

Pseudo-Simplicial Complexes from Maximal Locally Convex Functions*

Franz Aurenhammer and Hannes Krasser

Institute for Theoretical Computer Science,
Graz University of Technology,
Graz, Austria
{auren,hkrasser}@igi.tugraz.at

Abstract. We introduce and discuss pseudo-simplicial complexes in \mathbb{R}^d as generalizations of pseudo-triangulations in \mathbb{R}^2 . Our approach is based on the concept of maximal locally convex functions on polytopal domains.

1. Introduction

A pseudo-triangulation is a cell complex in the plane where each cell is a pseudo-triangle, i.e., a simple polygon with convex angles at exactly three vertices (the so-called corners). Figure 1 illustrates an example. Being an interesting and flexible generalization of triangulations, pseudo-triangulations have found their place in computational geometry. The scope of their applications is broad, and they enjoy rich combinatorial and geometric properties; see, e.g., [22], [11], [17], [10], [1], [16], and references therein. Unlike triangulations, pseudo-triangulations have so far eluded a meaningful generalization to higher dimensions. The definition of pseudo-simplices remained unclear, possibly because of the lack of an adequate definition of corners of a polytope.

In this paper we define pseudo-simplices and pseudo-simplicial complexes in d -space in a way consistent with pseudo-triangulations in the plane. Flip operations in pseudo-complexes are specified as combinations of flips in pseudo-triangulations [17], [1], [16] and of bistellar flips in simplicial complexes [12], [7], [6]. A certain class of pseudo-complexes is shown to be connected under such flips. The flip distance for this class is $O(n^{d+2})$, where n is the size of the underlying vertex set. Moreover, the class admits a

* This research was supported by the FWF Research Project P14755-MAT and the FWF Joint Research Project “Industrial Geometry” S9205-N12.




Fig. 1. Pseudo-triangulation of a simple polygon.

representation as a convex polytope in n -space, that generalizes the polytope constructions in [9], [3], and [1]. Our results are based on the concept of maximal locally convex functions on polyhedral domains, that allows us to unify several well-known structures, namely, pseudo-triangulations, constrained Delaunay triangulations [5], [21], and regular simplicial complexes [14], [7]. The concept is of interest in its own right, as it leads to a generalized notion of lower convex hulls of finite point sets.

Once available, pseudo-complexes give rise to various questions. Some of them find a direct answer, as, for instance, a consistent generalization of minimum pseudo-triangulations [22], [17] to pseudo-complexes in d -space. More elaborate is a question of interest in motion planning applications [11], of whether visible space between polyhedral objects can be tiled and maintained with pseudo-simplices and flips, respectively. A related important problem is whether any two simplicial complexes in d -space can be transformed into each other by flips. For Lawson flips [12], the answer is negative for $d \geq 5$, and unknown for dimensions 3 and 4; see [18] and [19]. We explore the impact of our results to these and other questions in a separate paper. The main intention of the present work is to lay theoretical foundations for a treatment of pseudo-complexes.

We give a short outline of the paper. Section 2 provides notation on polytopes and cell complexes in d -space. Beyond standard definitions, the concepts of terminal vertex and corner of a d -polytope are introduced. This allows us to define pseudo-simplices as polytopes having exactly $d + 1$ corners. Sections 3 and 4 are intended for a study of locally convex functions, which are maximal with respect to predefined value bounds at the vertices of the d -polytope that serves as their domain of definition. We show that any such maximal locally convex function f^* is piecewise linear, but possibly discontinuous for $d \geq 3$ at its domain boundary. A pointwise characterization of f^* is given, that identifies this function as a natural generalization of the lower convex hull of a finite point set in $(d + 1)$ -space. This characterization also allows us to describe the discontinuity behavior of f^* . In Section 5 we show that f^* generates cell complexes in its domain such that each d -dimensional cell is indeed a pseudo-simplex. This generalizes results in [1], that first observed that pseudo-triangulations are related to locally convex functions,

as well as in [2], that extended the situation to arbitrary polygonal domains. Section 6 demonstrates that pseudo-complexes are able to cope with flipping operations, namely, those resulting as minimal changes in f^* when value bounds at domain vertices vary. All known types of flips for simplicial complexes [12], [7], [6] and for pseudo-triangulations [17], [1], [16] can be obtained as special cases. We illustrate the different flips that arise in 3-space and discuss the occasional strange behavior of the resulting pseudo-complexes. The occurrence of vertices which are not vertices of the domain of f^* , and the existence of cells with tunnels, are two examples. Section 7 shows that the “surface theorem” for pseudo-triangulations in [1] partially generalizes to pseudo-complexes for $d \geq 3$. As a consequence, the class of pseudo-complexes that f^* generates in a given domain enjoys a representation as a high-dimensional convex polytope, a fact that generalizes results in [9], [3], and [1]. We also introduce two subclasses of pseudo-complexes that constitute extensions to $d \geq 3$ of the class of triangulations, and the class of minimum pseudo-triangulations [22], [17], of a simple polygon, respectively.

We remark that pseudo-complexes, as defined in this paper, do not contain internal vertices. That is, all their vertices lie on the boundary of the underlying domain. An extension that remedies this fact is straightforward, by predefining value bounds for f^* at certain points interior to the domain. Defining pseudo-complexes in this way may be of practical relevance but does not lead to new theoretical insights. We therefore refrained from doing so in the present paper, apart from some remarks in Sections 5 and 7.

2. Polytopes and Corners

We give some notation concerning polytopes in d -space \mathbb{R}^d . In particular, the proper definition of corners of a polytope is crucial for the subsequent investigations. Abbreviations in sans serif (bd, int, relint, conv, aff, dim) denote standard set-theoretic notions in convex geometry [4].

Let P be a bounded and closed subset of \mathbb{R}^d . Then P is called a d -polytope if P is an interior-connected d -manifold, with piecewise linear boundary $\text{bd } P$ that is structured as a $(d - 1)$ -dimensional cell complex. The components of $\text{bd } P$ of dimension j , $0 \leq j \leq d - 1$, are called the j -faces of P . (The dimension $\text{dim } X$ of a set $X \subset \mathbb{R}^d$ is the dimension of its affine hull $\text{aff } X$.) Faces of dimensions $d - 1$, 1, and 0, respectively, are also called *facets*, *edges*, and *vertices* of P . We denote with $\text{vert } P$ the set of vertices of P . A d -polytope P is called *boundary-connected* if $\text{bd } P$ is a connected set. P is called *simple* if P is homeomorphic to a closed ball in \mathbb{R}^d .

A face F of a d -polytope P is *concave* if there exists some line segment $L \subset P$ such that $\text{relint } L \cap \text{relint } F$ is a single point. Otherwise, F is *nonconcave*. The facets of P are nonconcave faces, whereas faces of lower dimensions may be of either type. In Fig. 2, the edges xz and yz are concave, whereas the edge xy is nonconcave.

A *terminal* of a d -polytope P is a point $x \in P$ such that no line segment $L \subset P$ has $x \in \text{relint } L$. All terminals of P belong to $\text{vert } P$. In fact, they are just the nonconcave vertices of P . A terminal v of P is called a *corner* of P if there exists a hyperplane through v that has all edges of P incident to v on a fixed side. For $d = 2$, the vertices of P automatically fulfill this condition, because we required P to be interior-connected, and thus P is a simple polygon. Therefore “terminal” and “corner” are identical notions




Fig. 2. Simple 3-polytope with different edge and vertex types.

in this case. Vertices of P which are not corners of P are called *noncorners* of P . In the 3-polytope in Fig. 2, vertex x is a corner, vertex y is a terminal but a noncorner, and vertex z is not a terminal and therefore a noncorner.

If v is a corner of P then v is also a corner of all faces of P incident to v . Moreover, all vertices of the convex hull $\text{conv } P$ are corners of P . So every d -polytope P has at least $d + 1$ corners, and if P is convex then all its vertices are corners. The converse of the last statement is true only for $d = 2$; Fig. 3 gives a counterexample for $d = 3$.

A d -polytope P is termed a *pseudo-simplex* if P is boundary-connected and has exactly $d + 1$ corners. Clearly, every simplex is a pseudo-simplex. For $d = 2$ and $d = 3$, respectively, pseudo-simplices are also called *pseudo-triangles* and *pseudo-tetrahedra*. Our definition of a pseudo-triangle is equivalent to the classical definition, see, e.g., [17] and [1], which requires exactly three polygon vertices with an internal angle smaller than π . A *pseudo-complex* is a cell complex in d -space whose cells (d -faces) are all pseudo-simplices. For $d = 2$, pseudo-complexes are pseudo-triangulations. See Fig. 1 where a pseudo-triangulation with seven cells (pseudo-triangles) is shown.

3. Local Convexity

The theory of maximal locally convex functions is the key to a derivation of pseudo-complexes. For $d = 2$, the relationship between these two concepts has been observed and exploited in [1] and [2]. Its generalization to d -space has its peculiarities, and requires careful treatment of locally convex functions.

We consider real-valued functions on \mathbb{R}^d whose domains are simple d -polytopes. For $x, y \in \mathbb{R}^d$, a function f on the line segment xy is called *convex* if $f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y)$, for $0 \leq \lambda \leq 1$. f is called *linear* on xy if equality holds. Let a simple d -polytope D be given. A function f on D is called *locally convex* if f is convex on each line segment $L \subset D$. A locally convex function is convex if D is convex.

Let \mathbf{h} be a real-valued vector that assigns a $(d + 1)$ st coordinate h_i (called *height*) to each vertex $v_i \in \text{vert } D$. In the discussion below, we assume that the domain D and the height vector \mathbf{h} are fixed. Our interest is in the maximal locally convex function f^* on D which fulfills $f^*(v_i) \leq h_i$ for each $v_i \in \text{vert } D$. The function f^* is unique, because f^* is the pointwise maximum of all locally convex functions which satisfy the constraints in \mathbf{h} .

Moreover, f^* is continuous on $\text{int } D$ by its local convexity. f^* need not be continuous on $\text{bd } D$, however; see Section 4.

Figure 1 illustrates an example for $d = 2$. The domain D is a simple polygon whose vertices v_i are labeled with upper height bounds h_i for f^* . The function f^* is linear on each of the shown subpolygons of D . By its local convexity, f^* does not assume the maximal allowed value h_i at certain noncorners of D , namely, at the four vertices with labels 5 or larger.

A vertex $v_i \in \text{vert } D$ is termed *maximal* if $f^*(v_i) = h_i$. Clearly, if v_i is not maximal then we have $f^*(v_i) < h_i$. We state several basic properties of maximal locally convex functions.

Lemma 1 (Linearity Lemma). *For each point $z \in D$ that is not a maximal vertex of D , there exists some line segment $L \subset D$ with $z \in \text{relint } L$ and such that f^* is linear on L .*

Proof. Assume that such a line segment L does not exist for a point $z \in D$. Consider an open ball B in \mathbb{R}^d that is centered at z . For all line segments $xy \subset B \cap D$ with $z \in \text{relint } xy$, by assumption we have $f^*(z) < \lambda f^*(x) + (1-\lambda)f^*(y)$, for $z = \lambda x + (1-\lambda)y$. We construct a function f on D with the following properties: f is convex on all the line segments xy above and linear on at least one among those, and f coincides with f^* on the complement of B in D . By construction, f is locally convex on D . Moreover, if B is chosen small enough then f satisfies the constraints in \mathbf{h} , unless z is a maximal vertex. However, $f(z) > f^*(z)$ holds, contradicting the maximality of f^* . We conclude that z is a maximal vertex. \square

Corollary 1. *All terminals of D are maximal vertices.*

Proof. Let v be a terminal of D . By definition of a terminal, there exists no line segment $L \subset D$ with $v \in \text{relint } L$. Thus, by Lemma 1, v is a maximal vertex. \square

It is well known that any convex function f induces a face-to-face cell complex in its domain [4]. Namely, a j -face of f is a maximal connected subset G of f 's domain such that, for each $x \in \text{relint } G$, the space of directions where f has a partial derivative at x has dimension j . As the maximal locally convex function f^* is convex on each convex subset of its domain D , f^* induces such a cell complex in D as well. By convention, the intersection of a face of f^* with a face of $\text{bd } D$ will be considered a face of f^* . So, for example, all vertices of D are included in the set of 0-faces of f^* .

Lemma 2. *f^* is linear in the (relative) interior of each of its j -faces, for $1 \leq j \leq d$.*

Proof. We use induction on the dimension j . Let F be a 1-face of f^* . By Lemma 1, for each point $x \in \text{relint } F$ there is some line segment $L_x \subset D$ with $x \in \text{relint } L_x$ and f^* linear on L_x . Since F is a 1-face of f^* , we cannot have $L_x \cap F = x$. This implies that F is a line segment itself, and that f^* is linear on $\text{relint } F$.

Consider a j -face G of f^* , for $2 \leq j \leq d$. Assume inductively that f^* is linear on the relative interior of each k -face, $k < j$, incident to G . We show that f^* is linear on




Fig. 3. Splitting the Schönhardt polytope with f^* .

relint G . Let B be some j -ball with $B \subset \text{relint } G$. Then f^* is piecewise linear on B , because f^* could be increased on B without change on the relative boundary of G , otherwise, contradicting the maximality of f^* . (Note that this is true even if f^* is discontinuous.) This implies that f^* is piecewise linear on relint G . However, then f^* has to be linear on relint G , by the definition of a j -face of f^* . \square

By Lemma 2, f^* is a piecewise linear function on D . The continuity of f^* on int D implies that the faces of f^* have piecewise linear boundaries, and therefore are j -polytopes, for $0 \leq j \leq d$. The d -faces of f^* are also called *cells*.

Lemma 3. *The cells of f^* are boundary-connected d -polytopes.*

Proof. Let C be a cell of f^* . For every line segment $L \subset D$, f^* is linear on relint $(L \cap C)$, by Lemma 2. Moreover, f^* is convex on the entire segment L , by its local convexity on D . This implies that $L \cap C$ is connected. The assumption that D is a simple d -polytope now yields that C is boundary-connected. \square

In Fig. 1 we see how f^* partitions a simple polygon D into polygonal cells. (Here and in Fig. 3, numbers denote vertex heights.) Observe that each polygonal cell has exactly three corners. Figure 3 illustrates a three-dimensional cell complex induced by f^* when D is the Schönhardt polytope [20]. All six vertices of D are corners and therefore are maximal. The complex consists of three nontetrahedral cells, each having exactly four corners. Note the occurrence of 0-faces of f^* in the relative interior of concave edges of D .

4. Local Characterization of f^*

Our next aim is to give a local characterization of f^* by means of visibility. Intuitively speaking, f^* locally behaves like a convex hull of a discrete point set in \mathbb{R}^{d+1} . We shall see that maximal locally convex functions constitute a natural generalization of lower convex hulls.

4.1. Hull Theorem

We start with a technical lemma. Recall that f^* is defined by its domain D and its height vector \mathbf{h} . For a given simplex Δ with $\text{vert } \Delta \subset D$, let ℓ^Δ be the linear function on Δ with $\ell^\Delta(x) = f^*(x)$ for all $x \in \text{vert } \Delta$. For two points $x, y \in D$, we say that x sees y if the line segment xy is contained in D .

Lemma 4. *Let $z \in \text{int } D$ be a point that sees $t \leq d + 1$ affinely independent points $x_1, \dots, x_t \in D$. Let $\Delta = \text{conv } \{x_1, \dots, x_t\}$. If $z \in \Delta$ then $f^*(z) \leq \ell^\Delta(z)$ holds.*

Proof. If $z = x_i$ for some i , then $f^*(z) = \ell^\Delta(z)$ holds by definition of ℓ^Δ . The interesting case is $z \notin \{x_1, \dots, x_t\}$. In this case there exists a (sufficiently small) $(t - 1)$ -simplex $\Delta_\varepsilon \subset D$ with vertices on the t line segments zx_1, \dots, zx_t and with $z \in \Delta_\varepsilon \setminus \text{vert } \Delta_\varepsilon$. Moreover, there is a face F of Δ_ε and an index i such that the line L containing zx_i intersects $\text{relint } F$ in a point y . ($y = z$ is possible if $z \in \text{bd } \Delta_\varepsilon$.) Now assume $f^*(z) > \ell^\Delta(z)$. As f^* is convex on each connected part of $L \cap D$, we get $f^*(y) > \ell^F(y)$. Repeating this argument with Δ_ε replaced by F yields the existence of an edge e of Δ_ε and a point $x \in \text{relint } e$ with $f^*(x) > \ell^e(x)$. Thus f^* is not convex on e , a contradiction. \square

The following theorem asserts that f^* can be defined pointwise, by means of its values at visible 0-faces. Some notation is needed for a precise formulation of this fact. For any point $x \in D$, interpret the pair $(x, f^*(x))$ as a point x^* in \mathbb{R}^{d+1} . For a point set $A \subset D$, let $A^* = \{x^* \mid x \in A\}$. For a point set $B \subset \mathbb{R}^{d+1}$, define low_B as the (convex) function whose graph is the lower convex hull of B , i.e., the part of $\text{bd conv } B$ visible from $-\infty$ on the $(d + 1)$ st coordinate axis. Finally, denote with $V(x)$ the set of all 0-faces of f^* that a point $x \in D$ can see.

Theorem 1 (Hull Theorem). *For every point $x \in \text{int } D$ we have $f^*(x) = \text{low}_{V(x)^*}(x)$.*

Proof. Let $x \in \text{int } D$. Consider the set S of all d -simplices that contain x and that are spanned by $d + 1$ 0-faces in $V(x)$. Any simplex Δ in S that yields the smallest function value $\ell^\Delta(x)$ defines $\text{low}_{V(x)^*}(x)$. So, by Lemma 4, we have $f^*(x) \leq \ell^\Delta(x) = \text{low}_{V(x)^*}(x)$. On the other hand, there is a simplex Δ_C in S , spanned by 0-faces of a cell C of f^* that contains x , and defining $f^*(x) = \ell^{\Delta_C}(x)$, by the linearity of f^* on $\text{int } C$ (Lemma 2) and the maximality of f^* . This implies $f^*(x) \geq \text{low}_{V(x)^*}(x)$. The theorem follows. \square

4.2. Continuity Behavior of f^*

We next study the continuity behavior of f^* . Somewhat unexpectedly, f^* may be discontinuous at certain subsets of $\text{bd } D$. The polytope D in Fig. 2 serves as an example. As heights for D we choose 1 for vertex y and values smaller than 0 for the remaining vertices. Then, by Lemma 4, $f^*(p) < 0$ for all $p \in \text{int } D$. On the other hand, we have $f^*(y) = 1$, because y is a terminal and therefore is a maximal vertex by Corollary 1.

Thus f^* is discontinuous at y (as well as in the relative interior of the three nonconcave edges incident to y).

The following lemma indicates that f^* is “well behaved” at the corners of its cells. We call a 0-face v of f^* *complete* if v is a corner of all its incident cells. For instance, each corner of D is a complete 0-face of f^* . For the case $d = 2$, the notion of completeness of a vertex in a cell complex has been introduced in [1], and has been used extensively in [2].

Lemma 5.

- (a) For each cell C of f^* we have $(\text{int } C)^* \subset \text{aff } Z^*$, where Z denotes the set of corners of C .
- (b) A 0-face where f^* is discontinuous cannot be a corner of any cell.
- (c) Complete 0-faces of f^* are characterized by being vertices of D that are maximal and where f^* is continuous.

Proof. Let z be a corner of the cell C . We claim (and prove below) that $f^*(z)$ is the continuous extension of f^* from $\text{int } C$ to z . By Lemma 2, this claim implies assertion (a). Further, by the continuity of f^* on $\text{int } D$, the claim yields that f^* is continuous at each 0-face that is a corner of at least one cell. This shows assertion (b). In particular, f^* is continuous at its complete 0-faces. However, each complete 0-face x of f^* is a maximal vertex of D , because there exists no line segment for x as in Lemma 1, by definition of a corner of a cell. Conversely, if a 0-face y of f^* is not complete and thus is a noncorner of some cell C , then either f^* is discontinuous at y , or we have $y^* \in \text{aff}(\text{int } C)^*$ by assertion (a). As y is not a corner of $\text{conv } C$ (y would be a corner of C , otherwise), y is not a maximal vertex of D in the latter case. This proves assertion (c).

To prove the claim, let $N(z)$ be a (small) neighborhood of z in $\text{int } C$. Consider some point $x \in N(z)$. By Theorem 1, we have $f^*(x) = \text{low}_{V(x)^*}(x)$, and this value is determined by a d -simplex Δ that contains x and that is spanned by 0-faces of C . Since z is a corner of C , there exists a hyperplane E through z such that $N(z)$ lies on a fixed side of E . Therefore $z \in \text{vert } \Delta$ holds, and the claim follows. \square

Corollary 2. No 0-face of f^* lies in $\text{int } D$, or in $\text{relint } F$ for any nonconcave j -face F of D , $j \geq 1$.

Proof. Let x be a 0-face of f^* . Assume $x \in \text{int } D$, or $x \in \text{relint } F$ for some nonconcave j -face F of D with $j \geq 1$. In both cases there exists a d -ball B with center x and such that $B \cap D$ is a convex set. In the latter case, x lies on $\text{bd}(B \cap D)$, and therefore x has to be a corner of all its incident cells. In the former case, $x \in \text{int}(B \cap D)$ holds, and we use the fact that f^* is a convex function on $B \cap D$ to see that x has the same property, namely, being a complete vertex of f^* . So x is a maximal vertex of D , by Lemma 5(c). However, this contradicts the choice of x . \square

As a consequence of Corollary 2, f^* is linear on each nonconcave edge of D . In particular, for $d = 2$ where D is a simple polygon, f^* is linear on all the edges of D , because

edges are facets in this case, and facets are nonconcave faces. Also, no noncorner x of D is a terminal for $d = 2$. This implies the following property.

Corollary 3. *For $d = 2$, f^* is continuous on the entire domain D .*

Proof. Let x be a vertex of D . Then x is either a complete vertex of f^* (and f^* is continuous at x by Lemma 5(c)), or x is not complete and a noncorner of D . In the latter case, as x cannot be a terminal of D , there exist line segments $L \subset D$ with $x \in \text{relint } L$ and where f^* is linear. Consequently, f^* is continuous at x in that case, too. The assertion follows. \square

The theorem below reveals that, for convex domains, f^* is just another means for describing lower convex hulls.

Theorem 2. *If the domain D is convex then, for every height vector \mathbf{h} , the graph of f^* is exactly the lower convex hull of the point set $\text{vert } D$ lifted by \mathbf{h} .*

Proof. Assume D is convex. Then all faces of D are nonconcave, so each 0-face of f^* is a vertex of D by Corollary 2. However, each vertex v of D is a corner of D . Therefore, v is a complete vertex of f^* , and v is maximal by Lemma 5(c). Finally, every point $x \in D$ sees all vertices of D . The statement now follows from Theorem 1. \square

5. Pseudo-Complexes

For a given simple d -polytope D and a height vector \mathbf{h} for $\text{vert } D$, let $\mathcal{PC}(D, \mathbf{h})$ denote the polytopal cell complex induced by f^* in D . In this section and in Section 6 we derive a list of structural properties of $\mathcal{PC}(D, \mathbf{h})$, which culminates in the finding that $\mathcal{PC}(D, \mathbf{h})$ is a d -dimensional pseudo-complex that admits bistellar flipping operations.

From Sections 3 and 4 we know that each vertex x of $\mathcal{PC}(D, \mathbf{h})$ either belongs to $\text{vert } D$ or x is the intersection of a j -face of $\mathcal{PC}(D, \mathbf{h})$ with a concave $(d - j)$ -face of D . Accordingly, x will be termed a *primary* or a *secondary* vertex. We extend this terminology to the cells of $\mathcal{PC}(D, \mathbf{h})$ by distinguishing whether or not all corners of a cell are primary vertices. (Noncorners do not influence the status of a cell.) By Theorem 2, secondary vertices and thus secondary cells do not arise if the domain D is convex. The same is true for nonconvex domains if $d = 2$, by Corollary 2. Moreover, as we shall see below, secondary cells are prevented by certain choices of the height vector \mathbf{h} .

To ease the exposition, we restrict attention to nondegenerate height vectors as follows. Let ε be a real number, and let \mathbf{h}_ε be a vector obtained from \mathbf{h} by perturbing each coordinate of \mathbf{h} by at most ε . We call \mathbf{h} *generic* (for D) if an $\varepsilon > 0$ exists such that $\mathcal{PC}(D, \mathbf{h}_\varepsilon) = \mathcal{PC}(D, \mathbf{h})$ holds for all possible vectors \mathbf{h}_ε . This property of \mathbf{h} will be implicitly assumed henceforth.

We start by studying the structure of $\mathcal{PC}(D, \mathbf{h})$ for a special class of height vectors. Denote by H the point set in \mathbb{R}^{d+1} that results from lifting $\text{vert } D$ by \mathbf{h} . We

call \mathbf{h} *convex* (for D) if $\text{low}_H(v_i) = h_i$ holds for each $v_i \in \text{vert } D$. In particular, \mathbf{h} is called *parabolic* if $h_i = |v_i|^2$ holds for each i . Observe that, for arbitrary \mathbf{h} , we have $\text{low}_H(x) \leq f^*(x)$ for all $x \in D$, because low_H is a locally convex function that satisfies \mathbf{h} , and f^* is maximal for these constraints. If \mathbf{h} is convex then $\mathcal{PC}(D, \mathbf{h})$ shows several nice properties.

Theorem 3. *Let \mathbf{h} be a convex height vector. Then all cells of $\mathcal{PC}(D, \mathbf{h})$ are primary cells, and are pseudo-simplices. All primary vertices of $\mathcal{PC}(D, \mathbf{h})$ are complete. Moreover, f^* is continuous on the entire domain D .*

Proof. As \mathbf{h} is convex, $\text{low}_H(v_i) = h_i$ holds for each $v_i \in \text{vert } D$. By $\text{low}_H(v_i) \leq f^*(v_i)$ this implies $f^*(v_i) = h_i$. That is, all primary vertices of $\mathcal{PC}(D, \mathbf{h})$ are maximal. Moreover, by the assumption that \mathbf{h} is generic, we get that the set $(\text{vert } D)^*$ is in strictly convex position in \mathbb{R}^{d+1} . Therefore, for each cell C of $\mathcal{PC}(D, \mathbf{h})$, f^* on $\text{int } C$ depends on exactly $d + 1$ points in $(\text{vert } D)^*$. On the other hand, f^* on $\text{int } C$ is determined by the values of f^* at the corners of C , by Lemma 5(a). We conclude that C has exactly $d + 1$ corners which all belong to $\text{vert } D$. That is, C is a primary cell and a pseudo-simplex.

As $(\text{vert } D)^*$ is in strictly convex position, no point in $(\text{vert } D)^*$ is redundant for f^* on $\text{int } D$. So, by Lemma 5(a), each vertex v of D is a corner in some cell and, by Lemma 5(b), f^* is continuous at v . This implies that f^* is continuous on D . By Lemma 5(c) and the fact that primary vertices are maximal if \mathbf{h} is convex, all primary vertices are complete. \square

Remarks. By Theorem 3, locally convex functions generate pseudo-complexes if the height vector \mathbf{h} is convex. In fact, this is the case for arbitrary \mathbf{h} , as we will see in Section 6. If D (and with it \mathbf{h}) is convex then all cells are simplices, and $\mathcal{PC}(D, \mathbf{h})$ is a regular simplicial complex [14], [7] in D , by Theorem 2. If, in addition, \mathbf{h} is parabolic then the well-known Delaunay simplicial complex [8] for D is obtained. However, when \mathbf{h} is convex but D is not, then $\mathcal{PC}(D, \mathbf{h})$ is not simply the restriction to D of the cell complex defined by low_H (which is a regular simplicial complex in $\text{conv } D$). This is not even true for $d = 2$, where no secondary vertices arise.

The case $d = 2$ has been treated in [1]. In \mathbb{R}^2 , $\mathcal{PC}(D, \mathbf{h})$ is a constrained regular pseudo-triangulation of the simple polygon D , and a triangulation of D if \mathbf{h} is convex. (As all the vertices of $\mathcal{PC}(D, \mathbf{h})$ are primary for $d = 2$, they are all complete provided \mathbf{h} is convex, according to Theorem 3.) In particular, if \mathbf{h} is parabolic, then the constrained Delaunay triangulation [5] of D is obtained, which is the Delaunay triangulation [8] provided D is a convex polygon.

We point out that our discussion can be extended to the case where f^* is required to fulfill predefined height restrictions at points $v_i \in \text{int } D$, rather than at $\text{vert } D$ alone. (This has been done for $d = 2$ in [1].) It can be shown that such a point v_i is either a maximal vertex, or v_i is not a vertex of f^* at all; see the proof of Corollary 2. In the former case, $\mathcal{PC}(D, \mathbf{h})$ is simplicial in the neighborhood of v_i . All the results of the present paper extend easily. We confine ourselves to vertex-empty domains because this is sufficient for the intentions of this paper.

6. Bistellar Pseudoflips

We now investigate the properties of $\mathcal{PC}(D, \mathbf{h})$ for arbitrary height vectors \mathbf{h} , by defining flip operations in pseudo-complexes that result from controlled changes in \mathbf{h} . These operations generalize both the d -dimensional Lawson flip [12], [6] and the exchanging or removing flips in two-dimensional pseudo-triangulations [17], [1], [16]. The latter flip operations obey certain geodesics rules in polygons, whose meaning for $d \geq 3$ remained unclear. From now on, let $n = |\text{vert } D|$.

6.1. Moving Heights

Let \mathbf{h}_0 and \mathbf{h}_1 be two (generic) height vectors for $\text{vert } D$. Assume that \mathbf{h}_0 is convex, and that \mathbf{h}_0 is elementwise larger than \mathbf{h}_1 . Then f^* for \mathbf{h}_0 pointwise dominates f^* for \mathbf{h}_1 . We continuously deform $\mathcal{PC}(D, \mathbf{h}_0)$ into $\mathcal{PC}(D, \mathbf{h}_1)$ and study the changes in the structure of cells.

To this end, let $\mathbf{h}_\lambda = \lambda \mathbf{h}_1 + (1 - \lambda) \mathbf{h}_0$, for λ increasing from 0 to 1. Recall from Theorem 3 that, in $\mathcal{PC}(D, \mathbf{h}_0)$, all primary vertices are maximal and complete, and all cells are primary and are pseudo-simplices. $\mathcal{PC}(D, \mathbf{h}_\lambda)$ changes its shape exactly at values λ where \mathbf{h}_λ is not generic. Fix such a value λ . Consider a cell U of $\mathcal{PC}(D, \mathbf{h}_\lambda)$ which is not a cell of $\mathcal{PC}(D, \mathbf{h}_{\lambda-\varepsilon})$, for sufficiently small $\varepsilon > 0$. Denote with $\mathcal{PC}_{\lambda-\varepsilon}$ the restriction of $\mathcal{PC}(D, \mathbf{h}_{\lambda-\varepsilon})$ to U . The crucial observation is that, for $\lambda - \varepsilon$, f^* on $\text{int } U$ is determined by its values at those vertices that are complete with respect to the subcomplex $\mathcal{PC}_{\lambda-\varepsilon}$. This follows from Lemma 5(a) and (b). Therefore $\mathcal{PC}_{\lambda-\varepsilon}$ has exactly $d + 2$ complete vertices (apart from special cases which can be avoided by perturbing \mathbf{h}_0 slightly). In particular, U has at most $d + 2$ corners.

In $\mathcal{PC}(D, \mathbf{h}_{\lambda+\varepsilon})$ the polytope U is restructured into a cell complex $\mathcal{PC}_{\lambda+\varepsilon}$. The replacement of $\mathcal{PC}_{\lambda-\varepsilon}$ by $\mathcal{PC}_{\lambda+\varepsilon}$ is termed a *pseudoflip*. We face two different types of pseudoflips.

- (1) U has $d + 2$ corners, v_1, \dots, v_{d+2} . In $\mathcal{PC}_{\lambda+\varepsilon}$, each such vertex is still complete. Therefore, the set $\{v_1^*, \dots, v_{d+2}^*\}$ does not lie in a common hyperplane of \mathbb{R}^{d+1} , for $\lambda + \varepsilon$. The vertices v_1, \dots, v_{d+2} are not corners of a single cell of $\mathcal{PC}_{\lambda+\varepsilon}$. We call this flip an *exchanging pseudoflip*.
- (2) U has $d + 1$ corners, v_1, \dots, v_{d+1} . Let v_{d+2} be the additional complete vertex of $\mathcal{PC}_{\lambda-\varepsilon}$. As v_{d+2} is a noncorner of U we have $v_{d+2} \in \text{int conv } U$, and therefore f^* on $\text{int } U$ is determined by $\text{aff}\{v_1^*, \dots, v_{d+1}^*\}$ for $\lambda + \varepsilon$. That is, $\mathcal{PC}_{\lambda+\varepsilon}$ consists of a single primary cell, namely, U .
 - (a) If v_{d+2} is not a terminal of D then $f^*(v_{d+2})$ is determined by $\text{aff}\{v_1^*, \dots, v_{d+1}^*\}$ as well. The vertex v_{d+2} is not maximal any more.
 - (b) If v_{d+2} is a terminal of D then v_{d+2} stays maximal by Corollary 1, and f^* becomes discontinuous at v_{d+2} .

In both cases v_{d+2} is not a complete vertex of $\mathcal{PC}_{\lambda+\varepsilon}$. A flip of this type is called a *removing pseudoflip*.

Note that in (2)(a) the linear decrease of $f^*(v_{d+2})$ accelerates by this flip. This implies that v_{d+2} stays nonmaximal forever, such that—as in (2)(b) and in (1)—the

relative position of v_{d+2}^* with respect to $\text{aff}\{v_1^*, \dots, v_{d+1}^*\}$ does not change again. This gives:

Lemma 6. $\mathcal{PC}(D, \mathbf{h}_1)$ is obtained from $\mathcal{PC}(D, \mathbf{h}_0)$ by at most $\binom{n}{d+2}$ pseudoflips.

6.2. Anatomy of Pseudoflips

To study the structure of pseudoflips, we consider any complex $\mathcal{PC}(U, \mathbf{h})$ with exactly $d+2$ complete vertices v_1, \dots, v_{d+2} . By Lemma 5(c), full height is assumed at and only at v_1, \dots, v_{d+2} . Without loss of generality, let \mathbf{h} contain entries ∞ for all other vertices of U . Let \mathbf{h}^- be the vector obtained from \mathbf{h} by changing the signs of finite entries. Then, for any generic choice of heights h_1, \dots, h_{d+2} for v_1, \dots, v_{d+2} , one of the complexes $\mathcal{PC}(U, \mathbf{h})$ or $\mathcal{PC}(U, \mathbf{h}^-)$ has to arise, because the relative position of the $d+2$ points $\binom{v_i}{h_i}$ in \mathbb{R}^{d+1} already determines f^* .

As a consequence, each pseudoflip can be simulated by replacing $\mathcal{PC}(U, \mathbf{h})$ by $\mathcal{PC}(U, \mathbf{h}^-)$. Moreover, as $\mathcal{PC}(U, \mathbf{h}^-)$ has at most $d+2$ complete vertices, the cells of $\mathcal{PC}(U, \mathbf{h}^-)$ are pseudo-simplices, provided the same holds for $\mathcal{PC}(U, \mathbf{h})$. (If $\mathcal{PC}(U, \mathbf{h}^-)$ has only $d+1$ complete vertices then its only cell is the pseudo-simplex U .) Recalling from Theorem 3 that the original complex $\mathcal{PC}(D, \mathbf{h}_0)$ was a pseudo-complex, we get an inductive argument showing that the final complex $\mathcal{PC}(D, \mathbf{h}_1)$ is a pseudo-complex, too.

A removing pseudoflip transforms $\mathcal{PC}(U, \mathbf{h})$ into a single cell, and an exchanging pseudoflip transforms $\mathcal{PC}(U, \mathbf{h})$ into a complex with the same number, $d+2$, of complete vertices. (The inverse of a removing pseudoflip is also considered a valid pseudoflip; we call it an *inserting* pseudoflip.) $\mathcal{PC}(U, \mathbf{h})$ contains primary cells and, in general, also secondary cells, because secondary vertices may arise in the relative interior of faces of U . These faces are concave by Corollary 2. Neither the number of primary cells nor the number of secondary cells is bounded by a function of d . Already for $d=3$, $\Omega(k)$ primary cells and $\Omega(k^2)$ secondary cells may occur, for $k = |\text{vert } U|$; see Figs. 8 and 9. An upper bound for primary cells in this case is $O(k^2)$; see below.

For general d and arbitrary complexes $\mathcal{PC}(D, \mathbf{h})$, we can bound the number of primary cells using Lemma 7. The number of secondary cells can be shown to be finite but remains unclear.

Lemma 7. For given n , a maximal number of primary cells of $\mathcal{PC}(D, \mathbf{h})$ occurs if D is convex.

Proof. If $\mathcal{PC}(D, \mathbf{h})$ contains some primary vertex v that is a noncorner of a cell C , then applying an inserting pseudoflip by lowering the height of v splits C into more than one primary cell (and possibly several secondary cells) while leaving unaffected all cells of $\mathcal{PC}(D, \mathbf{h})$ different from C . Thus $\mathcal{PC}(D, \mathbf{h})$ contains a maximal number of primary cells only if all its primary vertices are complete. By Lemma 5(c), all primary vertices are maximal and continuous in this case, which implies that \mathbf{h} is convex. If \mathbf{h} is convex then all cells are primary by Theorem 3. However, the number of primary cells is maximized

if the visibility between primary vertices is not constrained by D , see Theorem 1, that is, if D is convex. \square

Combining Lemma 7 and Theorem 2, the number of primary cells does not exceed the maximal number of facets of a convex hull of n points in \mathbb{R}^{d+1} ; see [15]. Let us conclude the discussion in this section.

Theorem 4. $\mathcal{PC}(D, \mathbf{h})$ is a pseudo-complex for arbitrary (generic) \mathbf{h} . The number of primary cells of $\mathcal{PC}(D, \mathbf{h})$ is $O(n^{\lceil d/2 \rceil})$, for $n = |\text{vert } D|$. Given any two height vectors \mathbf{h} and \mathbf{h}' for D , the distance between $\mathcal{PC}(D, \mathbf{h})$ and $\mathcal{PC}(D, \mathbf{h}')$ by pseudoflips is $O(n^{d+2})$.

Remarks. The cells of $\mathcal{PC}(D, \mathbf{h})$ are not guaranteed to be *simple* pseudo-simplices. However, cells are always boundary-connected; see Lemma 3. Figure 7 illustrates a 3-polytope D where \mathbf{h} can be chosen such that $\mathcal{PC}(D, \mathbf{h})$ contains a cell with a tunnel.

Pseudoflips are bistellar operations that are “local” in the sense that each flip affects a subcomplex with at most $d + 2$ complete vertices. If the corresponding subdomain U is convex, then the classical bistellar flip [12], [6], also called a Lawson flip, is obtained. For $d = 2$, a pseudoflip is either an edge-exchanging flip [17] or an edge-removing (respectively, edge-inserting) flip [1] in a pseudo-triangulation. Removing pseudoflips of the type (2)(b) do not occur for $d = 2$, because all terminals are corners. In all these special cases, and in contrast to the general case, only $O(d)$ cells are affected by a single flip. A challenging open question is whether pseudoflip sequences do exist between pseudo-complexes in a given 3-polytope D such that cell sizes are always bounded by a constant; see Section 7.2.

6.3. Examples in \mathbb{R}^3

We illustrate some pseudoflips in \mathbb{R}^3 . From the foregoing discussion we know that the 3-polytope U where a flip takes place has at most five corners. They are labeled by numbers in the following figures. The pseudoflips are viewed best when imagining that corner 4 is the only corner whose height is altered.

Figure 4 shows a removing pseudoflip that is the simplest possible. The polytope U has four corners, 1, 2, 3, 4, and one noncorner, a . Before the flip, U contains two tetrahedral




Fig. 4. Removing pseudoflip.




Fig. 5. Exchanging pseudoflip.

cells $123a$ and $234a$ that share the triangular facet $23a$. The flip removes this facet and leaves a single cell, the pseudo-tetrahedron U . All involved cells are primary cells.

Figure 5 illustrates the simplest exchanging pseudoflip that is not a Lawson flip. Now the polytope U has five corners, 1, 2, 3, 4, 5. Again, before the flip, U contains two tetrahedral cells, namely, 1235 and 1345 . They are adjacent in the triangular facet 135 . After the flip, which destroys the facet 135 and creates the pseudo-triangular facet $234x$, two pseudo-tetrahedra arise as cells. Their corners are 1, 2, 3, 4 and 2, 3, 4, 5, respectively. The secondary vertex x arises as a noncorner of both cells. Still, all involved cells are primary cells.

The exchanging pseudoflip in Fig. 6 is more complicated. U contains two cells before the flip, the tetrahedron 1345 and the pseudo-tetrahedron with corners 1, 2, 3, 5 and the noncorner a . Both cells are primary cells. Their common triangular facet 135 is destroyed in the flip. After the flip, two new primary cells are present, namely, the pseudo-tetrahedra with corners 1, 2, 3, 4 and 2, 3, 4, 5, respectively. They are adjacent in the pseudo-triangular facet with corners 2, 3, 4. This facet (call it F) intersects the concave edge $1a$ at the secondary vertex x . However, F does not entirely split U , because F would define some secondary vertex on the *nonconcave* edge 15 , otherwise. Instead, the tetrahedron $145x$ arises as a secondary cell.

Figure 7 depicts an exchanging pseudoflip that creates a nonsimple cell. Three primary cells are present before the flip: The tetrahedra 1234 and 1235 , and the pseudo-tetrahedron with corners 2, 3, 4, 5 and noncorners a, b, c . These cells are pairwise adjacent in tri-




Fig. 6. Exchanging pseudoflip that generates a secondary cell.




Fig. 7. Exchanging pseudoflip that creates a tunnel.

angular facets which are destroyed in the flip. A single facet F with corners 1, 4, 5 and the secondary vertices x , y , z as noncorners is created. As xyz is a hole, F is not a valid pseudo-triangle, but rather a polygonal region with three corners. Two primary cells are adjacent in F . The cell with corners 1, 2, 4, 5 contains a tunnel, defined by the edges $2x$, ay , and bz .

An exchanging pseudoflip that gives rise to several secondary cells is illustrated in Fig. 8. Within U , the pseudo-tetrahedron with corners 1, 2, 3, 5 is separated by a




Fig. 8. Large exchanging pseudoflip.




Fig. 9. Large inserting pseudoflip.

triangular facet from the tetrahedron T with corner 4. The flip replaces this facet by various others. Three of them emanate from vertex 4, namely, the pseudo-triangles with corners 1, 2, 4 and 1, 3, 4, and 1, 4, 5, respectively. These facets separate three primary cells. Moreover, each such facet defines a secondary vertex on an edge of the former cell T . Vertex y is among them, defined by the facet with corners 1, 3, 4. This facet defines another secondary vertex, x . As the line segment xy does not lie on the boundary of U , triangular facets abx , bxy , xyd , ydc arise, which separate a polytope with six corners from the primary cells. This polytope splits into three secondary cells. A similar construction, where this polytope has k corners, shows that $\Theta(k)$ secondary cells can be created.

An inserting pseudoflip that generates a large number of primary cells and secondary cells is shown in Fig. 9. In the pseudo-tetrahedron U with corners 1, 2, 3, 4 and non-corners a, b, c , the edges $1b$ and $1c$ are nonconcave, and the edges $4b$ and $4c$ are concave. U splits into four primary cells, namely, the tetrahedra $123a$, $12ac$, $1abc$, and the pseudo-tetrahedron with corners 2, 3, 4, a and noncorners b, c . Note that ac is a nonconcave edge of the last cell, by the maximality of cells of f^* . Using a similar construction, where the number of noncorners of U is increased from 3 to k , we can create $\Theta(k)$ primary cells incident to the vertex 1.

In addition, $\Theta(k^2)$ secondary cells arise when a structure of size $\Theta(k)$ similar to the shape E is integrated into the boundary of U . The $\Theta(k)$ facets incident to the vertex 1 can be made to intersect $\Theta(k)$ concave boundary edges of U that stem from E . The resulting $\Theta(k^2)$ secondary vertices give rise to that many secondary cells, if the bold edges of E are chosen to be nonconcave.

7. Polytope Representation

Throughout this section let D be a fixed simple d -polytope that serves as a domain for the objects we consider. Let $n = |\text{vert } D|$. Consider the class of pseudo-complexes

$$\mathcal{R}(D) = \{\mathcal{P}\mathcal{C} \mid \exists \mathbf{h} \text{ with } \mathcal{P}\mathcal{C} = \mathcal{P}\mathcal{C}(D, \mathbf{h})\}.$$

The complexes in the class $\mathcal{R}(D)$ are called *constrained regular* (for D). We establish the existence of a convex polytope in \mathbb{R}^n that represents all the members of $\mathcal{R}(D)$. This generalizes the polytope constructions in [13] (the associahedron) and in [9] and [3] (the secondary polytope) which concern the regular simplicial complexes for convex domains D , as well as the polytope in [1] for constrained regular pseudo-triangulations of simple polygons D .

7.1. Function Lemma

Consider some piecewise linear function f on D . Call a point $x \in D$ *nonlinear* for f if there is no d -ball B centered at x such that f is linear on $B \cap D$. The subset of points of D where f is nonlinear will be denoted by $\mathcal{N}(f)$. Clearly, for any height vector \mathbf{h} for $\text{vert } D$ we have $\mathcal{N}(f^*) = \mathcal{P}\mathcal{C}(D, \mathbf{h})$, for the corresponding maximal locally convex function f^* . The following assertion shows that f^* can be “deformed” to being not locally convex. For two cell complexes $\mathcal{P}\mathcal{C}$ and $\mathcal{P}\mathcal{C}'$, we write $\mathcal{P}\mathcal{C} < \mathcal{P}\mathcal{C}'$ if each face of $\mathcal{P}\mathcal{C}$ is a subset of a face of $\mathcal{P}\mathcal{C}'$.

Lemma 8. *For every pseudo-complex $\mathcal{P}\mathcal{C} \in \mathcal{R}(D)$ and every height vector \mathbf{h} for $\text{vert } D$, there exists a unique piecewise linear function $f_{\mathcal{P}\mathcal{C}, \mathbf{h}}$ with $\mathcal{N}(f_{\mathcal{P}\mathcal{C}, \mathbf{h}}) < \mathcal{P}\mathcal{C}$ and such that $f_{\mathcal{P}\mathcal{C}, \mathbf{h}}(v_i) = h_i$ for each complete vertex v_i of $\mathcal{P}\mathcal{C}$.*

Proof. Let $\mathcal{P}\mathcal{C} \in \mathcal{R}(D)$. There is some height vector \mathbf{h}^* for $\text{vert } D$ such that $\mathcal{P}\mathcal{C} = \mathcal{P}\mathcal{C}(D, \mathbf{h}^*)$. Let v_1, \dots, v_k be the vertices of $\mathcal{P}\mathcal{C}$. We set up a system of k linear equations in variables t_1, \dots, t_k , with variable t_i corresponding to the height of vertex v_i , as follows.

If v_i is a complete vertex of $\mathcal{P}\mathcal{C}$ then we put $t_i = h_i^*$. Otherwise, there exists some cell C of $\mathcal{P}\mathcal{C}$ where v_i is a noncorner. Let v_1, \dots, v_{d+1} be the corners of C . By Lemma 5(a), we have $(\text{int } C)^* \subset \text{aff}\{v_1^*, \dots, v_{d+1}^*\}$. This gives a linear equation $t_i = a_1 t_1 + \dots + a_{d+1} t_{d+1}$ for the vertex v_i that is not complete. The resulting linear system $A \cdot \mathbf{t} = \mathbf{b}$ has a unique solution \mathbf{t} : by Lemma 5(c), the system describes the continuous extension of f^* from $\text{int } D$ to v_1, \dots, v_k , and f^* is unique for given D and \mathbf{h}^* .

Note that $\det A \neq 0$ by the uniqueness of \mathbf{t} . The vector \mathbf{b} coincides with \mathbf{h}^* at all entries for complete vertices. Now let \mathbf{h} be an arbitrary height vector for $\text{vert } D$, and let \mathbf{b}_h be the vector \mathbf{b} with entries for complete vertices replaced by those of \mathbf{h} . Because of $\det A \neq 0$, the system $A \cdot \mathbf{t} = \mathbf{b}_h$ has a unique solution $\mathbf{t} = A^{-1} \cdot \mathbf{b}_h$ as well. By construction, any solution \mathbf{t} defines a unique piecewise linear function $f_{\mathcal{P}\mathcal{C}, \mathbf{h}}$ with the desired properties. \square

Lemma 8 is a partial generalization to $d \geq 3$ of the “surface theorem” in [1]. That theorem, however, also holds for pseudo-triangulations which are not constrained regular.

7.2. Polytope Theorem

Using Lemma 8 the following theorem can be proved, in a way similar to a corresponding theorem for constrained regular pseudo-triangulations [1].

Theorem 5. *There exists a convex polytope $Q(D) \subset \mathbb{R}^n$ for D , whose vertices are in one-to-one correspondence with the pseudo-complexes in the class $\mathcal{R}(D)$. The edges of $Q(D)$ correspond to pseudoflips between the respective members of $\mathcal{R}(D)$. The diameter of $Q(D)$ is $O(n^{d+2})$.*

Proof. For a given member $\mathcal{PC} \in \mathcal{R}(D)$ and a given height vector \mathbf{h} for vert D , consider the unique piecewise linear function $f_{\mathcal{PC},\mathbf{h}}$ specified in Lemma 8. By definition, $f_{\mathcal{PC},\mathbf{h}}$ is locally convex exactly if $\mathcal{PC} = \mathcal{PC}(D, \mathbf{h})$. So assuming $\mathcal{PC} = \mathcal{PC}(D, \mathbf{h})$ implies $f_{\mathcal{PC},\mathbf{h}}(x) \leq f_{\mathcal{PC}',\mathbf{h}}(x)$, for all $\mathcal{PC}' \in \mathcal{R}(D)$ and all $x \in D$. Integrating both sides yields

$$\int_{x \in D} f_{\mathcal{PC},\mathbf{h}}(x) dx \leq \int_{x \in D} f_{\mathcal{PC}',\mathbf{h}}(x) dx.$$

Consider the left-hand side integral. The expressed volume is a linear (and homogenous) function of \mathbf{h} , because $f_{\mathcal{PC},\mathbf{h}}$ is piecewise linear and the value of $f_{\mathcal{PC},\mathbf{h}}$ at each vertex of \mathcal{PC} linearly depends on \mathbf{h} . Let $\mathbf{q}(\mathcal{PC})$ be the coefficient vector of this linear function. Then, from analogous observations for the integral on the right-hand side, we obtain $\mathbf{q}(\mathcal{PC}) \cdot \mathbf{h} \leq \mathbf{q}(\mathcal{PC}') \cdot \mathbf{h}$, for all $\mathcal{PC}' \in \mathcal{R}(D)$. We now interpret $\mathbf{q}(\mathcal{PC})$ as a point in \mathbb{R}^n , and consider the convex polytope

$$Q(D) = \text{conv}\{\mathbf{q}(\mathcal{PC}) \mid \mathcal{PC} \in \mathcal{R}(D)\}.$$

Define the equivalence classes of height vectors $\mathcal{H}(\mathcal{PC}) = \{\mathbf{h} \mid \mathcal{PC} = \mathcal{PC}(D, \mathbf{h})\}$. Then, for each $\mathbf{h} \in \mathcal{H}(\mathcal{PC})$, we have $\mathbf{q}(\mathcal{PC}) \cdot \mathbf{h} \leq \mathbf{q}(\mathcal{PC}') \cdot \mathbf{h}$, for all $\mathcal{PC}' \in \mathcal{R}(D)$. That is, the polytope $Q(D)$ lies in a halfspace of \mathbb{R}^n whose boundary contains $\mathbf{q}(\mathcal{PC})$. Thus $\mathbf{q}(\mathcal{PC})$ is a vertex of $Q(D)$.

Consider an edge e of $Q(D)$, and let e connect the vertices $\mathbf{q}(\mathcal{PC}_1)$ and $\mathbf{q}(\mathcal{PC}_2)$. Define $G = \mathcal{H}(\mathcal{PC}_1) \cap \mathcal{H}(\mathcal{PC}_2)$. By definition of these equivalence classes, no vector $\mathbf{h} \in G$ is generic for D . Moreover, we have $\mathbf{h} \cdot (\mathbf{q}(\mathcal{PC}_1) - \mathbf{q}(\mathcal{PC}_2)) = 0$. That is, the set G is orthogonal to e , and $\dim G = n - 1$. For $\mathbf{h} \in \text{relint } G$, there are exactly two pseudo-complexes \mathcal{PC}_1 and \mathcal{PC}_2 that yield the same function $f_{\mathcal{PC}_1,\mathbf{h}} = f_{\mathcal{PC}_2,\mathbf{h}}$. As a consequence, \mathcal{PC}_1 and \mathcal{PC}_2 differ by a single pseudoflip. The bound on the diameter of $Q(D)$ follows from Theorem 4. \square

For $d \geq 3$, not every simple d -polytope D can be *tetrahedrized*, i.e., partitioned into simplices without introducing additional vertices. The Schönhardt polytope [20] in Fig. 3 is a classical example. Moreover, even when D is convex and thus can be tetrahedrized,

not all its simplicial complexes have to belong to the class $\mathcal{R}(D)$; there exist convex 3-polytopes that admit nonregular tetrahedrizations [14]. Still, the constrained Delaunay simplicial complex $\mathcal{CD}(D)$ [21] exists for any tetrahedrizable d -polytope D , and we have $\mathcal{CD}(D) \in \mathcal{R}(D)$ because $\mathcal{CD}(D)$ arises from D when the parabolic height vector is applied. The assertion below, which is a corollary of Theorem 5, generalizes the well-known fact that every triangulation of a simple polygon D can be flipped to the constrained Delaunay triangulation [5] of D .

Corollary 4. *Let D be a simple d -polytope that is tetrahedrizable. Every pseudo-complex in $\mathcal{R}(D)$ can be flipped to the constrained Delaunay simplicial complex $\mathcal{CD}(D)$ for D . The number of pseudoflips is $O(n^{d+2})$.*

Remarks. Even when we start with a *simplicial* complex in $\mathcal{R}(D)$ and flip towards $\mathcal{CD}(D)$, the intermediate complexes need not be simplicial. They may well contain secondary vertices and cells for $d \geq 3$, but no secondary face remains in the final complex $\mathcal{CD}(D)$, by Theorem 3.

Unlike pseudo-triangles, pseudo-simplices for $d \geq 3$ are not tetrahedrizable, in general. A simple example is the Schönhardt polytope with a (sufficiently long) tetrahedron T attached to its top triangle. The top vertex of T sees only the remaining three vertices of T , and therefore T has to be part of every tetrahedrization of the resulting pseudo-tetrahedron. That is, the existence of such a tetrahedrization would imply that the Schönhardt polytope itself is tetrahedrizable.

Theorem 5 and Corollary 4 can be generalized by allowing the complexes to have vertices internal to D . (For example, choose any vertex set S and take $D = \text{conv } S$.) A prominent unsolved question is whether any two simplicial complexes on a vertex set S in \mathbb{R}^3 can be transformed into each other by Lawson flips. The answer is known to be negative for $d \geq 5$; see [19] that also gives a simplicial complex for a convex polytope in five dimensions where not a single Lawson flip can be applied. We plan to elaborate the consequences of pseudoflips for these and related questions in a separate paper.

7.3. Two Subclasses of $\mathcal{R}(D)$

There exist interesting subclasses of constrained regular pseudo-complexes. Define the class of *complete* pseudo-complexes as

$$\mathcal{C}(D) = \{\mathcal{PC} \mid \exists \mathbf{h} \text{ convex with } \mathcal{PC} = \mathcal{PC}(D, \mathbf{h})\}.$$

In each complex in $\mathcal{C}(D)$, all the vertices of D are maximal, and each such vertex is complete; see Section 5. The class $\mathcal{C}(D)$ contains the pseudo-complexes with a maximal number of primary cells, by Lemma 7, because if D is convex then all height vectors for D are convex as well. For $d = 2$, $\mathcal{C}(D)$ is the set of all possible triangulations of the simple polygon D : Every pseudo-triangulation of D can be generated by f^* when \mathbf{h} is chosen appropriately, and a pseudo-triangulation of D where all vertices of D are complete has to be a triangulation. For $d \geq 3$, however, $\mathcal{C}(D)$ may contain complexes that are not simplicial. Figure 3 in Section 3 gives an example in \mathbb{R}^3 .

At the other end of the spectrum, we may consider the class of *minimum* pseudo-complexes

$$\mathcal{M}(D) = \{\mathcal{PC}(D, \mathbf{h}) \mid h_i = \infty \text{ iff } v_i \text{ is a noncorner of } D\}.$$

Note that f^* is still finite on $\text{int } D$ for such height vectors \mathbf{h} , because D is contained in the convex hull of its corners. In each complex in $\mathcal{M}(D)$, the terminals of D are the only maximal vertices—the minimum possible by Corollary 1. Moreover, f^* is discontinuous at all terminals that are noncorners. By Lemma 5(c), this implies that the corners of D are the only complete vertices in such a complex. For $d = 2$, $\mathcal{M}(D)$ is precisely the class of minimum (or pointed) pseudo-triangulations of the polygon D . In [1], a pseudo-triangulation is defined to be minimum if it contains a minimum number of complete vertices. This definition is equivalent to the frequently used definition based on vertex pointedness [22], [17], but it generalizes nicely to higher dimensions.

The flip distance within the classes $\mathcal{M}(D)$ and $\mathcal{C}(D)$, respectively, is known to be $O(n)$ and $O(n^2)$ for $d = 2$; see [1] and [8].

8. Conclusions

We have introduced pseudo-simplicial complexes in \mathbb{R}^d and have discussed some of their basic properties. Our generalization nicely fits the existing concept of pseudo-triangulation in \mathbb{R}^2 , as far as the issues of flippability, regularity, and minimality are concerned. Already in \mathbb{R}^3 , however, pseudo-complexes show anomalies undesirable in practical applications. Secondary vertices not being part of the input do arise and increase the combinatorial complexity, and cells with tunnels add to the topological complexity.

Out of various interesting questions that are raised by the results of this paper, we consider the following three as most important:

- (1) Characterizing domain shapes and height vectors that prevent a prohibitive number of faces of a pseudo-complex.
- (2) Designing pseudoflip sequences that keep the number and the individual size of the involved cells small.
- (3) Developing efficient construction algorithms for pseudo-complexes.

For computing all primary cells in a pseudo-complex, the methods in [21] seem adaptable.

It may well be that a radically different approach to generalizing pseudo-triangulations to \mathbb{R}^d avoids certain unpleasant phenomena of pseudo-complexes as they stand now. Exploiting the relationship of pseudo-triangulations to structural rigidity concepts [22], [17], [10] seems a feasible alternative.

Acknowledgements

We are indebted to Günter Rote for pointing out errors in an earlier version of this paper. Thanks go to Paco Santos for discussions on the presented topic.

References

1. O. Aichholzer, F. Aurenhammer, P. Braß, H. Krasser. Pseudo-triangulations from surfaces and a novel type of edge flip. *SIAM Journal on Computing* **32** (2003), 1621–1653.
2. O. Aichholzer, F. Aurenhammer, T. Hackl. Pre-triangulations and liftable complexes. Manuscript, 2005.
3. L.J. Billera, P. Filliman, B. Sturmfels. Constructions and complexity of secondary polytopes. *Advances in Mathematics* **83** (1990), 155–179.
4. A. Brøndsted. *An Introduction to Convex Polytopes*. Springer-Verlag, New York, 1983.
5. L.P. Chew. Constrained Delaunay triangulations. *Algorithmica* **4** (1989), 97–108.
6. J.A. de Loera, F. Santos, J. Urrutia. The number of geometric bistellar neighbors of a triangulation. *Discrete & Computational Geometry* **21** (1999), 131–142.
7. H. Edelsbrunner, N.R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica* **15** (1996), 223–241.
8. 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, Singapore, 1995, pp. 225–265.
9. I.M. Gel'fand, M.M. Kapranov, A.V. Zelevinsky. Newton polyhedra of principal A-determinants. *Soviet Mathematics Doklady* **308** (1989), 20–23.
10. R. Haas, D. Orden, G. Rote, F. Santos, B. Servatius, H. Servatius, D. Souvaine, I. Streinu, W. Whiteley. Planar minimally rigid graphs and pseudo-triangulations. *Computational Geometry, Theory and Applications* **31** (2005), 31–61.
11. D. Kirkpatrick, J. Snoeyink, B. Speckmann. Kinetic collision detection for simple polygons. *International Journal of Computational Geometry & Applications* **12** (2002), 3–27.
12. C.L. Lawson. Properties of n -dimensional triangulations. *Computer Aided Geometric Design* **3** (1986), 231–246.
13. C.W. Lee. The associahedron and triangulations of the n -gon. *European Journal of Combinatorics* **10** (1989), 173–181.
14. C.W. Lee. Regular triangulations of convex polytopes. In P. Gritzmann, B. Sturmfels (eds.), *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science 4, AMS, Providence, RI, 1991, pp. 443–456.
15. P. McMullen. The maximal number of faces of a convex polytope. *Mathematica* **17** (1971), 179–184.
16. D. Orden, F. Santos. The polytope of non-crossing graphs on a planar point set. *Discrete & Computational Geometry* **33** (2005), 275–305.
17. G. Rote, F. Santos, I. Streinu. Expansive motions and the polytope of pointed pseudo-triangulations. In B. Aronov, S. Basu, J. Pach, M. Sharir (eds.), *Discrete & Computational Geometry—The Goodman–Pollack Festschrift*, Algorithms and Combinatorics, Springer, Berlin, 2003, pp. 699–736.
18. F. Santos. A point configuration whose space of triangulations is disconnected. *Journal of the American Mathematical Society* **13** (2000), 611–637.
19. F. Santos. Non-connected toric Hilbert schemes. *Mathematische Annalen* **332** (2005), 645–665.
20. E. Schönhardt. Über die Zerlegung von Dreieckspolyedern in Tetraeder. *Mathematische Annalen* **98** (1928), 309–312.
21. J.R. Shewchuk. Updating and constructing constrained Delaunay and constrained regular triangulations by flips. *Proceedings of the 19th Annual ACM Symposium on Computational Geometry*, 2003, pp. 181–190.
22. I. Streinu. A combinatorial approach to planar non-colliding robot arm motion planning. *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science* 2000, pp. 443–453.

Received June 7, 2004, and in revised form July 25, 2005. Online publication December 30, 2005.