CONCLUDING REMARKS

We have introduced the concept of pre-triangulations, a relaxation of triangulations that goes beyond the frequently used concept of pseudo-triangulations. Appearing less intuitive at first sight, pre-triangulations turned out to be more natural than pseudo-triangulations in questions concerning liftability and flippability of polygonal complexes – even in the most simple case where the underlying region is a convex polygon.

One of the crucial tools used in our developments is the completeness status of the vertices of a polygonal complex. (In)completeness of a vertex depends on the (non)convexity of its incident internal angles and thus on the order type of the vertices of the complex in the end. This enabled us to characterize complexes which exhibit the 'robust' lifting property we called combinatorial projectivity.

We did not address algorithmic issues in this paper. Deciding combinatorial projectivity of a given polygonal complex in an efficient way is one of them. In view of Theorem 2 this question reduces to identifying the M-skeleton of the complex or detecting its non-existence. As a simple approach we may try to construct – after ϵ-perturbing the vertices – a surface for the complex in question using random heights for its complete vertices. The surface construction mainly means resolving a system of n linear equations [1] where n is the number of vertices. If the system has a solution then the edges of the resulting unique surface correspond to the M-skeleton with probability one (see Lemma 6). Otherwise, the M-skeleton does not exist.

Other relevant algorithmic questions are finding short flip sequences within particular complex classes, and implementing the operation FLIP efficiently. These questions are discussed in detail in a separate paper [4]. Emphasis is laid on constructing locally convex functions, and Delaunay-type complexes in particular.

14. REFERENCES

[1] O.Aichholzer, F.Aurenhammer, H.Krasser, P.Brass. *Pseudo-triangulations from surfaces and a novel type of edge flip.* SIAM Journal on Computing 32 (2003), 1621-1653.

[2] F.Aurenhammer. *Voronoi diagrams – a survey of a fundamental geometric data structure.* ACM Computing Surveys 23 (1991), 345-405.

[3] 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.

[4] F.Aurenhammer, T.Hackl, H.Krasser. *Short flip sequences for constructing the Delaunay triangulation.* Manuscript, 2005.

[5] 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.

[6] A.Brøndsted. An Introduction to Convex Polytopes. Springer Verlag, 1983.

[7] H.Crapo, W.Whiteley. *Plane self stresses and projected polyhedra.* Structural Topology 20 (1993), 55-78.

[8] H.Edelsbrunner. Algorithms in Combinatorial Geometry. Springer-Verlag, 1987.

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

[10] 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.

[11] B.Grünbaum. Convex Polytopes. Wiley Interscience, London, 1967.

[12] T.Hackl. Manipulation of Pseudo-Triangular Surfaces. Master Thesis, Institute for Theoretical Computer Science, University of Technology, Graz, Austria, 2004.

[13] C.Huemer. Flip Operations for Geometric and Combinatorial Objects. Master Thesis, Institute for Theoretical Computer Science, University of Technology, Graz, Austria, 2003.

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

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

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

[17] D.Orden, F.Santos. *The polytope of non-crossing graphs on a planar point set.* Discrete & Computational Geometry 33 (2005), 275-305.

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

[19] F.P.Preparata, M.I.Shamos. Computational Geometry: An Introduction. Springer-Verlag, New York, 1985.

[20] G.Rote, A.Schulz. *A pointed Delaunay pseudo-triangulation of a simple polygon.* Proc. 21st European Workshop on Computational Geometry, 2005, 77-80.

[21] E.Steinitz. *Polyeder und Raumeinteilungen.* Enzyklopaedie der Math. Wiss. III AB 12, Leipzig, 1916.

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