Contributions to Discrete Mathematics

Volume 7, Number 1, Pages 68–81
ISSN 1715-0868

DEFORMATIONS OF ASSOCIAHEDRA AND VISIBILITY GRAPHS

SATYAN L. DEVADOSS, RAHUL SHAH, XUANCHENG SHAO, AND EZRA WINSTON

Abstract. Given an arbitrary polygon P with holes, we construct a polytopal complex analogous to the associahedron based on convex diagonalizations of P. This polytopal complex is shown to be contractible, and a geometric realization is provided based on the theory of secondary polytopes. We then reformulate a combinatorial deformation theory and present an open problem based on visibility which is a close cousin to the Carpenter's Rule theorem of computational geometry.

1. Introduction

The associahedron is a convex polytope whose face poset is based on nonintersecting diagonals of a convex polygon. This polytope and its generalizations continue to appear in a vast number of mathematical fields, including homotopy theory, representation theory, mathematical physics, geometric group theory, and computational biology. The vertices of the associahedron are enumerated by the famous Catalan numbers, corresponding to triangulations of a convex n-gon, bracketings on $n-1$ letters, or the set of rooted binary trees with $n-1$ leaves; Stanley offers over a hundred combinatorial and geometric bijections of this famous number [18]. In this paper, given an arbitrary polygon, we construct a polytopal complex analogous to the associahedron based on convex diagonalizations of P.

There are numerous polytopes which generalize the associahedron, such as those involving cluster algebras [5], graph associahedra [4], Coxeter systems [17], and generalized permutohedra [16]. Most notable to this paper, Orden and Santos [15] construct polytopes with face posets of noncrossing graphs of planar point sets. Our focus on nonconvex polygons is a related one involving constrained edges, resulting not in a polytope but a polytopal complex. Section 2 begins by providing the basic definitions of our construction, yielding a product structure on the facets. Section 3 focuses on two results, one showing the polytopal complex to be contractible, and the other providing a geometric realization based on the theory of secondary

Received by the editors June 21, 2010, and in revised form March 8, 2012.
2000 *Mathematics Subject Classification.* Primary 52B11, Secondary 57Q10, 68R05.
Key words and phrases. Visibility graph, associahedron, secondary polytope.

polytopes. This theory, spearheaded by the work of Gelfand et al. [11], is based on certain classes of triangulations of point sets, yielding numerous connections outside of mathematics [9].

Finally, Section 4 reformulates the combinatorial deformation theory in terms of visibility. This leads to an open problem (with partial solution) which can be viewed as a close cousin to the Carpenter's Rule theorem of computational geometry [6]. Here, instead of convexifying polygons while fixing edge lengths, we ask for convexification without losing internal visibility of vertices.

2. Associahedral Complex from Polygons

2.1. Let P be a simple planar polygon with labeled vertices. Unless mentioned otherwise, assume the vertices of P in general position, with no three collinear vertices. A *diagonal* of P is a line segment contained in the interior of P connecting two vertices. A *diagonalization* of P is a partition of P into smaller polygons using noncrossing[1] diagonals of P. Let a *convex* diagonalization of P be one which divides P into smaller convex polygons.

Definition 1. Let $\pi(P)$ be the poset of all convex diagonalizations of P where for $a \prec a'$ if a is obtained from a' by adding new diagonals.

It was independently proven by Lee [12] and Haiman (unpublished) that there exists a convex polytope $\mathcal{K}_n$ of dim $n-3$, called the *associahedron*, whose face poset is isomorphic to $\pi(P)$. Almost twenty years before this result was discovered, the associahedron had originally been defined by Stasheff for use in homotopy theory in connection with associativity properties of H-spaces [19]. Figure 1 shows examples of associahedra with labelings of certain faces. Classically, the associahedron is based on all bracketings of $n-1$ letters and denoted as K_{n-1}; we use the script notation $\mathcal{K}_n$ with an index shift for ease of notation in our polygonal context.

We now consider extending the associahedron for the case of arbitrary simple polygons P. A *polytopal complex* S is a finite collection of convex polytopes (containing all the faces of its polytopes) such that the intersection of any two of its polytopes is a (possibly empty) common face of each of them. The dimension of the complex S, denoted $\dim(S)$, is the largest dimension of a polytope in S.

Theorem 2. *For a polygon P with n vertices, there exists a polytopal complex $\mathcal{K}_P$ whose face poset is isomorphic to $\pi(P)$. Moreover, $\mathcal{K}_P$ is a subcomplex of the associahedron $\mathcal{K}_n$ of dimension $n-3-d(P)$, where $d(P)$ is the minimum number of diagonals required to diagonalize P into convex polygons.*

Proof. Let $p_1, \ldots, p_n$ be the vertices of P labeled cyclically. For a convex n-gon Q, let $q_1, \ldots, q_n$ be its vertices again with clockwise labeling. The

[1] Mention of *diagonals* will henceforth mean noncrossing ones.

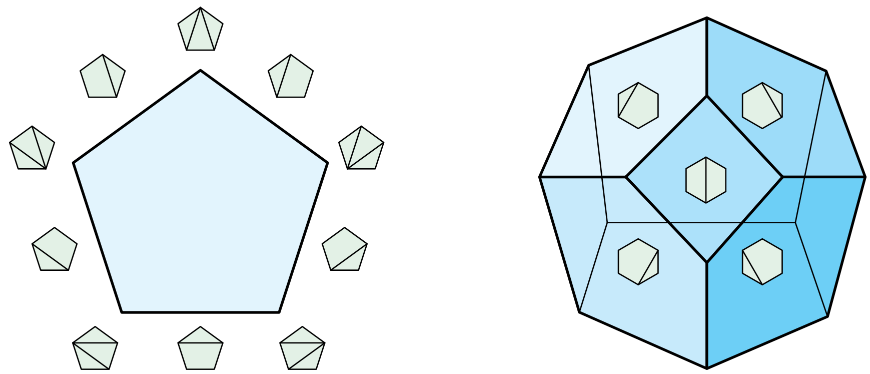

(a) (b)

FIGURE 1. Associahedra (a) $\mathcal{K}_5$ and (b) $\mathcal{K}_6$.

natural mapping from P to Q (taking p_i to q_i) induces an injective map $\phi : \pi(P) \longrightarrow \pi(Q)$. Assign to $t \in \pi(P)$ the face of K_n that corresponds to $\phi(t) \in \pi(Q)$. It is trivial to see that $\phi(t_1) \prec \phi(t_2)$ in $\pi(Q)$ if $t_1 \prec t_2$ in $\pi(P)$.

Moreover, for any $\phi(t)$ in $\pi(Q)$ and any diagonal (q_i, q_j) which does not cross the diagonals of $\phi(t)$, we see that (p_i, p_j) does not cross any diagonal of t. So if a face f of $\mathcal{K}_n$ is contained in a face corresponding to $\phi(t)$, then there exists a diagonalization $t' \in \pi(P)$ where $\phi(t')$ corresponds to f and $t' \prec t$. Since the addition of any noncrossing diagonals to a convex diagonalization is still a convex diagonalization, the intersection of any two faces[2] is also a face in $\mathcal{K}_P$. So $\mathcal{K}_P$ satisfies the requirements of a polytopal complex and (due to the map ϕ) is a subcomplex of $\mathcal{K}_n$.

The dimension of a polytopal complex is defined as the maximum dimension of any face. In the associahedron $\mathcal{K}_n$, a face of dimension k corresponds to a convex diagonalization with $n - 3 - k$ diagonals. The result follows since ϕ is an injection. □

Example. Figure 2(a) shows the polytopal complex $\mathcal{K}_P$ for the deformed hexagon P, made from two line segments glued to opposite vertices of a square. Note how this complex appears as a subcomplex of $\mathcal{K}_6$ from Figure 1(b). Figure 2(b) displays the labeling of the complex by $\pi(P)$, where the number of diagonals in each diagonalization is constant across dimensions of the faces.

Remark. In contrast to the convex case, the nonconvex case depends on the detailed geometry of the polygon. Consider the *bow tie* hexagon of Figure 2: if the two indentations were not on the same vertical line, and the bottom indentation could be pushed up to a peak higher than the valley of the top indentation, the resulting complex $\mathcal{K}_P$ would be entirely different.

[2]Such an intersection could possibly be empty if diagonals are crossing.

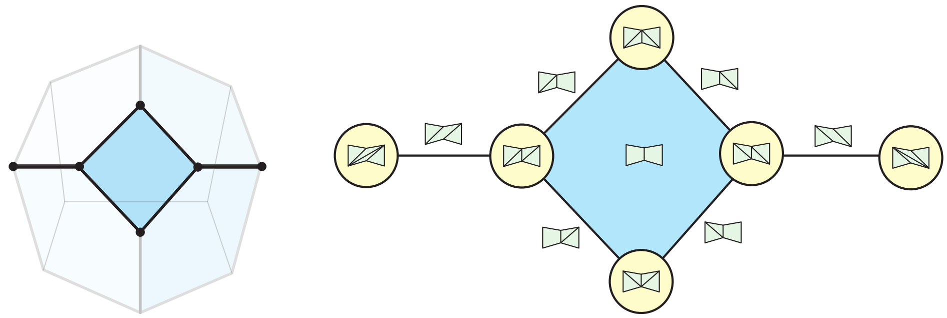

(a) (b)

FIGURE 2. A polytopal complex and its labeling.

2.2. We begin this section by considering arbitrary (not just convex) diagonalizations of P and the resulting geometry of $\mathcal{K}_P$. Let $\Delta = \{d_1, \ldots, d_k\}$ be a set of noncrossing diagonals of P, and let $\mathcal{K}_P(\Delta)$ be the collection of faces in $\mathcal{K}_P$ corresponding to all diagonalizations of P containing Δ.

Lemma 3. *$\mathcal{K}_P(\Delta)$ is a polytopal complex.*

Proof. If t is a diagonalization of P containing Δ, then any t' in $\pi(P)$ must also contain Δ if $t' \prec t$. Thus there must be a face in $\mathcal{K}_P(\Delta)$ that corresponds to t'. Furthermore, consider faces f_1 and f_2 in $\mathcal{K}_P(\Delta)$ corresponding to diagonalizations t_1 and t_2 of P. Then the intersection of f_1 and f_2 must correspond to a diagonalization including Δ since $f_1 \cap f_2$ is the (possibly empty) face corresponding to all convex diagonalizations that include every diagonal of t_1 and t_2. $\square$

Theorem 4. *If diagonals $\Delta = \{d_1, \ldots, d_k\}$ divide P into (not necessarily convex) polygons $Q_0, \ldots, Q_k$, then $\mathcal{K}_P(\Delta)$ is isomorphic to the Cartesian product $\mathcal{K}_{Q_0} \times \cdots \times \mathcal{K}_{Q_k}$.*

Proof. We use induction on k. When $k = 1$, any face $f \in \mathcal{K}_P(d)$ corresponds to a convex diagonalization of Q_0 paired with a convex diagonalization of Q_1. Thus, a face of $\mathcal{K}_{Q_0} \times \mathcal{K}_{Q_1}$ exists for each pair of faces (f_0, f_1), for $f_0 \in \mathcal{K}_{Q_0}$ and $f_1 \in \mathcal{K}_{Q_1}$. For $k > 1$, order the diagonals such that d_k divides P into polygons $Q_* = Q_0 \cup \cdots \cup Q_{k-1}$ and Q_k. A face in $\mathcal{K}_P(\Delta)$ corresponds to a convex diagonalization t_1 of Q_* and a convex diagonalization t_2 of Q_k. The pair $(t_1, t_2) \in \mathcal{K}_P(\Delta)$ corresponds to a face in $\mathcal{K}_{Q_*}(\Delta \setminus d_k) \times \mathcal{K}_{Q_k}$. By the induction hypothesis, $\mathcal{K}_{Q_*}(\Delta \setminus d_k)$ is isomorphic to $\mathcal{K}_{Q_0} \times \mathcal{K}_{Q_1} \times \cdots \times \mathcal{K}_{Q_{k-1}}$. $\square$

Remark. This product structure on the faces of $\mathcal{K}_P$ provides a generalization of the *mosaic operad* related to the real moduli space of curves [7].

For a polytopal complex $\mathcal{K}_P$, the *maximal* elements of its face poset $\pi(P)$ are analogous to facets of convex polytopes. A face f of $\mathcal{K}_P$ corresponding to a diagonalization $t \in \pi(P)$ is a *maximal face* if there does not exist $t' \in \pi(P)$

such that $t \prec t'$. Thus a maximal face of $\mathcal{K}_P$ has a convex diagonalization of P using the *minimal* number of diagonals. Figure 3 shows a polygon P along with six minimal convex diagonalizations of P. As this shows, such diagonalizations may not necessarily have the same number of diagonals. In other words, the polytopal complex is not *pure*.

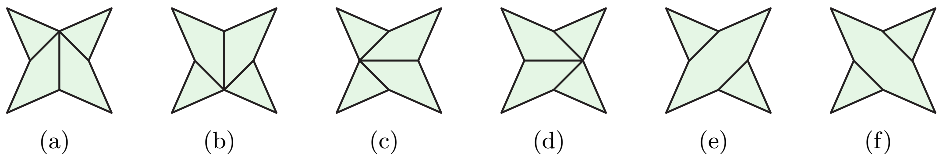

FIGURE 3. Six minimal convex diagonalizations of a polygon.

Example. Figure 4 shows an example of $\mathcal{K}_P$ for the polygon P from Figure 3. It is a polyhedral subcomplex of the 5-dimensional convex associahedron $\mathcal{K}_8$. We see that $\mathcal{K}_P$ is made of six maximal faces, four squares (where each square is a product $\mathcal{K}_4 \times \mathcal{K}_4$ of line segments) and two $\mathcal{K}_6$ associahedra. Each of these six faces correspond to the minimal convex diagonalizations from Figure 3.

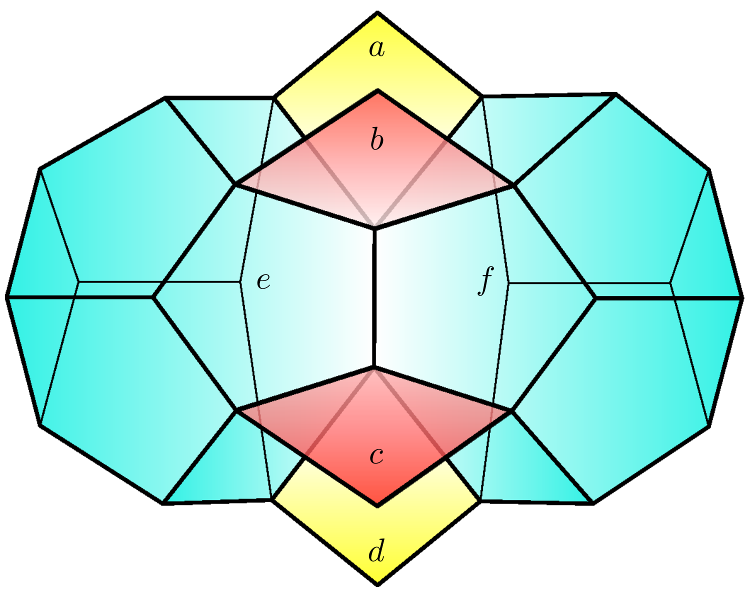

FIGURE 4. The complex $\mathcal{K}_P$ from the polygon in Figure 3.

3. Topological and Geometric Properties

3.1. We now prove the polytopal complex $\mathcal{K}_P$ is contractible. In 1998, Edelman and Reiner [10] showed a similar result: For a planar point set A and an arbitrary planar simplicial complex P which uses only vertices in A, they considered the *Baues poset* Baues(P, A) and observed that its order complex is contractible. When P is restricted to be a nonconvex polygon with vertex set A, contractibility of $\mathcal{K}_P$ becomes a special case. The technique of their proof is advanced, using a version of *deletion-contraction* from matroid theory along with topological analysis. More recently, Braun and Ehrenborg [3] studied an analogous complex $\theta(P)$ for nonconvex polygons P, seen as the combinatorial dual of $\mathcal{K}_P$. Their central result showed $\theta(P)$ to be homeomorphic to a ball, akin to showing $\mathcal{K}_P$ contractible, based on discrete Morse theory and a pairing lemma of Linusson and Shareshian [13].

Compared to both of these approaches, our proof is much shorter, using simple techniques based solely on the geometry of reflex vertices. A vertex of a polygon is called *reflex* if the diagonal between its two adjacent vertices cannot exist. Note that every nonconvex polygon has a reflex vertex.

Lemma 5. *For any reflex vertex v of a nonconvex polygon P, every element of $\pi(P)$ has at least one diagonal incident to v.*

Proof. Assume otherwise and consider an element of $\pi(P)$. In this convex diagonalization, since there is no diagonal incident to v, there exists a unique subpolygon containing v. Since v is reflex, this subpolygon cannot be convex, which is a contradiction. □

Lemma 6. *Let $F = \{f_1, f_2, \ldots, f_k\}$ be a set of faces of $\mathcal{K}_P$ such that $\bigcap_F f_i$ is nonempty. If $\bigcap_{F'} f_i$ is contractible for* every *$F' \subset F$, then $\bigcup_F f_i$ is contractible.*

Proof. We prove this by induction on the number of faces in F. A single face is trivially contractible. Now assume $\bigcup_F f_i$ is contractible. For a face $f_{k+1} \notin F$ of $\mathcal{K}_P$, let $G = \{f_{k+1}\} \cup F$ so that $\bigcap_G f_i$ is nonempty. Since f_{k+1} intersects the intersection of the face F and since this intersection is nonempty and contractible, we can deformation retract $\bigcup_G f_i$ onto $\bigcup_F f_i$ and hence maintain contractibility. □

Theorem 7. *For any polygon P, the polytopal complex $\mathcal{K}_P$ is contractible.*

Proof. We prove this by induction on the number of vertices. For the base case, note that $\mathcal{K}_P$ is a point for any triangle P. Now let P be a polygon with n vertices. If P is convex, then $\mathcal{K}_P$ is the associahedron $\mathcal{K}_n$ and we are done. For P nonconvex, let v be a reflex vertex of P. Since each diagonal d of P incident to v separates P into two smaller polygons Q_1 and Q_2, by our hypothesis, $\mathcal{K}_{Q_1}$ and $\mathcal{K}_{Q_2}$ are contractible. Theorem 4 shows that $\mathcal{K}_P(d)$ is isomorphic to $\mathcal{K}_{Q_1} \times \mathcal{K}_{Q_2}$, resulting in $\mathcal{K}_P(d)$ to be contractible.

Let $\Delta = \{d_1, \ldots, d_k\}$ be the set of all diagonals incident to v. Since Δ is a set of *noncrossing* diagonals, then $\bigcap_\Delta \mathcal{K}_P(d_i)$ is nonempty. Furthermore, for any subset $\Delta' \subset \Delta$, we have $\mathcal{K}_P(\Delta') = \bigcap_{\Delta'} \mathcal{K}_P(d)$. By Theorem 4, this is a product of contractible pieces, and thus itself is contractible. Therefore, by Lemma 6, the union of the complexes $\mathcal{K}_P(d_i)$ is contractible. However, since Lemma 5 shows that this union is indeed $\mathcal{K}_P$, we are done. □

Remark. The previous constructions and arguments can be extended to include planar polygons with holes. These *generalized polygons* are bounded, connected planar regions whose boundary is the disjoint union of simple polygonal loops. We state the following and leave the straightforward proof to the reader.

Corollary 8. *For a generalized polygon R with $h+1$ boundary components, there exists a polytopal complex $\mathcal{K}_R$ whose face poset is isomorphic with $\pi(R)$. The dimension of $\mathcal{K}_R$ is $n+3h-d(R)-3$, where $d(R)$ is the minimum number of diagonals required to diagonalize R into convex polygons. Moreover, $\mathcal{K}_R$ is contractible.*

Figure 5 shows an example of the associahedron of a pentagon with a triangular hole, whose maximal faces are 8 cubes and 3 squares. Similar to the polygonal case, the complexes $\mathcal{K}_R$ can be obtained by gluing together different polytopal complexes $\mathcal{K}_P$, for various polygons P. The geometry of R is crucial for the geometry $\mathcal{K}_R$: as the size of the internal triangle increases inside the pentagon, the complex $\mathcal{K}_R$ will deform as well.

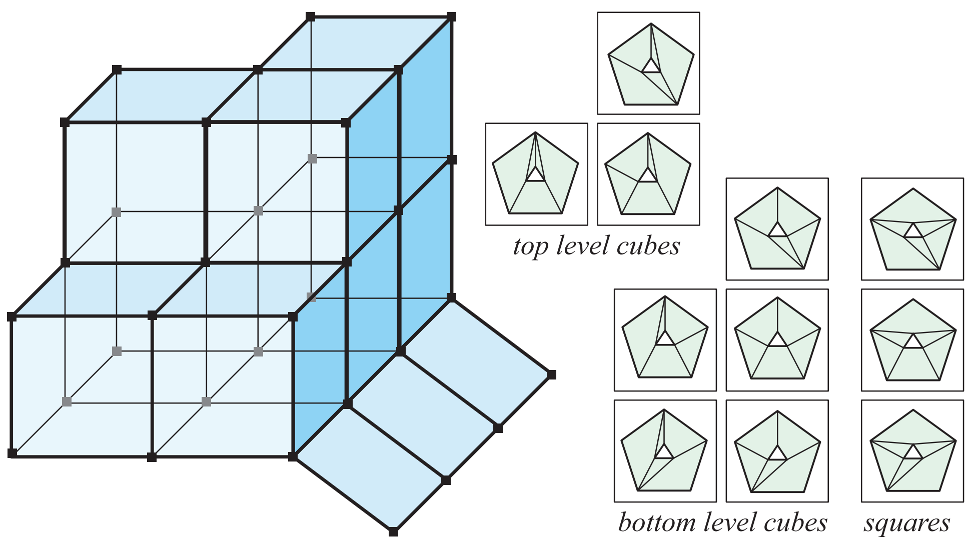

FIGURE 5. The associahedral complex for a pentagon with a triangular hole, along with its 11 maximal faces.

Remark. The polytopal complex depends on the detailed geometry of the generalized polygon. If the size of the internal triangle of Figure 5 is substantially increased, the resulting complex $\mathcal{K}_R$ would be vastly different.

3.2. We now turn to geometric realizations of $\mathcal{K}_P$, with integer coordinates for its vertices. There are numerous realizations of the classical associahedron and its generalizations, such as those given by Devadoss [8], Loday [14], and Postnikov [16]. For nonconvex polygons, since $\mathcal{K}_P$ is a subcomplex of

$\mathcal{K}_n$ by Theorem 2, any such realization extends to a realization of $\mathcal{K}_P$ with integer coordinates. However, since our interests are in the deformations of the underlying polygons in the plane, we turn to a realization based on secondary polytopes which is more tailored for our situation.

Secondary polytopes were developed by Gelfand, Kapranov, and Zelevinsky [11], of which we consider one case: Let P be a polygon with vertices $p_1, \ldots, p_n$. For a triangulation T of P, let

$$\phi(p_i) = \sum_{p_i \in \Delta \in T} \operatorname{area}(\Delta)$$

be the sum of the areas of all triangles Δ which contain the vertex p_i. Let the *area vector* of T be

$$\Phi(T) = (\phi(p_1), \ldots, \phi(p_n)).$$

The *secondary polytope* of $\Sigma(P)$ of a polygon P is the convex hull of the area vectors of all triangulations of P. In particular, when P is a convex n-gon, the secondary polytope $\Sigma(P)$ is a realization of the associahedron $\mathcal{K}_n$.

We show that the secondary polytope of a nonconvex polygon has all its area vectors on its hull, as is the case for a convex polygon. However, since the secondary polytope of a nonconvex polygon is *not* a subcomplex of the secondary polytope for convex polygon, this result is not trivial.

Theorem 9. *For any polygon P, and any triangulation T of P, the area vectors $\Phi(T)$ lie on the hull of $\Sigma(P)$.*

Proof. Fix a triangulation T of P. We first show that there is a height function ω on P which raises the vertices of T to a locally convex surface in $\mathbb{R}^3$, that is, a surface which is convex on every line segment in P. Choose an edge e of P to be its base so that the dual tree of T is rooted at e. Starting from the root and moving outward, assign increasing numbers m_i to each consecutive triangle Δ_i in the tree. Define a height function

$$\omega(p_i) = \min\{m_k \mid p_i \in \Delta_k\}$$

for each vertex p_i of P. Observe that for every pair of adjacent triangles Δ_1 and Δ_2 (in the dual tree), we can choose the value m_i to be large enough such that the planes containing $\omega(\Delta_1)$ and $\omega(\Delta_2)$ are distinct and meet in a convex angle.

In order to show that $\Phi(T)$ lies on the hull of $\Sigma(P)$, we construct a linear function $\rho(v)$ on $\Sigma(P)$ such that $\rho(\Phi(T))$ is a unique minimum of this function on $\Sigma(P)$. For any v in $\Sigma(P)$, define $\rho(v) = \langle \omega(T), v \rangle$ to be the inner product of the vectors $v \in \mathbb{R}^n$ and

$$\omega(T) = (\omega(p_1), \ldots, \omega(p_n)).$$

For a triangle Δ of T with vertices p_i, p_j, p_k, the volume in $\mathbb{R}^3$ enclosed between Δ and the lifted triangle $\omega(\Delta)$ can be written as

$$\frac{\omega(p_i) + \omega(p_j) + \omega(p_k)}{3} \operatorname{area}(\Delta).$$

The volume between the surface on which the $\omega(p_i)$'s lie and the plane is

$$\begin{aligned}\sum_{\Delta\in T}\frac{\omega(p_i)+\omega(p_j)+\omega(p_k)}{3}\ \mathrm{area}(\Delta) &= \sum_{i=1}^{n}\left[\frac{\omega(p_i)}{3}\sum_{p_i\in\Delta\in T}\mathrm{area}(\Delta)\right]\\ &= \sum_{i=1}^{n}\frac{\omega(p_i)}{3}\phi(p_i)\\ &= \frac{1}{3}\langle\omega(T),\Phi(T)\rangle.\end{aligned}$$

Since ω lifts T to a locally convex surface S, we know that w will lift any $T'\neq T$ to a surface S' above S. Thus $\langle\omega(T),\Phi(T)\rangle < \langle\omega(T),\Phi(T')\rangle$, implying all vertices of $\Sigma(P)$ lie on the hull. □

Corollary 10. *If P is nonconvex, then a subset of the faces of $\Sigma(P)$ yield a realization of $\mathcal{K}_P$.*

Proof. For any face f of $\mathcal{K}_P$, let $T_1,\ldots,T_k$ be the triangulations corresponding to the vertices of f. We use the same argument as the theorem above to show there exists a height function ω such that $\langle\omega,\Phi(T)\rangle$ is constant for any $T\in\{T_1,\ldots,T_k\}$ and $\langle\omega,\Phi(T)\rangle < \langle\omega,\phi(T')\rangle$ for any $T'\notin\{T_1,\ldots,T_k\}$. □

4. Visibility Graphs

4.1. This final section places these polytopal complexes in a larger setting, from the viewpoint of continuous and discrete deformations. As a convex n-gon is transformed continuously in the plane to an n-gon with a unique triangulation, its associated polytopal complex goes through a discrete deformation, starting from the associahedron $\mathcal{K}_n$ polytope and ending at a topological point. Since our objects remain contractible during this process, as given by Theorem 7, the deformation can be considered a *discrete* analog of a deformation retract.

In order to understand the underlying combinatorics, we show that the natural setting for study comes from the notion of *visibility* and a computational geometric perspective. In this section, we only consider simple polygons with vertices labeled $\{1,\ldots,n\}$ in this cyclic order. As before, assume the vertices of P in general position, with no three collinear vertices. The *visibility graph* $\mathcal{V}(P)$ of a labeled polygon P is the labeled graph with the same vertex set as P, with e as an edge of $\mathcal{V}(P)$ if e is an edge or diagonal of P. We say two polygons P_1 and P_2 are *$\mathcal{V}$-equivalent* if $\mathcal{V}(P_1)=\mathcal{V}(P_2)$.

There is a natural relationship between the graph $\mathcal{V}(P)$ and the polytopal complex $\mathcal{K}_P$: if polygons P_1 and P_2 are $\mathcal{V}$-equivalent then $\mathcal{K}_{P_1}$ and $\mathcal{K}_{P_2}$ yield the same complex. We wish to classify polygons under a stronger relationship than $\mathcal{V}$-equivalence. For a polygon P, let (x_i,y_i) be the coordinate of its i-th vertex in $\mathbb{R}^2$. We associate a point $\gamma(P)$ in $\mathbb{R}^{2n}$ to P where

$$\gamma(P)=(x_1,y_1,x_2,y_2,\ldots,x_n,y_n).$$

Since P is labeled, it is obvious that γ is injective but not surjective.

Definition 11. Two polygons P_1 and P_2 are *$\mathcal{V}$-isotopic* if there exists a continuous map $f : [0,1] \longrightarrow \mathbb{R}^{2n}$ such that $f(0) = \gamma(P_1)$, $f(1) = \gamma(P_2)$, and for every $t \in [0,1]$, $f(t) = \gamma(P)$ for some simple polygon P where $\mathcal{V}(P) = \mathcal{V}(P_1)$.

It follows from the definition that two polygons that are $\mathcal{V}$-isotopic are $\mathcal{V}$-equivalent, whereas the converse is not necessarily true. For a polygon P with n vertices, let $\mathcal{D}(P)$ be the $\mathcal{V}$-isotopic equivalence class containing the polygon P and let $\mathcal{D}$ be the set of all such equivalence classes of polygons with n vertices. We give $\mathcal{D}$ a poset structure: for two n-gons P_1 and P_2, the relation $\mathcal{D}(P_2) \prec \mathcal{D}(P_1)$ is given if the following two conditions hold:

1. $\mathcal{V}(P_1)$ is obtained by adding one more edge to $\mathcal{V}(P_2)$.
2. There exists a continuous map $f : [0,1] \longrightarrow \mathbb{R}^{2n}$, such that $f(0) = \gamma(P_1)$, $f(1) = \gamma(P_2)$, and for every $t \in [0,1/2)$, $f(t) = \gamma(P)$ for some polygon P with $\mathcal{V}(P) = \mathcal{V}(P_1)$, while for every $t \in (1/2,1]$, $f(t) = \gamma(Q)$ for some polygon Q with $\mathcal{V}(Q) = \mathcal{V}(P_2)$.

If P_1 and P_2 are $\mathcal{V}$-isotopic, let $\mathcal{D}(P_1) = \mathcal{D}(P_2)$. Taking the transitive closure of $\preceq$ yields the *deformation poset* $\mathcal{D}$. A natural ranking exists on $\mathcal{D}$ based on the number of edges of the visibility graphs.

Example. The top image of Figure 6 shows a subdiagram of the Hasse diagram for $\mathcal{D}$ for 6-gons, where we have forgone the labeling on the vertices. A polygonal representative for each equivalence class is drawn along with its underlying visibility graph. Each element of $\mathcal{D}$ corresponds to a polytopal complex $\mathcal{K}_P$ as displayed in the bottom image. Notice that as the polygon deforms and loses visibility edges, its associated complex collapses into a vertex of $\mathcal{K}_6$.

4.2. It is easy to see that the deformation poset $\mathcal{D}$ is connected: notice that $\mathcal{D}$ has a unique *maximum* element corresponding to the convex polygon. Given any polygon P in the plane, one can move its vertices, deforming P into convex position, making each element of $\mathcal{D}$ connected to the maximum element. Since the vertices of P are in general position, we can insure that the visibility graph of P changes only one diagonal at a time during the deformation. However, the visibility graph of the deforming polygon might gain and lose edges, moving up and down the poset $\mathcal{D}$.

We are interested in the combinatorial structure of the deformation poset beyond connectivity. The maximum element of $\mathcal{D}$ corresponds to the convex n-gon (with $\binom{n}{2}$ edges in its visibility graph) whereas the minimal elements (which are not unique in $\mathcal{D}$) correspond to polygons with unique triangulations (with $2n-3$ edges in each of their visibility graphs). This implies that the height of the deformation poset is $\binom{n}{2} - 2n + 4$. We pose the following problem and close this paper with a discussion of partial results.

Visibility Deformation Problem. *Show that every maximal chain of $\mathcal{D}$ has length $\binom{n}{2} - 2n + 4$.*

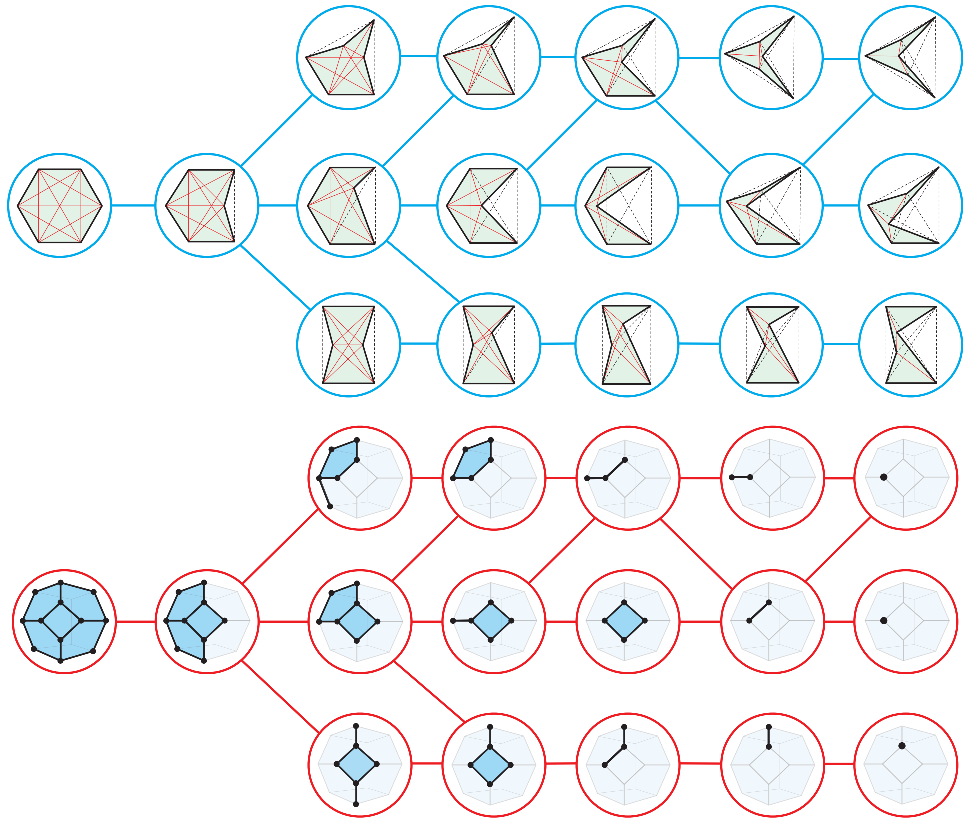

FIGURE 6. A subdiagram of the Hasse diagram of $\mathcal{D}$ for 6-gons along with the corresponding subcomplexes of $\mathcal{K}_6$.

More loosely, does there exist a deformation of any simple polygon into a polygon with a unique triangulation such that throughout the deformation, the visibility of the polygon monotonically decreases? And moreover, does there exist a deformation of any simple polygon into a convex polygon such that throughout the deformation, the visibility of the polygon monotonically increases? This latter question was recently given a positive answer in [1] based on a novel idea of *visibility-increasing* edges. Indeed, this can be viewed as a close cousin to the Carpenter's Rule theorem [6], but instead of convexifying polygons with fixed edge lengths, we ask for convexification without losing internal visibility of vertices.

We close with a result which holds for *star-shaped* polygons. A polygon P is star-shaped if there exists a point $p \in P$ such that p is visible to all points of P.

Theorem 12. *Let P be a star-shaped polygon. There exists a chain in $\mathcal{D}$ from P to the maximum element.*

Proof. Let x be a point in the *kernel* of P, the set of points which are visible to all points of P. Choose an ε-neighborhood around x contained in the kernel. For any $a \in P$, let $p(a)$ be a point on the boundary of P which is the intersection of the ray from x passing through a with the boundary. Let a' be the point on the ray from x passing through a such that $d(a', x) = \varepsilon \cdot r(a)$, where

$$r(a) = \frac{d(a,x)}{d(p(a),x)}.$$

Let ϕ be the map from a to a'. We thus construct a linear map $f : P \times [0,1] \to P$ where $f(P,0) = P$ and $f(P,1) = \phi(P)$ and where

$$\frac{\partial}{\partial t} f(a,t) = r(a).$$

For any two visible vertices a and b of P, consider the triangle abx. There cannot be any vertices of P contained in the triangle. If for any vertex c of P, the ray from x passing through c intersects the line segment ab at a point z, then $d(c,x) > d(z,x)$ and thus for no $t \in [0,1)$ can $d(f(c,t),x) \leq d(f(z,t),x)$. So no visibility is lost during the transformation, but notice that $\phi(P)$ is a circle. However, if we apply $f(a,t)$ only to the vertices of P and map any point z on an edge (a,b) of P to z' on the edge between $f(a,t)$ and $f(b,t)$, we find that we get a polygon at every t. Moreover, the edge is always further from c than $f(z,t)$ for every $t \in [0,1]$, and thus visibility is still maintained. □

A natural approach is to discretize this problem into moving vertices of the polygon one by one. In other words, for any polygon, does there exist *one vertex* which can be moved that increases visibility? Based on this work, Aichholzer et al. [2] have recently provided an elegant counterexample to this claim, seen in Figure 7. A partial collection of the visibility edges of

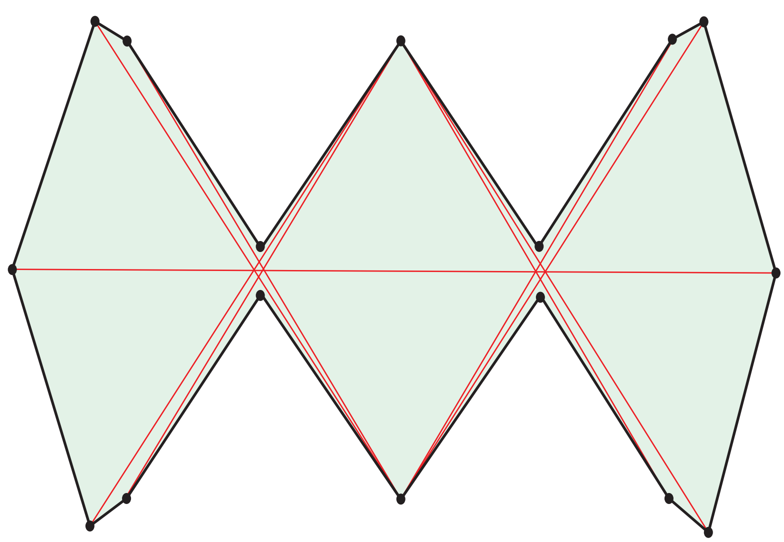

FIGURE 7. No vertex may be moved to increase visibility.

this polygon is given in red. No vertex of this polygon may be moved which strictly increases visibility.

Acknowledgments

We thank Oswin Aichholzer, Joseph O'Rourke, and Victor Reiner for helpful discussions. We are also grateful to Williams College and to the NSF for partially supporting this work with grant DMS-0353634. The first author is also thankful to Lior Pachter, Bernd Sturmfels, and UC Berkeley during his sabbatical leave, where this work was finished.

References

1. O. Aichholzer, G. Aloupis, E. Demaine, M. Demaine, V. Dujmovic, F. Hurtado, A. Lubiw, G. Rote, A. Schulz, D. Souvaine, and A. Winslow, *Convexifying polygons without losing visibilities*, Proceedings of the 23^{rd} Canadian Conference on Computational Geometry, 2011.
2. O. Aichholzer, M. Cetina, R. Fabila-Monroy, J. Leanos, G. Salazar, and J. Urrutia, *Convexifying monotone polygons while maintaining internal visibility*, Encuentros de Geometria Computacional XIV, 2011, pp. 35–38.
3. B. Braun and R. Ehrenborg, *The complex of noncrossing diagonals of a polygon*, Journal of Combinatorial Theory Series A **117** (2010), 642–649.
4. M. Carr and S. Devadoss, *Coxeter complexes and graph associahedra*, Topology and its Applications **153** (2006), 2155–2168.
5. F. Chapoton, S. Fomin, and A. Zelevinsky, *Polytopal realizations of generalized associahedra*, Canadian Mathematical Bulletin **45** (2002), 537–566.
6. R. Connelly, E. Demaine, and G. Rote, *Straightening polygonal arcs and convexifying polygonal cycles*, Discrete and Computational Geometry **30** (2003), 205–239.
7. S. Devadoss, *Tessellations of moduli spaces and the mosaic operad*, Contemporary Mathematics **239** (1999), 91–114.
8. ———, *A realization of graph associahedra*, Discrete Mathematics **309** (2009), 271–276.
9. S. Devadoss and J. O'Rourke, *Discrete and Computational Geometry*, Princeton University Press, 2011.
10. P. Edelman and V. Reiner, *Visibility complexes and the Baues problem for triangulations in the plane*, Discrete and Computational Geometry **20** (1998), 35–59.
11. I. Gelfand, M. Kapranov, and A. Zelevinsky, *Discriminants, resultants, and multidimensional determinants*, Birkhäuser, 1994.
12. C. Lee, *The associahedron and triangulations of the n-gon*, European Journal of Combinatorics **10** (1989), 551–560.
13. S. Linusson and J. Shareshian, *Complexes of t-colorable graphs*, SIAM Journal of Discrete Mathematics **16** (2003), 371–389.
14. J.-L. Loday, *Realization of the Stasheff polytope*, Archiv der Mathematik **83** (2004), 267–278.
15. D. Orden and F. Santos, *The polytope of noncrossing graphs on a planar point set*, Discrete and Computational Geometry **33** (2005), 275–305.
16. A. Postnikov, *Permutohedra, associahedra, and beyond*, International Mathematical Research Notices **6** (2009), 1026–1106.
17. V. Reiner and G. Ziegler, *Coxeter-associahedra*, Mathematika **41** (1994), 364–393.
18. R. P. Stanley, *Enumerative Combinatorics*, vol. 2, Cambridge University Press, 1999.
19. J. Stasheff, *Homotopy associativity of H-spaces I*, Transactions of the AMS **108** (1963), 275–292.

Williams College, Williamstown, MA 01267
E-mail address: satyan.devadoss@williams.edu

University of California, Santa Barbara, CA 93106
E-mail address: rshah@math.ucsb.edu

Stanford University, Stanford, CA 94305
E-mail address: xhao@math.stanford.edu

Bard College, Annandale-on-Hudson, NY 12504
E-mail address: ezrawinston@gmail.com