Spatial Embedding of Pseudo-Triangulations

Oswin Aichholzer*
Institute for Software Technology
Graz University of Technology
Graz, Austria
oaich@ist.tugraz.at

Franz Aurenhammer
Hannes Krasser†
Institute for Theoretical Computer Science
Graz University of Technology
Graz, Austria
auren@igi.tugraz.at

Peter Braß‡
Institut für Informatik
Freie Universität Berlin
Berlin, Germany
brass@inf.fu-berlin.de

ABSTRACT

We show that pseudo-triangulations have natural embeddings in three-space. As a consequence, various concepts for triangulations, like flipping to optimality, (constrained) Delaunayhood, and a polytope representation carry over to pseudo-triangulations.

Categories and Subject Descriptors

F.2.2 [**Nonnumerical Algorithms**]: Geometric problems and computations—*computations on discrete structures*

General Terms

Algorithms, Theory

Keywords

Pseudo-triangulations, three-dimensional embedding, constrained regular complexes, polytope representation

1. INTRODUCTION

Triangulations of finite point sets are a versatile tool in geometric data processing, and are well-investigated objects in computational geometry; see e.g., the survey articles [11, 20, 10]. In recent years, an increasing trend concerning a generalization of triangulations, called pseudo-triangulations (or geodesic triangulations) can be observed. A *pseudo-triangulation* is a planar geometric graph whose faces are pseudo-triangles, i.e., polygons bounded by three concave chains. Interest in these structures originates from applications to ray shooting [14, 22] and visibility [31, 32]. Meanwhile, their usefulness as an efficient data structure in many other areas has been discovered, including kinetic collision detection [1, 25, 26], rigidity [39, 36, 23], and guarding [33, 38].

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 [24], on the number of possible pseudo-triangulations [34, 6], and on edge flipping operations in pseudo-triangulations [13, 36, 5] have been obtained recently. A geometric result, in [36] and generalized in [30], 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. As a main theorem, we prove that pseudo-triangulations have natural embeddings in polyhedral surfaces. We discuss various consequences of this theorem. The following is an outline of our results.

Surface theorem We start by revisiting flip operations in pseudo-triangulations, in Section 2. The edge-removing flip, which has been recently introduced (in [5] and independently in [30]), will turn out of particular importance for a proper treatment of pseudo-triangulation embeddings. Section 3 describes unique representations of pseudo-triangulations as polyhedral surfaces. 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 4 gives a surface interpretation of all the admissible flip types from Section 2. Special cases are tetrahedral surface flips, which are well known from flipping in convex hulls; see e.g. [27, 19].

Locally convex functions Section 5 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 per-

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

†Research supported by the FWF [Austrian Fonds zur Förderung der Wissenschaftlichen Forschung]

‡Research supported by DFG under grant Br 1465/5-2

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SoCG'03, June 8–10, 2003, San Diego, California, USA.
Copyright 2003 ACM 1-58113-663-3/03/0006 ...$5.00.

forming surface flips that 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 $O(n^2)$ surface flips suffice. The former result – where the optimum is a lower convex hull, and a regular triangulation in the projection – strengthens a result in [19]. They show that the optimum cannot be reached *within* the class of triangulations, unless insertion of vertices is allowed. 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 Delaunay flips; see e.g. [20].

Constrained regular cell complexes Extending the scenario, by fixing a set of edges which are not to be flipped, we arrive at the concept of constrained regular pseudo-triangulations. These are projections of unique maximal surfaces that are locally convex with respect to a constraining planar and connected graph. In contrast, constrained regular *triangulations* are ambiguous structures, apart from special cases like the constrained Delaunay triangulation [29, 15]. The afore-mentioned algorithmic results carry over to the constrained setting. Moreover, 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. This implies that the set of constrained regular pseudo-triangulations is connected under the flip operations in Section 2 – and generalizes the polytope constructions in [28, 21, 12] (the so-called associahedron, or secondary polytope) for regular triangulations.

Applications We expect various applications of our results, 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 [7] (a generalization of Voronoi diagrams, dual to regular triangulations), weighted alpha shapes [18] and dual complexes [17] (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.

2. BASIC PROPERTIES

Basics and notation For a (simple) polygon P, denote with $vert(P)$ the set of vertices of P. A *corner* of P is a vertex with interior angle less than π. The other vertices of P are called *non-corners*. The *geodesic* between two points $x, y \in P$ is the shortest curve that connects x and y and lies inside P. A *pseudo-triangle* is a polygon with exactly three corners.

Let S be a set of n points in the plane. We will assume – throughout this paper – that no three points in S are collinear. Let $conv(S)$ denote the convex hull of S. A *pseudo-triangulation* of S is a partition of $conv(S)$ into pseudo-triangles (faces) whose vertex set is exactly S. Two faces may have up to two edges in common. In case of such *double-adjacencies*, the union of the two adjacent pseudo-triangles is a pseudo-triangle itself.

Let $\mathcal{PT}$ be some pseudo-triangulation of S. A vertex of $\mathcal{PT}$ is called *pointed* if its incident edges lie in an angle smaller than π. $\mathcal{PT}$ contains exactly $3n - p - 3$ edges and $2n - p - 2$ pseudo-triangles, if 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. The edge rank expresses the excess in edges, compared to a minimum pseudo-triangulation. Its value is maximal if $\mathcal{PT}$ is a triangulation.

Pseudo-flips *Flips* are constant-size combinatorial changes which are commonly used to modify (pseudo)-triangulations. The standard edge flip, also called *Lawson flip* [27], exchanges diagonals in a convex quadrilateral formed by two adjacent faces of a triangulation. For minimum pseudo-triangulations, a geodesics-based generalization of the Lawson flip has been utilized [31, 39], namely the *exchanging flip* illustrated in Figure 2. Also, *vertex-removing* (resp. *inserting*) *flips* have been applied in certain situations, to triangulations [27, 35, 19] and later on to minimum pseudo-triangulations [39, 6]. Recently, *edge-removing flips* have been introduced in [5, 30]. See Figures 1 and 3 for the last two types.

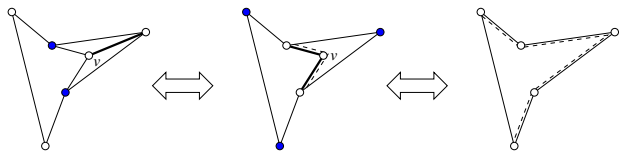

Figure 1: Edge-removing and vertex-removing flips

The edge-removing flip and its inverse, the edge-inserting flip, allow to decrease and increase the edge rank of a given pseudo-triangulation. All the mentioned flip types will have their equivalents for pseudo-triangular surfaces in three-space, described in Section 4.

Augmented polygons We will work in the general setting of pseudo-triangulations for simple polygons with internal points. This covers *pseudo-triangulations for point sets*: choose for the bounding polygon the convex hull of the given point set. As we shall see, certain questions concerning three-dimensional embeddings of pseudo-triangulations become more interesting when being defined in this setting.

Formally, let P be a polygon, and consider an n-point set $S \subset P$ with $vert(P) \subseteq S$. The pair (P, S) is called an *augmented* polygon. A result in [5] states that any two pseudo-triangulations 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. In this transformation, the underlying point set S is left unchanged, as no vertex-removing or vertex-inserting flips are used.

3. SPATIAL EMBEDDING

A *polyhedral surface* is the graph of a continuous and piecewise-linear function whose domain is (a polygonal subset of) the plane. 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 di-

agrams and power diagrams (even their generalizations to higher order), and Delaunay triangulations; see e.g. [8] and references therein. Not every triangulation, however, is regular but it is trivial that triangulations are projective.

We intend to show, in Subsection 3.1, that pseudo-triangulations have natural embeddings in polyhedral surfaces. We characterize the class of projective pseudo-triangulations in Subsection 3.2, and discuss regular pseudo-triangulations and their constrained variants in later sections.

3.1 Surface theorem

We start with defining a *status* for vertices, which will allow us to determine which vertices can be fixed in height when embedding a given pseudo-triangulation in 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 *rigid* if v is a corner in each of its incident pseudo-triangles. The status of v is *pending*, otherwise. For example, every corner of P is a rigid vertex, because no internal angle larger than π is possible. If $\mathcal{PT}$ is a minimum pseudo-triangulation then only the corners of P are rigid and, in particular, all points in $S \setminus vert(P)$ are pending. On the other end of the spectrum, if $\mathcal{PT}$ is a (full) triangulation then all its vertices are rigid.

Note that pending vertices have to be pointed, whereas rigid vertices are pointed if and only if they are corners of P. Observe further that no vertex status is affected by an exchanging flip, whereas an edge-removing flip alters the status of exactly one vertex from rigid to pending. A vertex-removing flip does not alter the status of any remaining vertex.

THEOREM 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 rigid vertex of $\mathcal{PT}$. For each choice of $\boldsymbol{h}$, there exists a unique polyhedral surface $\mathcal{F}$ above 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 rigid vertices of $\mathcal{PT}$, and let $v_{r+1}, \ldots, v_n$ be the pending ones. We set up a system of linear equations for the vertex heights $z_1, \ldots, z_n$ that completely describes the surface $\mathcal{F}$ and has a unique solution.

If v_i is a rigid 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 for the system $A \cdot \boldsymbol{z} = \boldsymbol{b}$, 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 [9].

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 rigid vertices) are present in A', and is retained because heights for rigid 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 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 3.2.

The existence of *some* projection surface is obvious when P contains no internal points, 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 3 in Section 7.

Theorem 1 holds for arbitrary (augmented) polygonal domains, including polygons with holes. We will use this generalization in Subsection 4.3. Non-corners of holes are corners of the domain, and thus are rigid 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 [9]. Moreover, the polytope representation of pseudo-triangulations we are going to describe in Section 7, could alternatively be obtained using Gale transform techniques; cf. [12]. 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 we will use later.

LEMMA 1. *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 (see e.g. [16]) asserts that, for a given planar straight-line graph G, there exists a projection surface (with a planar outer facet) if and only if G allows for a so-called stress, i.e., an assignment of a (possibly negative) tension to each edge of G such that each vertex of G is in equilibrium. For *convex* augmented polygons (P, S), this theorem combines with Theorem 1 when

height zero is assigned to each (necessarily rigid) boundary vertex of P.

COROLLARY 1. *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 non-pointed vertices.*

Interestingly, Corollary 1 also follows from recent results in [30].

3.2 Projective pseudo-triangulations

The surface $\mathcal{F}$ in Theorem 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 pending 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 $S = vert(P)$.

THEOREM 2. *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 pending 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 1, there exist unique surfaces $\mathcal{F}$ for $\mathcal{PT}$ and $\mathcal{F}'$ for $\mathcal{PT}'$. Since only pending 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 rigid 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 pending. 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 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. □

By Theorem 2, 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. This property does not characterize minimum pseudo-triangulations, though. In particular, not every pseudo-triangulation can be made minimum by removing edges. It can be made *minimal* in edge rank, and is termed thereafter in [1, 24, 37].

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 pending, and their elimination leaves a valid pseudo-triangle, P, without changing the status of any vertex of P. By Theorem 2, this implies that $\mathcal{MPT}$ is not projective.

4. FLIPS IN SURFACES

We next provide a surface interpretation for all the admissible flip operations mentioned in Section 2. To this end, consider some projective pseudo-triangulation $\mathcal{PT}$ of an augmented polygon (P, S). (Recall from Subsection 3.2 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 rigid in both structures. By Theorem 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.

In the next two subsections, we study the following *context-free* surface interpretations of exchanging and edge-removing flips. We will use the following notation. ∇_1 and ∇_2 are the two adjacent pseudo-triangles which are involved in such a flip, and R denotes their union. We consider the restriction of the surface flip to R, namely where $\mathcal{PT}$ and $\mathcal{PT}'$ are the two (projective) pseudo-triangulations of R before and after the flip, and where $\mathcal{F}$ and $\mathcal{F}'$ are the corresponding unique surfaces. The operation that takes $\mathcal{F}$ into $\mathcal{F}'$ is termed a context-free surface flip. We come back to its context-sensitive version in Subsection 4.3.

4.1 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 rigid 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 pending 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.

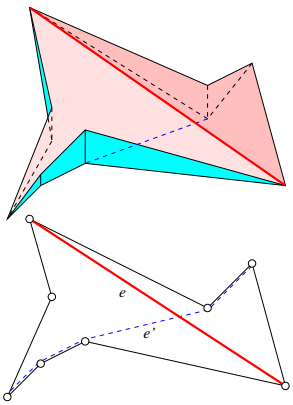

Figure 2: Convexifying flip

To view the change geometrically, see Figure 2. 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 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. [27].

4.2 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 rigid 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 rigid in $\mathcal{PT}$ to pending 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 pending 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 3. $\mathcal{F}'$ 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.

4.3 Context-sensitive view

The context-free flips described in Subsections 4.1 to 4.2 leave the heights of the rigid vertices of $\mathcal{PT}'$ unchanged. The pending vertices of $\mathcal{PT}'$ are lowered. Rigid vertices may, however, be pending 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.

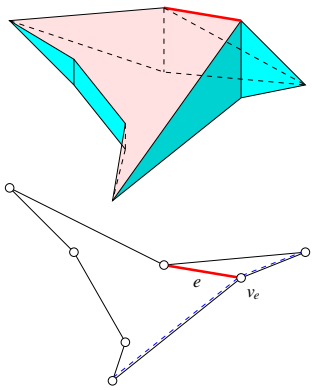

Figure 3: Planarizing flip

LEMMA 2. *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. Let (P, S) and $\mathcal{PT}$, respectively, be the augmented polygon and the pseudo-triangulation in question. Consider the polygon $R \subseteq P$ the context-free version of φ is restricted to, and denote this flip by φ_R. Furthermore, let $\mathcal{F}$ be the surface for $\mathcal{PT}$ before flipping.

Introduce new edges to $\mathcal{PT}$ such that a pseudo-triangulation, $\mathcal{PT}_1$, is obtained in which all corners of R are rigid vertices. Assign heights as determined by $\mathcal{F}$ to the rigid 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 pending in $\mathcal{PT}_2$) but leaves the heights of R's corners fixed. To see that no other pending 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 rigid 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 1. 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. □

Figure 4 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.

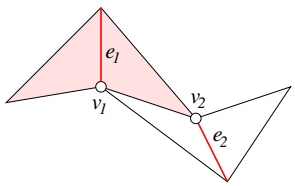

Figure 4: Oscillating flips

4.4 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 pending 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 2. 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. [27].

Remarks Each (non-tetrahedral) surface flip requires a system of linear equations to be solved; see Subsection 3.1. The corresponding matrix has a special structure, though. The first $n-p$ rows are unit rows, if p of the n vertices in the current pseudo-triangulation are pending. So the system can be solved in $O(n+p^3)$ time, which covers the geometrical part of the surface flip. The combinatorial part obviously can be done within this complexity.

5. 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 interior 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, in general.

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: 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, the surface $\mathcal{F}^*$ can be generated by $O(n^2)$ surface flips of the kind above. The bound $O(n^2)$ also holds when (P, S) is an arbitrary polygon without internal vertices.

5.1 Flipping to optimality

Consider some full triangulation $\mathcal{T}$ of the augmented polygon (P, S). All the n vertices of $\mathcal{T}$ are rigid 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 $\mathcal{F}(\mathcal{T}, \boldsymbol{h})$.

THEOREM 3. *(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 2, 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 1, these surfaces are pairwise different, so are the corresponding pseudo-triangulations. An exponential upper bound on the sequence length results. 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 rigid 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.

Results on flip distances in [5] imply 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 (as it is required in Theorem 3). We summarize:

THEOREM 4. *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.

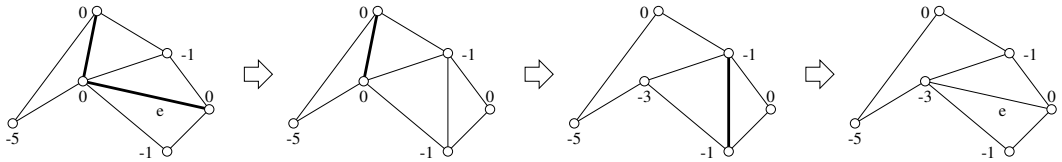

Figure 5: Edge e reappears

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

A principal problem, which complicates the construction of improving sequences of polynomial length (as in [20]), is that edges flipped away may reappear. Figure 5 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. For less general domains (P,S), the reappearance of edges can always be avoided; see Subsection 5.2.

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 6 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.

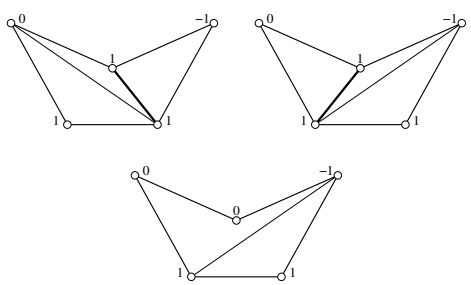

Figure 6: Local and global optima

5.2 Special domains

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. For this situation, reappearence of edges does not take place, if improving surface flips are applied to the initial triangular surface in a carefully chosen order. In particular, the following holds. (The proofs of the following two lemmas are omitted.)

LEMMA 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 $n = |S|$.*

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 which is taken from [19]. 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.

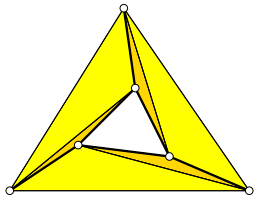

Figure 7: Local optimum for convex P

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 $vert(P) = S$.

LEMMA 4. *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 $\boldsymbol{h}$.*

6. (CONSTRAINED) REGULAR CELL COMPLEXES

6.1 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 [7] (a generalization of Voronoi diagrams) and have several applications [18, 17]. Figure 7 shows a well-known example of a non-regular triangulation.

We discuss the implications of Section 5 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 3 implies:

THEOREM 5. *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.*

COROLLARY 2. *The (lower part of the) convex hull of a three-dimensional n-point set M can be constructed by starting with 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. [20]

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 5 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. [20]. This implies that Theorem 5 is optimal in the worst case. Theorem 5 is a stronger version of a result in [19], 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 [36]. 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 pointedness of v forces every surface to be non-convex in the neighborhood of v.

6.2 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. [29, 15, 4]. We intend to show that our framework allows for a definition of constrained *regular* (pseudo-)triangulations.

To exploit our concepts in full 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. 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 $\bigcup_i S_i = S$.

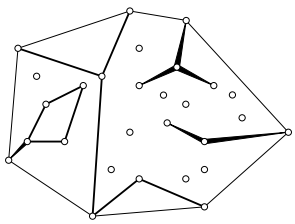

Figure 8: Cutting $conv(S)$ along G

Theorems 3 and 4 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 6. *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 4 applies.*

Remarks If G partitions $conv(S)$ into augmented *convex* polygons, then no pending 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 [29] of S and G is obtained.

An attempt to define constrained regular complexes as full triangulations results in ambiguity. Reflex edges not belonging to G arise, whose 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 6, where G delineates a non-convex polygon.

7. 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 6.2. 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 [28] (the associahedron) and in [21] (the secondary polytope, see also [12]), 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 6.2 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 rigid vertex of the pseudo-triangulation $\mathcal{PT}_i$. So, by the surface theorem (Theorem 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 from Section 6.2 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 3) thus implies $f_{\mathcal{PT},\boldsymbol{h}} \leq f_{\mathcal{PT}',\boldsymbol{h}}$ for all $\mathcal{PT}' \in \Pi(G)$. Integrating both sides yields

$$\int_{x \in conv(S)} f_{\mathcal{PT},\boldsymbol{h}}(x)\, dx \leq \int_{x \in conv(S)} f_{\mathcal{PT}',\boldsymbol{h}}(x)\, dx. \quad (1)$$

Consider the lefthand-side integral in (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 1, 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 (1), we obtain

$$\boldsymbol{q}(\mathcal{PT}) \cdot \boldsymbol{h} \leq \boldsymbol{q}(\mathcal{PT}') \cdot \boldsymbol{h}, \text{ for all } \mathcal{PT}' \in \Pi(G). \quad (2)$$

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 (2) implies that, for each $\boldsymbol{h} \in \mathsf{V}(G, \mathcal{PT})$, the point 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 6. $\mathsf{V}(G, \mathcal{PT})$ 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 (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 7. *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 the polytope $\mathcal{Q}(G)$ 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, by results in [5]. We obtain:

COROLLARY 3. *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 $O(n)$.*

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 [36]. In view of our extended repertoire of flips, this question arises naturally. We recently learned that [30] gave an affirmative answer for the point set case, by generalizing the construction in [36]. Recent results on flip distances [5] imply that the diameters of the polytopes in [36] and [30] are $O(n \log^2 n)$ and $O(n \log n)$, respectively. What is the relation of these polytopes to the polytopes in Theorem 7 and Corollary 3?

8. CONCLUSION

We have demonstrated that pseudo-triangulations possess strong geometric properties. Various algorithmic consequences have been discussed, which we believe to be of value in practical applications. 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 4, opens a way to define pseudo-complexes in higher dimensions, and flips therein. Also, the techniques used to prove the surface theorem (Theorem 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.

9. ACKNOWLEDGMENTS

The authors gratefully acknowledge discussions on the presented topic with Günter Rote and Francisco Santos.

10. 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. Adapting pseudo-triangulations with a near-linear number of edge flips. Manuscript, 2002.

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

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

[8] 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.

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

[10] 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.

[11] 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.

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

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

[14] 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.

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

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

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

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

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

[20] 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.

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

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

[23] R.Haas, F.Santos, B.Servatius, D.Souvaine, I.Streinu, W.Whiteley. Planar minimally rigid graphs have pseudo-triangular embeddings. Manuscript, 2002.

[24] 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, 2002. To appear.

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

[26] 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.

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

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

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

[30] D.Orden, F.Santos. The polyhedron of non-crossing graphs on a planar point set. Manuscript, 2002.

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

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

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

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

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

[36] G.Rote, F.Santos, I.Streinu. Expansive motions and the polytope of pointed pseudo-triangulations. Manuscript, FU-Berlin, 2001.

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

[38] B.Speckmann, C.D.Toth. Allocating vertex π-guards in simple polygons via pseudo-triangulations. Abstr. 18th Europ. Workshop Computational Geometry, 2002, 12-15.

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