A POLYTOPAL GENERALIZATION OF SPERNER'S LEMMA

ELISHA PETERSON AND FRANCIS EDWARD SU

version Mar 15 2001

Abstract. We prove a generalization of Sperner's Lemma to simplicial polytopes and polytopes with simplicial ridges. We give two proofs for simplicial polytopes. The first proof is non-constructive, relying on a degree argument and an interesting observation about how to distributing $(n-k)$ points in a polytope in such a way that any simplex contains at most one such point. The second proof is constructive, and relies on a "path-following" argument. It can be extended to handle polytopes with simplicial ridges.

1. Introduction

Sperner's Lemma is a combinatorial statement about labellings of triangulated simplices whose claim to fame is its equivalence with the topological fixed point theorem of Brouwer. In this paper we address a generalization of Sperner's Lemma that was conjectured by Atanassov in 1996.

Consider a polytope in $\mathbb{R}^k$ defined by

$$P = \operatorname{conv}(\mathbf{v}_1, ..., \mathbf{v}_n),$$

the convex hull of n vertices $\mathbf{v}_1, \ldots, \mathbf{v}_n \in \mathbb{R}^k$ such that no vertex is in the interior of P. This defines a bounded, convex region P we shall call an (n, k)*-polytope.* If $n = k + 1$, P is called a *simplex.* A hyperplane is said to *support* P if one of its closed half-spaces contains P. A subset F of P is a *face* of P if F is the intersection of P with a supporting hyperplane. (Thus the empty set is also a face of P.) Faces of dimension $(k-1)$ and $(k-2)$ are called *facets* and *ridges*, respectively. A *triangulation* T of the polytope P is a finite collection of distinct simplices (called *elementary simplices*) such that: (i) the union of the simplices of T is P, (ii) every face of a simplex in T is in T, and (iii) any two simplices in T intersect in a face common to both. We call the vertices $\mathbf{v}_1, \ldots, \mathbf{v}_n$ of P *main vertices* to distinguish them from the vertices of the simplices in any triangulation.

Let T be a triangulation of P, and suppose that the vertices of T have a labelling satisfying these conditions: each main vertex of P is assigned a unique label from the set $L = \{1, 2, \ldots, n\}$, and each other vertex v is assigned a label of one of the main vertices spanning a minimal face that contains v. Such a labelling is called a *Sperner labelling* of T. An elementary k-simplex in the triangulation is called a *fully-labeled simplex* or simply a *full cell* if all its labels are distinct. The following result was proved by Sperner [14] in 1928:

Sperner's Lemma. *Any triangulation of a k-simplex with a Sperner labelling must contain odd number of full cells. (In particular, there is at least one).*

.

The equivalence of this lemma with the Brouwer theorem was established in 1929 [10]. Constructive proofs of Sperner's lemma [2, 11, 12] emerged in the 1960's, and these were used to develop constructive methods for locating fixed points [16, 18]. More recently, Sperner's lemma has been used to solve *fair division* problems in game theory [15].

Generalizations of Sperner's lemma have proceeded in several directions. [3, 4, 7] In 1996 Atanassov [1] conjectured a generalization of Sperner's lemma for polytopes:

Conjecture (Atanassov)**.** *Let T be a triangulation of a (n,k)-polytope P with a Sperner labelling. Then T contains at least $(n-k)$ full cells.*

Atanassov proved the conjecture for the case where $k=2$. Note that Sperner's lemma is the case $n=k+1$.

In this paper we prove this conjecture for all (n,k)-polytopes with simplicial ridges (Theorem 23). Such polytopes include the class of *simplicial polytopes* (polytopes whose facets are simplicial) as well as all polytopes of dimension 3 and lower.

Just as the original lemma of Sperner had non-constructive and constructive proofs, we provide a non-constructive and constructive proof of this conjecture for simplicial polytopes. Section 2 uses a degree argument and a notion of *spanning pebbles* to provide a non-constructive proof for simplicial polytopes. Section 3 develops background on path-following arguments in polytopes that is closely related to classical path-following arguments for Sperner's lemma [2, 11, 16]. This is applied in Section 4 to give a constructive proof for simplicial polytopes and in Section 5 to prove the conjecture for polytopes with simplicial ridges. We end with some applications and further questions in Section 6.

2. A Non-Constructive Proof for Simplicial Polytopes

The non-constructive proof that we give in this section is an extension of "degree" arguments for proving the usual Sperner lemma.

Let P be a simplicial (n,k)-polytope with triangulation T. Since P is simplicial, the main vertices $\mathbf{v}_1,\ldots,\mathbf{v}_n$ of P are in general position. Consider the piecewise linear (PL) map $f: P \to P$ that maps each vertex of T to the main vertex of P that shares the same label, and is linear on each elementary simplex.

Let ∂P denote the boundary of P (i.e., the union of the facets). Because of the Sperner labelling, the map $\partial f := f|_{\partial P}$ has degree 1. Hence $\partial f_*: H_{n-1}(\partial P) \to H_{n-1}(\partial P)$ is the identity map on homology groups. Thus the map f induces the following commutative diagram of homology sequences:

$$\begin{array}{ccccccc}
H_n(P)\cong 0 & \longrightarrow & H_n(P,\partial P) & \xrightarrow{\partial} & H_{n-1}(\partial P) & \longrightarrow & H_{n-1}(P)\cong 0 \\
 & & \downarrow{\scriptstyle f_*} & & \downarrow{\scriptstyle \partial f_* = id} & & \\
H_n(P)\cong 0 & \longrightarrow & H_n(P,\partial P) & \xrightarrow{\partial} & H_{n-1}(\partial P) & \longrightarrow & H_{n-1}(P)\cong 0
\end{array}$$

where the rows are exact. Since $H_i(P)=0$, the maps ∂ are isomorphisms, which implies that f_* is the identity map. Hence f has degree 1.

Therefore the number of pre-images of a point in the image of f is 1 (when counted with sign) as long as the map f is non-degenerate at that point. Thus if we can find a set of points in P such that any k-simplex spanned by $(k+1)$ main vertices of P contains at most one such point in its interior, then the pre-image of each such point will correspond to a full cell in P.

Definition 1. A set of *spanning pebbles* of a (n, k)-polytope P is a set of $(n - k)$ points (*pebbles*) such that each k-simplex spanned by $(k + 1)$ main vertices of P contains at most one pebble.

It is not clear that a set of $(n - k)$ spanning pebbles should exist for a given (n, k)-polytope. In fact it is easy to construct examples for which they do not exist. However, we show that they do exist for simplicial (n, k)-polytopes:

Theorem 2. *Every simplicial (n, k)-polytope contains a set of $(n - k)$ spanning pebbles.*

Proof. As above, let $\mathbf{v}_1, ..., \mathbf{v}_n \in \mathbb{R}^k$ be the vertices of the given (n, k)-polytope. Set

$$\mathbf{v}_0 = \frac{1}{k}(\mathbf{v}_{n-k+1} + \cdots + \mathbf{v}_n). \tag{1}$$

Let diam(P) denote the diameter of the polytope P, i.e., the maximum pairwise distance between any two points in P. Let H be the minimum distance between any vertex $\mathbf{v} \in V$ and any hyperplane spanned vertices in $V \setminus \{\mathbf{v}\}$. Since there at most finitely many such distances and the vertices are in general position, H exists and is positive. Set

$$\varepsilon = \frac{H}{2\,\mathrm{diam}(P)}.$$

Define

$$\mathbf{p}_i = \varepsilon \mathbf{v}_0 + (1 - \varepsilon)\mathbf{v}_i \tag{2}$$

for $1 \le i \le n - k$. We claim that these points $\mathbf{p}_i$ are spanning pebbles.

Lemma 3. *Each $\mathbf{p}_i$ has the property that $\mathbf{p}_i \notin$ conv$(\mathbf{v}_1, ..., \hat{\mathbf{v}}_i, ..., \mathbf{v}_n)$. (Here $\hat{\mathbf{v}}_i$ means that $\mathbf{v}_i$ is omitted.)*

Proof. This follows because

$$\frac{\|\mathbf{v}_i - \mathbf{p}_i\|}{\|\mathbf{v}_i - \mathbf{v}_0\|} = \varepsilon = \frac{H}{2\,\mathrm{diam}(P)} \le \frac{H}{2\|\mathbf{v}_i - \mathbf{v}_0\|},$$

and this implies that $\|\mathbf{v}_i - \mathbf{p}_i\| < H/2$. Hence $\mathbf{p}_i$ is closer to $\mathbf{v}_i$ than $\mathbf{v}_i$ is to the nearest hyperplane spanned by vertices in $V \setminus \{\mathbf{v}_i\}$. Since $\mathbf{v}_i$ not in the convex hull of vertices in $V \setminus \{\mathbf{v}_i\}$, neither is $\mathbf{p}_i$. □

Now consider any k-simplex S spanned by $\mathbf{s}_1, ..., \mathbf{s}_{k+1} \in V$. We claim that S cannot contain more than one pebble.

Suppose by way of contradiction that it did. Say $\mathbf{p}_{i'}$ and $\mathbf{p}_{j'}$ are contained in S for distinct i', j', $1 \le i', j' \le k$. Lemma 3 implies that $\mathbf{v}_{i'}, \mathbf{v}_{j'}$ must both be vertices of S. Without loss of generality, let $\mathbf{s}_1 = \mathbf{p}_{i'}$ and $\mathbf{s}_2 = \mathbf{p}_{j'}$. Let A be a matrix filled with columns consisting of $\mathbf{v}_0$ and the vertices of S, adjoined with a row of 1's:

$$A = \begin{bmatrix} \mathbf{s}_1 & \mathbf{s}_2 & \cdots & \mathbf{s}_{k+1} & \mathbf{v}_0 \\ 1 & 1 & \cdots & 1 & 1 \end{bmatrix}.$$

This is a $(k + 1) \times (k + 2)$ matrix that has rank $(k + 1)$ because the $\mathbf{s}_i$ are affinely independent. So ker A is 1-dimensional. Note that $\mathbf{p}_{i'} \in S$ implies that it is a convex combination of the first $(k+1)$ columns of B. Thus there exist constants $0 \leq x_1, x_2, ..., x_{k+1} \leq 1$ satisfying

$$\begin{bmatrix} \mathbf{p}_{i'} \\ \\ 1 \end{bmatrix} = A \begin{bmatrix} 1-\varepsilon \\ 0 \\ \vdots \\ 0 \\ \varepsilon \end{bmatrix} = A \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_{k+1} \\ 0 \end{bmatrix} \tag{3}$$

where the first equality follows from (2). Similarly, $\mathbf{p}_{j'} \in S$ implies

$$\begin{bmatrix} \mathbf{p}_{j'} \\ \\ 1 \end{bmatrix} = A \begin{bmatrix} 0 \\ 1-\varepsilon \\ \vdots \\ 0 \\ \varepsilon \end{bmatrix} = A \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_{k+1} \\ 0 \end{bmatrix}$$

for some constants $0 \leq y_1, y_2, ..., y_{k+1} \leq 1$.

The above equations show that $(x_1+\varepsilon-1, x_2, x_3 \ldots, x_{k+1}, -\varepsilon)^T$ and $(y_1, y_2+\varepsilon-1, y_3, \ldots, y_{k+1}, -\varepsilon)^T$ are both in $\ker A$. But since $\ker A$ is 1-dimensional, and the last coordinates of these vectors are equal, all entries of these vectors are identical. In particular, $x_1 + \varepsilon - 1 = y_1$.

We now claim that $x_1 + \varepsilon < 1$ and hence $y_1 < 0$, which would show that $\mathbf{p}_{j'}$ could not have been in S after all, a contradiction.

To establish the claim, use equations (2) and (3) to express $\mathbf{v}_0$ as an affine combination of the vertices of S:

$$\mathbf{v}_0 = \frac{1}{\varepsilon}((x_1 + \varepsilon - 1)\mathbf{s}_1 + x_2\mathbf{s}_2 + \ldots x_{k+1}\mathbf{s}_{k+1}). \tag{4}$$

Since $x_2, \ldots, x_{k+1} \geq 0$ and $\mathbf{v}_0$ is not in the interior of S, it must be the case that $x_1 + \varepsilon \leq 1$. If $x_1 + \varepsilon = 1$, then by (4), $\mathbf{v}_0$ is on a facet of S. In light of (1) and the fact that the the vertices of the polytope P are in general position, the vertices of S must be $\{\mathbf{v}_{n-k+1}, \ldots, \mathbf{v}_n, \mathbf{v}_{i'}\}$. By assumption $\mathbf{v}_{j'}$ must be distinct from every member of this set, although Lemma 3 implies that it must be a vertex of S since $\mathbf{p}_{j'}$ is in the convex hull of S. Hence equality cannot hold, and $x_1 + \varepsilon < 1$ as desired. $\square$

We remark that in choosing $\mathbf{v}_0$ in equation (1) we could have used any point which is a strict convex combination of $\mathbf{v}_{n-k+1}, \ldots, \mathbf{v}_n$.

We note that since a set of $(n-k)$ spanning pebbles do not necessarily exist for every polytope, this proof technique does not generalize. In the next section, we give a proof for simplicial polytopes which does allow for some generalization.

3. Path-following in Polytopes

The usual Sperner lemma had a constructive proof which relied on a "path-following" argument, developed in the late 1960's by [2, 11, 12] and others. See [15] for a simple version of this proof technique. The path-following argument relied on tracing a path through simplices in the triangulation based on the labels encountered. In this section we adapt these ideas for Sperner-labelled polytopes, which are used in the subsequent sections to provide proofs for simplicial polytopes and polytopes with simplicial ridges.

Fix a label set $L = \{1, 2, \ldots, n\}$ used in the Sperner labelling of a triangulation of the simplicial (n, k)-polytope P.

Definition 4. Let $L(\sigma)$, the *label set* of σ, denote the set of distinct labels of vertices of σ.

Definition 5. A k-simplex $\sigma \in T$ is a *full cell* if the vertex labels of σ are all distinct. Similarly, a $(k-1)$-simplex $\tau \in T$ is a *full facet* if the vertex labels of τ are all distinct.

Definition 6. Let K be a subset of the label set $L = \{1, 2, \ldots, n\}$. The *derived graph* G_K is a graph with nodes that are simplices of T whose label set contains K and is at least size k. Formally, σ is a node of G_K if $|L(\sigma)| \geq k$ and $K \subset L(\sigma)$. Two nodes of G_K are adjacent in G_K if (as simplices) one is a face of the other.

Thus the nodes of G_K have label sets containing K and are either full cells, full facets, or k-simplices with exactly one repeated label. We remark that Freund and Todd [8] use a similar graph in their proof of a combinatorial lemma of Tucker.

For our purposes we shall be concerned only with situations where K is of size k, $(k-1)$ or $K = \emptyset$ (the empty set). For simplicity, we let G denote the graph $G_\emptyset$. It is easy to see that for any K, G_K is a subgraph of G.

The following lemma shows that for K of size $(k-1)$, G_K is a graph whose connected components are either loops or paths that connect pairs of full facets on the boundary of P.

Lemma 7. *The nodes of the graph G_K are either of degree 1 or 2 for any K of size $k-1$. A node σ is of degree 1 if and only if σ is a full facet on the boundary of P.*

Proof. Recall that each node σ of G_K has a label set containing K and is either a full cell, full facet, or k-simplex with exactly one repeated label.

If σ is a full cell and if $|K| = k-1$, then $L(\sigma)$ consists of labels in K and two other labels l_1, l_2. There are exactly two facets of σ whose label sets contain K; these are the full facets with label sets $K \cup l_1$ and $K \cup l_1$, respectively. Thus σ has degree 2.

If σ is a full facet with label set containing K, then it is the face of exactly two k-simplices, unless σ is on the boundary of P, in which case it is the face of exactly one k-simplex. Thus σ is degree 1 or 2 in G_K, and degree 1 when σ is a full facet on the boundary of P.

If σ is a k-simplex with exactly one repeated label, then it must possess exactly two full facets. Since $K \subset L(\sigma)$, these full facets must also have label sets that contain K. Hence these two full facets are the neighbors of σ in G_K, so σ has degree 2. $\square$

The next lemma shows that for K of size k, G_K is a graph whose connected components are either loops or paths. Such paths have endpoints that are either full facets on the boundary of P or full cells in the interior.

Lemma 8. *The nodes of the graph G_K are either of degree 1 or 2 for any K of size k. A node σ is of degree 1 if and only if σ is a full cell or full facet on the boundary of P.*

Proof. Recall that each node σ of G_K has a label set containing K and is either a full cell, full facet, or k-simplex with exactly one repeated label. The arguments for the latter two cases are identical to those in the proof of Lemma 7.

If σ is a full cell and if $|K| = k$, then σ has exactly one facet with label set K, so it has degree 1 in G_K. $\square$

Lemma 9. *The nodes of the graph G are of degree 1, 2, or $k+1$. A node σ is of degree 1 if and only if σ is a full facet on the boundary of P. A node σ is of degree $k+1$ if and only if σ is a full cell.*

Proof. As noted before, each node σ of G is either a full cell, full facet, or k-simplex with exactly one repeated label. The arguments for the latter two cases are identical to those in the proof of Lemma 7 by setting $K = \emptyset$.

If σ is a full cell, then every facet of σ is a full facet, hence the degree of σ is $(k+1)$ in G. $\square$

Lemma 10. *If σ is adjacent to τ in G, then $L(\sigma) \subset L(\tau)$, unless σ is a full cell, in which case $L(\tau) \subset L(\sigma)$.*

Proof. Suppose σ is a k-simplex with exactly one repeated label. Then it is adjacent to two full facets with exactly the same label set, so the conclusion holds.

Otherwise, if σ is a full facet, then it is adjacent to two k-simplices that contain it as a facet. Hence $L(\sigma) \subseteq L(\tau)$ for τ adjacent to σ.

Finally, if σ is a full cell, any simplex τ adjacent to σ in G is contained in σ as a facet, so $L(\tau) \subset L(\sigma)$ in that case. $\square$

Lemma 11. *Suppose G is connected. If G contains at least one full cell as a node, then all the labels in G are carried by the full cells.*

Proof. Since G is connected and contains at least one full cell, each simplex σ that is not a full cell is connected to a full cell τ via a path in G that does not intersect any other full cell in G. Call this path $\{\sigma = \sigma_1, \sigma_2, ..., \sigma_p = \tau\}$,

Lemma 10 shows that adjacent nodes in G carry exactly the same labels unless one of them is a full cell. Hence $L(\sigma_1) = L(\sigma_2) = ... = L(\sigma_{p-1}) \subset L(\tau)$. Therefore labels carried by the full cells contain all labels carried by any other node of the graph. $\square$

Definition 12. Let G' denote the *full cell graph*, whose nodes are full cells in G. Two full cells σ, τ are adjacent in G' if there exists a path from σ to τ in G that does not intersect any other full cell.

Lemma 13. *Any two adjacent nodes in G' are full cells in T whose label sets contain at least k labels in common.*

Proof. Let σ_1 and σ_2 be adjacent nodes in G'. By construction they must simplices connected by a path in G; let τ be any such node along this path. By Lemma 10, $L(\tau) \subset L(\sigma_1)$ and $L(\tau) \subset L(\sigma_2)$, so $L(\sigma_1) \cap L(\sigma_2)$ contains at least the k labels in $L(\tau)$. $\square$

Proposition 14. *Suppose G' is connected. Let $L(G')$ denote the set of all labels carried by simplices in G'. Then*

$$|G'| \geq |L(G')| - k.$$

Proof. We induct on $|G'|$. If $|G'| = 1$, the one full cell in G' has $k+1$ labels. Hence $|L(G')| - k = (k+1) - k = 1$, so the statement holds.

We now assume the statement holds for full cell graphs with less than j nodes, and show it holds for full cell graphs G' with $|G'| = j$.

Assume G' has j full cells. We claim that it is possible to remove a vertex v from G' and leave G' connected. This is true because G' contains a maximal spanning tree, and the removal of any leaf from this tree will leave the rest of the nodes in G' connected by a path in this tree.

Now G' with v and all its incident edges removed is a new graph (denoted by $G' - v$) with $j - 1$ nodes. By the inductive hypothesis, $|G' - v| \geq |L(G' - v)| - k$.

Clearly $|G'| - 1 = |G' - v|$, and $|L(G' - v)| \geq L(G') - 1$ because v has at least k labels in common with some vertex in $G' - v$, by Lemma 13. Hence $|G'| - 1 \geq L(G') - 1 - k$. Adding 1 to both sides, gives the desired conclusion. $\square$

This is useful in proving the following.

Theorem 15. *Let T be a Sperner-labelled triangulation of an (n, k)-polytope P. If the derived graph G has a component $\bar{G}$ that carries all the labels of G, then T contains at least $(n - k)$ full cells.*

Proof. Use $\bar{G}$ to construct the full cell graph $\bar{G}'$ as above. Note that if $\bar{G}$ is connected then $\bar{G}'$ is also connected. By Lemma 11, $L(\bar{G}) = L(\bar{G}')$. Using Proposition 14, we have

$$|\bar{G}'| \geq |L(\bar{G})| - k = n - k \ ,$$

which shows there are at least $(n - k)$ full cells in $\bar{G}$, and hence in G itself. $\square$

Thus to prove our conjecture for a given (n, k)-polytope it suffices to find some component $\bar{G}$ of G for which $L(\bar{G}) = n$. This is the central idea of the proofs in the next two sections.

4. Simplicial Polytopes

We now use path-following ideas to provide a proof of Atanassov's conjecture for simplicial facets. It will be constructive in the sense described at the end of this section.

Theorem 16. *If P is a simplicial polytope, there is some component $\bar{G}$ of G which visits every main facet of P, and hence carries all labels of G.*

Proof. Let F be a main facet of the polytope P. Let $\chi(G, F)$ count the number of nodes of G that are simplices on the facet F. This may be thought of as the number of endpoints of paths in G that terminate on the facet F.

Consider two "adjacent" main facets F_1, F_2 of P, whose intersection is a ridge of the polytope P spanned by $(k - 1)$ main vertices of P. Let K be the label set of these main vertices; by the Sperner labeling the elements of K are all distinct. Consider the derived graph G_K, and note that endpoints of paths in G_K must be full facets in F_1 or F_2, since no other main facet of P has a label set containing K (because of the Sperner labelling).

Since every facet of P is simplicial, all the full facets in F_1 and F_2 contain K in their label set. Thus all the nodes of G that are full facets in F_1 and F_2 must also be nodes in the graph G_K. Since G_K is a subgraph of G and consists of paths with endpoints that pair up full facets in F_1 and F_2, we see that

$$\chi(G, F_1) \equiv \chi(G, F_2) \mod 2.$$

In fact, the same argument shows $\chi(\bar{G}, F_1) \equiv \chi(\bar{G}, F_2) \mod 2$ for any connected component $\bar{G}$ of G.

Since F_1 and F_2 were arbitrary, the same argument holds for any two adjacent facets. This yields the somewhat surprising conclusion that the parity of $\chi(\bar{G}, F_i)$ is independent of the facet F_i. We denote this parity by $\rho(\bar{G})$. Clearly $\chi(G, F_i)$ is also independent of facet and we can define $\rho(G)$ similarly.

Since $\chi(G, F_i)$ is the sum of $\chi(\bar{G}, F_i)$ over all connected components $\bar{G}$ of G, it follows that $\rho(G) \equiv \sum \rho(\bar{G}) \mod 2$ over all connected components $\bar{G}$ of G. Moreover, $\rho(G) \equiv 1 \mod 2$ because of the usual Sperner's lemma applied to (any) simplicial facet F.

Hence there must be some $\bar{G}$ such that $\rho(\bar{G}) \equiv 1 \mod 2$, i.e., this $\bar{G}$ visits every main facet of P. Because the facets of P are simplicial, this means that $|L(\bar{G})| = n$. □

Theorem 17. *Any Sperner-labelled triangulation of a simplicial (n, k)-polytope must contain at least $(n - k)$ full cells.*

Proof. This follows immediately from Theorem 16 and Theorem 15. □

Although the above argument in Theorem 16 appears non-constructive, the derived graph G provides a constructive path-following procedure for locating a full cell in the triangulation that contains the label set of any particular facet F. The procedure relies on the fact that there are constructive methods for locating all the full cells in a Sperner-labelled simplex.

The procedure begins by starting from the main facet F and using a constructive proof of Sperner's lemma to locate a full facet on F. Starting from that full facet, trace a path in the derived graph G. If it does not eventually hit a full cell in G, then it must terminate in a full facet in F. Such paths can only pair up an even number of full facets in F; since there are an odd number of such facets, at least one full facet in F has a path that leads to a full cell in G.

Finding a set of $(n-k)$ full cells in G can be done in a somewhat less constructive fashion. We know from the above discussion that some component $\bar{G}$ of G carries all labels and hits every facet F. Thus by following paths from full cells in F, one can find such a component. Since this component can have a quite complicated tree structure, one might have to exhaust the paths in this tree to find at least $(n - k)$ full cells; (this could be quite inefficient).

5. Polytopes with Simplicial Ridges

In this section, we extend the proof of the last section to polytopes whose ridges are simplicial. Such polytopes include the class of polytopes with simplicial facets as well as all polytopes of dimension 3 and lower.

Throughout this section we let the symbol $\equiv$ denote equivalence mod 2.

We first extend the definition of χ from the last section.

Definition 18. Let $\chi(G, F, L)$ count the number of nodes of the graph G ending on the facet F that contain the label set L. Recall that $L(F)$ denotes the label set of F.

Theorem 19. *Let R be a ridge of P, let F be a facet of P, and suppose that R is a simplicial facet of F. Then*

$$\chi(G, F, L(R)) \equiv 1.$$

Proof. Recall that $\chi(G, F, L(R))$ counts the number of full facets in F that contain the label set $L(R)$. Since the facet F is itself a polytope, it has a derived graph G^F, which Lemma 8 shows are paths connecting fully-labelled simplices in R and fully-labelled simplices in F containing the label set $L(R)$. Thus

$$\chi(G, F, L(R)) \equiv \chi(G^F, R, L(R)).$$

But since R is simplicial, the usual Sperner's lemma applied to R shows that $\chi(G^F, R, L(R)) \equiv 1$. □

We now show that $\chi(G, F, L(R))$ does not depend on either F or R, but only on G.

Theorem 20. *Let R, R' be ridges of P, let F be a facet of P, and suppose that R and R' are simplicial facets of F. Then*

$$\chi(G, F, L(R)) \equiv \chi(G, F, L(R')).$$

Proof. Suppose first that R and R' are adjacent and share a common facet C (it has dimension $k-3$, if you are keeping track). We claim that

$$\sum_x \chi(G, F, L(C) \cup x) \equiv 0,$$

where x runs over all labels in $L(F)$ that are not in C. This follows because the sum counts only fully-labelled simplices σ in F that contain $L(C)$, and every such σ appears exactly *twice* in this sum; if $L(\sigma) = L(C) \cup a \cup b$, then σ is counted once each in $\chi(G, F, L(C) \cup a)$ and $\chi(G, F, L(C) \cup b)$.

On the other hand, if $L^* = L(C) \cup x$ is not the label set of any ridge of P, then any fully-labelled simplex in on the boundary of P that contains L^* must be contained in the facet F. Since L^* is of size $k-1$, by Lemma 7, we see that $\chi(G, F, L^*) \equiv 0$.

Thus the only terms surviving the above sum correspond to label sets of the two ridges R, R' that are facets of F and share a common face C, i.e.,

$$\chi(G, F, L(R)) + \chi(G, F, L(R')) \equiv 0,$$

which yields the desired conclusion for neighboring ridges R, R'. Since any two ridges of a facet F are connected by a chain of adjacent ridges, the general conclusion holds. □

Theorem 21. *Let R be ridges of P, let F, F' be a facet of P, and suppose that R is a common facet of both F, F'. Then*

$$\chi(G, F, L(R)) \equiv \chi(G, F', L(R)).$$

Proof. Since $L(R)$ is a label set of size $(k-1)$, consider the derived graph $G_{L(R)}$. By Lemma 7, this graph consists of paths connecting fully-labelled simplices on the boundary of P that contain the label set $L(R)$. Hence the number of such simplices is even.

On the other hand, because of the Sperner labelling of P, all such simplices must lie in facets of P that contain R. There are exactly two such facets, F and F'. Hence

$$\chi(G, F, L(R)) + \chi(G, F', L(R)) \equiv 0,$$

which gives the desired conclusion. □

Since all facets are connected by a sequence of adjacent facets, the two theorems above show that the parity of $\chi(G, F, L(R))$ is independent of both F and R. Hence we may define the *parity* of G to be $\rho(G) \equiv \chi(G, F, L(R))$.

Now we may prove

Theorem 22. *If P is a polytope with simplicial ridges, there is some component $\bar{G}$ of G which carries all labels of P.*

Proof. Fix some facet F and ridge R of P, such that R is a facet of F. Since $\chi(G, F, L(R))$ is the sum of $\chi(\bar{G}, F, L(R))$ over all connected components $\bar{G}$ of G, it follows that $\rho(G) \equiv \sum \rho(\bar{G})$ over all connected components $\bar{G}$ of G. Moreover, Theorem 19 shows that $\rho(G) \equiv 1$.

Hence there must be some $\bar{G}$ such that $\rho(\bar{G}) \equiv 1$, i.e., this $\bar{G}$ carries the labels of R. Since R and F were arbitrary, and $\rho(\bar{G})$ is independent of R and F, $\bar{G}$ must carry all labels of P. □

This implies

Theorem 23. *Any Sperner-labelled triangulation of an (n, k)-polytope with simplicial ridges must contain at least $(n - k)$ full cells.*

Proof. This follows immediately from Theorem 22 and Theorem 15. □

6. CONCLUSION

Insert some possible applications; discuss related applications in game theory and fair division [15, 13].

Pose the general Atanassov conjecture as an open question.

REFERENCES

1. Atanassov, K. T., "On Sperner's Lemma," *Stud. Sci. Math. Hungarica* **32** (1996), 71-74.
2. Cohen, D. I. A., "On the Sperner Lemma," *J. Comb. Theory* **2** (1967), 585-587.
3. Eaves, C., G. van der Laan, D. Talman, and Z. Yang, "Balanced Simplices on Polytopes," preprint.
4. Ehrenborg, R., and G. Hetyei, "Generalizations of Baxter's Theorem and Cubical Homology," *J. Comb. Theory, Series A* **69** (1995), 233-287.
5. Fan, K., "Combinatorial Properties of Certain Simplicial and Cubical Vertex Maps," *Arch. Math.* **11** (1960), 368-377.
6. Freund, R. M., "Variable Dimension Complexes," *Math. Oper. Research* **9** (1984), 479-509.
7. Freund, R. M., "Combinatorial Analogs of Brouwer's Fixed-Point Theorem on a Bounded Polyhedron," *J. Comb. Theory, Series B* **47** (1989), 192-219.
8. Freund, R.M., Todd, M.J., 1981. A constructive proof of Tucker's combinatorial lemma, *J. Comb. Theory Series A* **30**, 321-325.
9. Herings, J. J. and D. Talman, "Intersection Theorems with a Continuum of Intersection Points," *J. Optimization Theory and Apps.* **2** (1998), 311-335.
10. B. Knaster, C. Kuratowski, and S. Mazurkiewicz, Ein Beweis des Fixpunktsatzes für n-dimensionale Simplexe, *Fund. Math.* **14** (1929) 132-137.
11. Kuhn, H. W., "Simplicial Approximation of Fixed Points," *Proc. Natl. Acad. Sci.* **61** (1968), 1238-1242.

12. Scarf, H., "The Approximation of Fixed Points of a Continuous Mapping," *SIAM J. Appl. Math.* **15** (1967), 1328-1343.
13. Simmons, F. W. and F. E. Su, "Consensus-halving via a theorems of Borsuk-Ulam and Tucker," preprint.
14. Sperner, E., Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes, *Abh. Math. Sem. Hamburg* **6** (1928), 265-272.
15. Su, F. E., "Rental Harmony: Sperner's Lemma in Fair Division," *Amer. Math. Monthly* **106** (1999), 930-942.
16. Todd, M. J., *The Computation of Fixed Points and Applications*, Lecture Notes in Econ. and Math. Syst., Springer-Verlang, New York (1976).
17. Tucker, A. W., "Some Topological Properties of Disk and Sphere," in *Proceedings of the first Canadian Mathematical Congress*, University of Toronto Press, Toronto (1945) 285-309.
18. Yang, Z., *Computing Equilibria and Fixed Points*, Kluwer Academic Publishers, Boston (1999).
19. Ziegler, G. M., *Lectures on Polytopes*, Springer-Verlag, New York (1995).