A POLYTOPAL GENERALIZATION OF SPERNER'S LEMMA

JESUS A. DE LOERA, ELISHA PETERSON, AND FRANCIS EDWARD SU

version July 2 2001

Abstract. We prove the following conjecture of K.T. Atanassov [2]:

Let T be a triangulation of a d-dimensional polytope P with n vertices $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$. Label the vertices of T by $1, 2, \ldots, n$ in such a way that a vertex of T belonging to the interior of a face F of P can only be labeled by j if $\mathbf{v}_j$ is on F. Then there are at least $(n-d)$ full dimensional simplices of T, each labeled with $d+1$ different labels. We provide two proofs of this result: a non-constructive proof introducing the notion of a *pebble set* of a polytope, and a constructive proof using a path-following argument. We also show a new proof of a KKM-type intersection theorem from [10] and indicate interesting relations to minimal simplicial covers of convex polyhedra and their chamber complexes.

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 [8, 16]. In this paper we prove a generalization of Sperner's Lemma that settles a conjecture proposed by K.T. Atanassov [2].

Consider a convex polytope P in $\mathbb{R}^d$ defined by n vertices $\mathbf{v}_1, \ldots, \mathbf{v}_n \in \mathbb{R}^d$. For brevity, we will call such polytope an (n, d)*-polytope.* Throughout the paper we will follow the terminology of the book [20]. By a *triangulation* T of the polytope P we mean a finite collection of distinct 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. The points $\mathbf{v}_1, \ldots, \mathbf{v}_n$ are called *vertices of* P to distinguish them from vertices of T, the triangulation. Similarly, a simplex spanned by vertices of P will be called a *simplex of* P to distinguish it from simplices involving other vertices of T. If S is a subset of P, then the *carrier* of S, denoted $\operatorname{carr}(S)$, is the smallest face F of P that contains S. In that case we say S is *carried* by F. A *cover* C of a convex polytope P is a collection of full dimensional simplices in P such that $\cup_{\sigma \in C} \sigma = P$. The *size of a cover* is the number of simplices in the cover.

.

Let T be a triangulation of P, and suppose that the vertices of T have a labelling satisfying these conditions: each vertex of P is assigned a unique label from the set $\{1, 2, \ldots, n\}$, and each other vertex v of T is assigned a label of one of the vertices of P in $\text{carr}(v)$. Such a labelling is called a *Sperner labelling* of T. A d-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 [16] in 1928:

Sperner's Lemma. *Any Sperner labelling of a triangulation of a d-simplex must contain an odd number of full cells; in particular, there is at least one.*

Constructive proofs of Sperner's lemma [4, 9, 12] emerged in the 1960's, and these were used to develop constructive methods for locating fixed points [18, 19]. Sperner's lemma and its variants continue to be useful in applications. For example, recently they have been used to solve *fair division* problems in game theory [14, 17]. The main purpose of this paper is to present a solution of the following conjecture.

Conjecture (Atanassov)**.** *Any Sperner labelling of a triangulation of an (n, d)-polytope must contain at least $(n - d)$ full cells.*

In 1996, K.T. Atanassov [2] stated the conjecture and gave a proof for the case where $d = 2$. Note that Sperner's lemma is exactly the case $n = d + 1$. In this paper we prove this conjecture for all (n, d)-polytopes. Here is the central result of our paper:

Theorem 1. *Any Sperner labelling of a triangulation T of an (n, d)-polytope P must contain at least $(n - d)$ full cells. Moreover, the collection of full cells in T corresponds to a cover of P under the piecewise linear map that sends each vertex of T to the vertex of P that shares the same label.*

We provide a non-constructive and a constructive proof of Theorem 1. The non-constructive proof in Section 2 is obtained via a degree argument and the notion of a pebble set. Section 3 develops background on path-following arguments in polytopes that is closely related to classical path-following arguments for Sperner's lemma [4, 9, 18]. This is applied to give a constructive proof for simplicial polytopes. In Section 4 we extend the construction to prove the conjecture for arbitrary polytopes. The final section of the paper is devoted to two interesting consequences of Theorem 1 and its proofs. From our first proof we derive the following corollary:

Corollary 2. *Let $c(P)$ denote the* covering number *of an (n, d)-polytope P, which is the size of the smallest cover of P. Then, $c(P) \geq n - d$. This result is best possible as the equality is attained for stacked polytopes.*

We also get a slight strenghtening of Theorem 10 of [10]. We need to recall the notion of chamber complex of a polytope P (see [1]): let Σ be the set of all d-simplices of P. Denote by $\text{bdry}(\sigma)$ the boundary of simplex σ. Consider the set of open polyhedra $P - \cup_{\sigma \in \Sigma} \text{bdry}(\sigma)$. A *chamber* is the closure of one of these components. The chamber complex of P is the polyhedral complex given by all chambers and their faces.

Corollary 3. *Let P be an (n, d)-polytope with vertices $\mathbf{v}_1, \dots, \mathbf{v}_n$. Let $\{C^j \mid j = 1, \dots, n\}$ be a collection of closed sets covering the (n, d)-polytope P, such that each face F is covered by $\cup\{C^h \mid \mathbf{v}^h \in F\}$.*

Then, for each $p \in P$, there exists a subset $J_p \subset \{1, 2, \dots, n\}$ such that (1) p lies in the convex hull of the vertices $\mathbf{v}_j$ with $j \in J_p$, (2) J_p has cardinality $d + 1$, (3) $\cap_{j \in J_p} C^j \neq \emptyset$, and (4) if p and q are interior points of the same chamber of P, then $J_p = J_q$. There are at least $c(P)$, the covering number, different such subsets, and the simplices of P indicated by the labels in these subsets form a cover of P.

Figure 1 illustrates with an example the content of the above corollary:

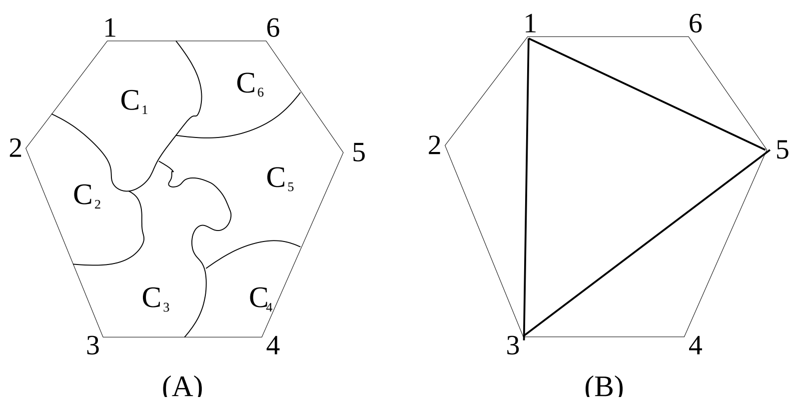

FIGURE 1. Part (A) shows several closed sets covering a hexagon and their four intersection points. The points of intersection correspond to a cover of the hexagon, in this case a triangulation, illustrated in part (B).

2. A Non-Constructive Proof using Pebble Sets

The non-constructive proof of Theorem 1 that we give in this section is an extension of "degree" arguments for proving the usual Sperner's lemma. We first establish a proposition that yields the covering property, then show how the construction a *pebble set* will yield a lower bound for the number of full cells.

Let P be an (n, d)-polytope with Sperner-labelled triangulation T. Consider the piecewise linear (PL) map $f : P \to P$ that maps each vertex of T to the vertex of P that shares the same label, and is linear on each d-simplex of T.

Proposition 4. *The map $f : P \to P$ defined as above is surjective, and thus the collection of full cells in T forms a cover of P under f.*

Before proving this result, we recall a few facts about the *degree* of a map f between manifolds. If $f(x) = y$ and the Jacobian determinant of f exists and is non-zero at x, then x is called a *regular point* of f, and the *sign* of x is the sign (± 1) of this determinant. The point y is called a *regular value* of f if *every* pre-image of y is a regular point. Any regular value y has finitely many pre-images, and the sum of the signs of its pre-images is, in fact, independent of the choice of y and is known as the *degree* of the map f.

The degree is a homotopy invariant of mappings between manifolds (relative to their boundaries) and may also be computed as the multiplicative factor induced by the map f on the corresponding top homology groups (relative to their boundaries). See [5, Ch.1] or [13, Sec.38] for expositions of the topological degree of simplicial maps, or [7] for the general theory.

Thus for the map $f : P \to P$ defined above, the interior points of simplices of T are regular points; interior points of chambers of P are regular values. Observe that the sign of a regular point x depends essentially on the orientation of the labels of the simplex of T that contains x.

Proof. We shall show, by induction on the dimension d of P, that f has degree 1.

If $d = 0$, then P is a point and the statement is clearly true. So assume that the above statement holds for all polytopes of dimension less than d.

Given a d-dimensional polytope P, let ∂P denote the boundary of P (i.e., the union of the facets). If $d = 1$ then ∂P consists of two points on which the map $\partial f := f|_{\partial P}$ is the identity (due to the Sperner labelling), and therefore $\partial f_* : H_{d-1}(\partial P) \to H_{d-1}(\partial P)$ is the identity map on (reduced) homology groups. For $d > 1$, we use the Sperner labelling on the facets and the inductive hypothesis to show that the map ∂f_* is the identity map. Specifically, if F is any facet of P, let $A = \partial P \setminus F$. Then the map ∂f induces the following commutative diagram of (reduced) homology groups:

$$\begin{array}{ccccccc} H_{d-1}(A) & \to & H_{d-1}(\partial P) & \xrightarrow{\pi_*} & H_{d-1}(\partial P, A) & \to & H_{d-2}(A) \\ & & \downarrow \partial f_* & & \downarrow f_* & & \\ H_{d-1}(A) & \to & H_{d-1}(\partial P) & \xrightarrow{\pi_*} & H_{d-1}(\partial P, A) & \to & H_{d-2}(A) \end{array}$$

where the rows are exact. Since $H_i(F) \cong 0$ for $i \in \{d-1, d-2\}$, the maps π_* are isomorphisms. By excision, $H_{d-1}(\partial P, A) \cong H_{d-1}(F, \partial F)$, and this isomorphism reveals f_* to be induced by f on the facet F (a polytope of one lower dimension) which by the inductive hypothesis has degree 1 and must therefore be the identity. Since the maps π_* are isomorphisms and f_* is the identity map, $\partial f_* : H_{d-1}(\partial P) \to H_{d-1}(\partial P)$ is also the identity map.

The map ∂f_* appears in the another commutative diagram of (reduced) homology groups induced by f:

$$\begin{array}{ccccccc} H_d(P) & \to & H_d(P,\partial P) & \xrightarrow{\partial_*} & H_{d-1}(\partial P) & \to & H_{d-1}(P) \\ & & \Big\downarrow f_* & & \Big\downarrow \partial f_* = id & & \\ H_d(P) & \to & H_d(P,\partial P) & \xrightarrow{\partial_*} & H_{d-1}(\partial P) & \to & H_{d-1}(P) \end{array}$$

where the rows are exact. Since $H_i(P) \cong 0$ for $i \in \{d, d-1\}$, the maps ∂_* are isomorphisms, which implies that f_* is the identity map. Hence f has degree 1.

Therefore the number of pre-images of any regular point y in the image of f is 1 (when counted with sign). So the map f is surjective, and full cells in T gives a cover of P under f. ☐

Thus if we can find a set of points in P such that any d-simplex spanned by $(d+1)$ 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 (in fact, an odd number, because the number of pre-images, counted with sign, is 1). Thus finding full cells in Theorem 1 corresponds precisely to looking for the following kind of finite point set:

Definition. A *pebble set* of a (n, d)-polytope P is a finite set of points (*pebbles*) such that each d-simplex of P contains at most one pebble in its interior.

It is worth remarking now two facts about pebble sets. First, the larger the pebble set, the more full cells we can identify, i.e., the number of full cells is at least the cardinality of the largest size pebble set in P. Second, by the definition of chamber, only one pebble can exist within a chamber and when choosing a pebble p we have the freedom to replace it by any point p' in the interior of the same chamber because p and p' are contained by the same set of d-simplices. We now show that a pebble set of size $(n-d)$ exists for any (n, d)-polytope P by a "facet-pivoting" construction.

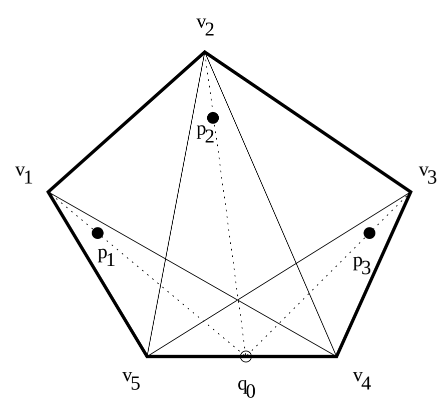

FIGURE 2. A pebble set with pebbles $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3$.

In the simplest situation, if one of the facets of P is a simplex, call this simplex the *base facet*. Choose any point $\mathbf{q}_0$ (the *basepoint*) in the interior of this base facet. Now for each vertex $\mathbf{v}_i$ not in the base facet, choose a point $\mathbf{p}_i$ along a line between $\mathbf{q}_0$ and $\mathbf{v}_i$ but very close to $\mathbf{v}_i$. Exactly how close will be specified in the proof. The collection of all such points $\{\mathbf{p}_i\}$ forms a pebble set; it is size $(n-d)$ because the simplicial base facet has d vertices. See, for example, Figure 2 for the case of a pentagon; it is a $(5,2)$-polytope with pebble set $\{\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3\}$.

If none of the facets are simplicial, then one must choose a non-simplicial facet as base. In this case, choose a pebble set $\{\mathbf{q}_i\}$ for the base facet (an inductive hypothesis is used here) and then use any one of them for a basepoint $\mathbf{q}_0$ to construct $\mathbf{p}_i$ as above. The remaining pebbles are obtained from the other $\mathbf{q}_i$ by perturbing them so they are interior to P. See Figure 3.

Theorem 5. *Any (n,d)-polytope contains a pebble set of size $(n-d)$.*

Proof. We induct on the dimension d. For dimension $d=1$, a polytope is just a line segment spanned by two vertices. Hence $n-d=1$ and clearly any point in the interior of the line segment forms a pebble set.

For any other dimension d, let $V=\{\mathbf{v}_1,\ldots,\mathbf{v}_n\}\in\mathbb{R}^d$ denote the vertices of the given (n,d)-polytope P. Choose any facet F of P as a "base facet", and suppose without loss of generality that it is the convex hull of the last k vertices $\mathbf{v}_{n-k+1},\ldots,\mathbf{v}_n\in\mathbb{R}^d$, $k\geq d$. Then F is a $(d-1)$-dimensional polytope with k vertices, and by the inductive hypothesis, F has a pebble set Q_F with $(k-d+1)$ pebbles $\mathbf{q}_0,\mathbf{q}_1,\ldots,\mathbf{q}_{k-d}$. (If F is a simplex, then $k=d$ and Q_F consists of one point $\mathbf{q}_0$, which can be taken to be any point on the interior of F.)

Let $\operatorname{diam}(P)$ 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 the convex hull of the vertices in $V\setminus\{\mathbf{v}\}$. Since there are finitely many such distances and the vertices are in convex position, H exists and is positive. Set

$$\varepsilon=\frac{H}{2\operatorname{diam}(P)}. \tag{1}$$

Using $\mathbf{q}_0$, let $Q=\{\mathbf{p}_1,\ldots,\mathbf{p}_{n-k}\}$ denote the collection of $(n-k)$ points defined by

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

for $1\leq i\leq n-k$, where ε is a small positive constant given by (1). Thus points in Q lie along straight lines extending from $\mathbf{q}_0$ and very close to the vertices of P not in F.

Because $\mathbf{q}_i$ is in F, it lies on the boundary of P and borders exactly one chamber of P (since by induction it is interior to a single chamber in the facet F). Ignoring $\mathbf{q}_0$ momentarily, for $1\leq i\leq(k-d)$, let $\mathbf{q}_i^\circ$ denote a

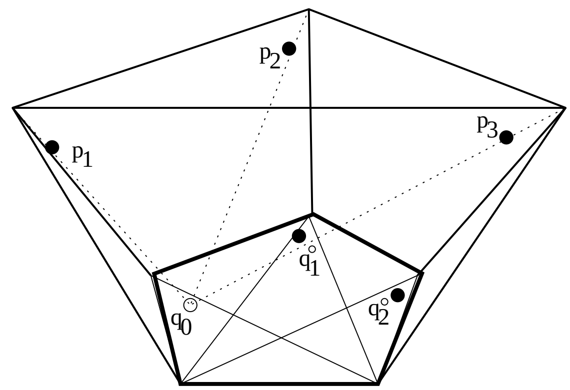

FIGURE 3. A pebble set with pebbles $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{q}_1^\circ, \mathbf{q}_2^\circ$. The pebbles $\mathbf{q}_1^\circ, \mathbf{q}_2^\circ$ lie just above $\mathbf{q}_1, \mathbf{q}_2$ (not shown) on the base of the polytope. Note how $\mathbf{q}_0, \mathbf{q}_1, \mathbf{q}_2$ arise from the pebble set construction in Figure 2.

point obtained by "pushing" $\mathbf{q}_i$ into the interior of the unique chamber that it borders. Let $Q_F^\circ = \{\mathbf{q}_1^\circ, ..., \mathbf{q}_{k-d}^\circ\}$. We shall show that

$$Q \cup Q_F^\circ = \{\mathbf{p}_1, ..., \mathbf{p}_{n-k}, \mathbf{q}_1^\circ, ..., \mathbf{q}_{k-d}^\circ\}$$

is a pebble set for the polytope P. Note that if P has a simplicial facet F, then with this facet as base, the set Q suffices; for then Q_F° is empty and the construction (2) yields the required number of pebbles by choosing any $\mathbf{q}_0$ in the interior of a simplicial facet F.

First we prove some important facts about the $\mathbf{p}_i$ and $\mathbf{q}_i^\circ$.

Lemma 6. *Let S be a d-simplex spanned by vertices of P. If S contains $\mathbf{p}_i$, it must also contain $\mathbf{v}_i$ as one of its vertices.*

Proof. By construction, each $\mathbf{p}_i$ has the property that $\mathbf{p}_i$ is not not in the convex hull of $V \setminus \{\mathbf{v}_i\}$. This follows because

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

hence $\|\mathbf{v}_i - \mathbf{p}_i\| \leq H/2$, implying that the distance of $\mathbf{p}_i$ from the convex hull of $V \setminus \{\mathbf{v}_i\}$ is greater than or equal to $H/2$.

Since the convex hull of $V \setminus \{\mathbf{v}_i\}$ does not contain $\mathbf{p}_i$, if S is to contain $\mathbf{p}_i$ it must contain $\mathbf{v}_i$ as one of its vertices. □

Lemma 7. *Let S be a non-degenerate d-simplex spanned by vertices of P. Then $\mathbf{q}_i^\circ$ is in S if and only if $\mathbf{q}_i$ is in $S \cap F$.*

Proof. Since $\mathbf{q}_i^\circ$ is in the unique chamber of P that $\mathbf{q}_i$ borders, any non-degenerate simplex containing $\mathbf{q}_i$ must contain $\mathbf{q}_i^\circ$. Conversely, any simplex S containing $\mathbf{q}_i^\circ$ must contain its chamber and therefore contains $\mathbf{q}_i$. Since $\mathbf{q}_i$ is in F, then $\mathbf{q}_i$ is in $S \cap F$. □

The next three lemmas will show that $Q \cup Q_F^\circ$ is a pebble set for P.

Lemma 8. *Any d-simplex S spanned by vertices of P contains no more than one pebble of Q.*

Proof. If S is degenerate (i.e., the convex hull of those vertices is not full dimensional), then it clearly contains no pebbles of Q because the $\mathbf{p}_i$ are by construction in the interior of a chamber. So we may assume that S is non-degenerate.

Let $\mathbf{s}_1, ..., \mathbf{s}_{d+1} \in V$ denote the vertices of S. Suppose by way of contradiction that S contained more than one point of Q. Then $\mathbf{p}_{i'}$ and $\mathbf{p}_{j'}$ are contained in S for distinct i', j', where $1 \le i', j' \le n-k$. Lemma 6 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 whose columns consist of $\mathbf{q}_0$ and the vertices of S, adjoined with a row of 1's:

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

This is a $(d+1) \times (d+2)$ matrix that has rank $(d+1)$ because the $\mathbf{s}_i$ are affinely independent (by the non-degeneracy of S). So the kernel of A, $\ker(A)$ is 1-dimensional. Note that $\mathbf{p}_{i'} \in S$ implies that it is a convex combination of the first $(d+1)$ columns of A. On the other hand, by construction, it is also a convex combination of $\mathbf{s}_1$ and $\mathbf{q}_0$. Thus there exist constants $0 \le x_1, x_2, ..., x_{d+1} \le 1$ satisfying

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

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

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

for some constants $0 \le y_1, y_2, ..., y_{d+1} \le 1$. The above equations show that $(x_1 + \varepsilon - 1, x_2, x_3 \dots, x_{d+1}, -\varepsilon)^T$ and $(y_1, y_2 + \varepsilon - 1, y_3, \dots, y_{d+1}, -\varepsilon)^T$ 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 $\mathbf{q}_0$ as an affine combination of the vertices of S:

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

Since $x_2, \dots, x_{d+1} \geq 0$ and $\mathbf{q}_0$ 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), $\mathbf{q}_0$ is on a facet of S. This means that it is spanned by d vertices on the facet F of P. Thus the vertices of S must include those d vertices but by Lemma 6, $\mathbf{v}_{i'}$ and $\mathbf{v}_{j'}$ as well. Since $\mathbf{v}_{i'}, \mathbf{v}_{j'}$ were not on the facet F (because S is non-degenerate), we obtain a contradiction since S cannot contain more than $d+1$ vertices. Thus the equality $x_1 + \varepsilon = 1$ cannot hold, and $x_1 + \varepsilon < 1$ as desired. □

Lemma 9. *Any d-simplex S spanned by vertices of P contains no more than one pebble in Q_F°.*

Proof. Since $S \cap F$ is a simplex in F that contains at most one point of Q_F, then by Lemma 7, S can contain at most one point of Q_F°. □

Lemma 10. *Any d-simplex S spanned by vertices of P cannot contain pebbles of Q and Q_F° simultaneously.*

Proof. Suppose S contained a point $\mathbf{q}_i^\circ$ of Q_F°. Then by Lemma 7, $S \cap F$ contains $\mathbf{q}_i$ of Q_F. Since Q_F was a pebble set for the facet F, $S \cap F$ cannot also contain $\mathbf{q}_0$.

If S also contained a pebble $\mathbf{p}_{i'}$ of Q, then by Lemma 6, S contains $\mathbf{v}_{i'}$ as a vertex. Since $S \cap F$ contains $\mathbf{q}_i$ which is interior to a chamber of F, S must also contain d vertices of F. Since $\mathbf{q}_0$ is in F (but not in $S \cap F$), $\mathbf{q}_0$ is expressible as a linear combination (but not convex combination) of those d vertices. This linear combination, when substituted for $\mathbf{q}_0$ in (2), would show that the pebble $\mathbf{p}_{i'}$ is not a convex combination of $\mathbf{v}_{i'}$ and those d vertices. This contradicts the fact that $\mathbf{p}_{i'}$ was in S to begin with. □

Together, the three lemmas above show that S cannot contain more than one point of $Q \cup Q_F^\circ$, which concludes the proof of Theorem 5. □

Together, Proposition 4 and Theorem 5 prove Theorem 1.

3. Graphs for Path-following and Simplicial Polytopes.

Sperner's lemma has a number of constructive proofs which rely on "path-following" arguments (see, for example, the survey of Todd [18]). Path-following arguments work by using a labelling to determine a path through simplices in a triangulation, in which one endpoint is known and the other endpoint is a full cell. In this section we adapt these ideas for Sperner-labelled polytopes, which are used in the next section to give a constructive "path-following" proof of Theorem 1.

Let P be an (n, d)-polytope with triangulation T and a Sperner labelling using the label set $L = \{1, 2, \dots, n\}$. We define some further terminology and notation that we will use from now on. Let $L(\sigma)$, the *label set* of σ, denote the set of distinct labels of vertices of σ. Let $L(F)$ denote the label set of a face F of P. As defined earlier, a d-simplex σ in T is a *full cell* if the vertex labels of σ are all distinct. Similarly, a $(d-1)$-simplex τ in T is a *full facet* if the vertex labels of τ are all distinct. Note that a full facet on the boundary of P can be regarded as a full cell in that facet.

Definition. Given a Sperner-labelled triangulation T of a polytope P, we define three useful graphs:

1. The *nerve graph* G is a graph with nodes that are simplices of T whose label set is of size at least d. Formally, σ is a node of G if $|L(\sigma)| \geq d$. Two nodes in G are adjacent if (as simplices) one is a face of the other.
2. If K is a subset of the label set $L = \{1, 2, \ldots, n\}$ of size $(d-1)$, the *derived graph* G_K is the subgraph of the nerve graph G consisting of nodes in G whose label sets contain K.
3. Let G' denote the *full cell graph*, whose nodes are full cells in the nerve graph 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. If $\bar{G}$ is a connected component of G, construct the full cell graph $\bar{G}'$ similarly.

Thus the nodes of G and G_K are either full cells, full facets, or d-simplices with exactly one repeated label. The full cell graph G' only has full cells as nodes.

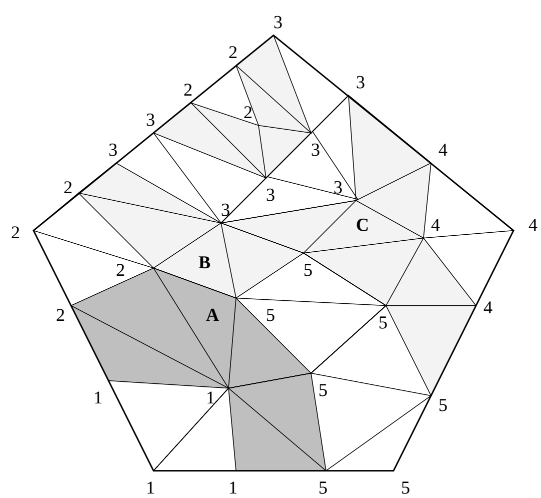

FIGURE 4. A triangulated $(5, 2)$-polytope (a pentagon) with Sperner labelling. If $K = \{1\}$, nodes of G_K consist of the dark-shaded simplices, nodes of G consist of dark and light-shaded simplices, and nodes of G' consist of the three full cells marked by A, B, C.

Example. The pentagon in Figure 4 has dimension $d = 2$. Let $K = \{1\}$, a set of cardinality $(d-1)$. Then the derived graph G_K consists of 1-simplices and 2-simplices that are darkly shaded, and it is a subgraph of the nerve graph G consisting of the dark and light-shaded 1-simplices and 2-simplices in Figure 4. In Figure 4, G' is a 3-node graph with nodes A, B, C, the full cells. In G', A is adjacent to B, and B is adjacent to C, but A is not adjacent to C.

As the example shows, the nerve graph G branches in $(d+1)$ directions at full cells, while the derived graph G_K is the subgraph consisting of paths or loops that "follow" the labels of K along the boundary of the simplices in G. We prove these assertions.

Lemma 11. *The nodes of the derived graph G_K are either of degree 1 or 2 for any K of size $d-1$. A node σ is of degree 1 if and only if σ is a full facet on the boundary of P. Hence G_K is a graph whose connected components are either loops or paths that connect pairs of full facets 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 d-simplex with exactly one repeated label.

If σ is a full cell, since $|K| = d-1$ we see that $L(\sigma)$ 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 d-simplices, unless σ is on the boundary of P, in which case it is the face of exactly one d-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 d-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. □

Lemma 12. *The nodes of the nerve graph G are of degree 1, 2, or $d+1$. 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 $d+1$ 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 d-simplex with exactly one repeated label. The arguments for the latter two cases are identical to those in the proof of Lemma 11 by letting K be the empty set.

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

The nerve graph G may have several components, as in Figure 4. In Theorem 16, we will establish an interesting relation between the labels carried by a component $\bar{G}$ and the number of full cells it carries. First we show that all the labels in a component are carried by the full cells.

Lemma 13. *If σ is adjacent to τ in G, then $L(\sigma) \subseteq L(\tau)$, unless σ is a full cell, in which case $L(\tau) \subset L(\sigma)$. Thus adjacent nodes in $\bar{G}$ carry exactly the same labels unless one of them is a full cell.*

Proof. Suppose σ is a d-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 d-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. □

Lemma 14. *Suppose $\bar{G}$ is connected component of G. If $\bar{G}$ contains at least one full cell as a node, then all the labels in $\bar{G}$ are carried by its full cells.*

For example, in Figure 4, G has two components. One of them has no full cells. In the other component, all of its labels $\{1, 2, 3, 4, 5\}$ are carried by its full cells A, B, C.

Proof. Since $\bar{G}$ 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 $\bar{G}$ that does not intersect any other full cell in $\bar{G}$. Call this path $\{\sigma = \sigma_1, \sigma_2, ..., \sigma_p = \tau\}$. By Lemma 13, $L(\sigma_1) = L(\sigma_2) = ... = L(\sigma_{p-1}) \subset L(\tau)$. Therefore labels carried by the full cells contain all labels carried by any other node of the graph. □

Since the label information in a nerve graph is found in its full cells, it suffices to understand how the full cells connect to each other.

Lemma 15. *Any two adjacent nodes in G' are full cells in T whose label sets contain at least d 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. Repeated application of Lemma 13 yields $L(\tau) \subset L(\sigma_1)$ and $L(\tau) \subset L(\sigma_2)$, so $L(\sigma_1) \cap L(\sigma_2)$ contains at least the d labels in $L(\tau)$. □

We will say the full cell graph G' is a *fully d-labelled graph* because it clearly satisfies four properties:
(a) all nodes in the graph are assigned $(d + 1)$ labels (simply assign to a node σ of G' the label set $L(\sigma)$),
(b) all edges are assigned d labels (assign an edge (σ_1, σ_2) of G' the d labels specified in Lemma 15),
(c) the label set of an edge (σ, τ) (denoted by $L(\sigma, \tau)$) is contained in $L(\sigma) \cap L(\tau)$, and
(d) if τ, τ' are nodes each adjacent to σ, then $L(\tau, \sigma) \neq L(\tau', \sigma)$.

Proposition 16. *Suppose G' is a connected fully d-labeled graph. Let $L(G')$ denote the set of all labels carried by simplices in G' and $|G'|$ the number of* *nodes in G'. Then*

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

We shall use this theorem for graphs G' arising as a full cell graph of one connected component of a nerve graph G. In Figure 4, the full cell graph G' has just one connected component, and $L(G') = 5$, $|G'| = 3$ and $d = 2$, and indeed $3 \geq 5 - 2$.

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

We now assume the statement holds for fully d-labeled graphs with less than j nodes, and show it holds for fully d-labeled 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. Note that this new graph is still fully d-labeled, so by the inductive hypothesis, $|G' - v| \geq |L(G' - v)| - d$.

Clearly $|G' - v| = |G'| - 1$, and $|L(G' - v)| \geq L(G') - 1$ because v has at least d labels in common with some vertex in $G' - v$, by Lemma 15. Hence $|G'| - 1 \geq L(G') - 1 - d$. Adding 1 to both sides gives the desired conclusion. □

This will prove the following useful result.

Theorem 17. *Let T be a Sperner-labelled triangulation of an (n, d)-polytope P. If the nerve graph G has a component $\bar{G}$ that carries all the labels of G, then T contains at least $(n - d)$ full cells.*

Proof. Use $\bar{G}$ to construct the full cell graph $\bar{G}'$ as above, which is a fully d-labelled graph. Note that if $\bar{G}$ is connected then $\bar{G}'$ is also connected. By Lemma 14, $L(\bar{G}) = L(\bar{G}')$. Using Proposition 16, we have

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

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

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

We now use path-following ideas to outline a proof of Atanassov's conjecture in the special case where the polytope is simplicial. This will motivate the proof of Theorem 1 for arbitrary (n, d)-polytopes in the subsequent section.

Theorem 18. *If P is a simplicial polytope, there is some component $\bar{G}$ of the nerve graph G which meets every facet of P, and hence carries all labels of G.*

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

Consider two "adjacent" facets F_1, F_2 of P, whose intersection is a ridge of the polytope P spanned by $(d-1)$ vertices of P. These vertices have distinct labels; let K be their label set. The derived graph G_K consists of loops or paths whose endpoints in G_K must be full facets in F_1 or F_2, since the Sperner labelling guarantees that no other facet of P has a label set containing K.

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, since paths in G_K are connected, this argument shows that

$$\chi(\bar{G}, F_1) \equiv \chi(\bar{G}, F_2) \mod 2$$

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 $\chi(\bar{G}, F)$ is independent of the facet F. We denote this parity by $\rho(\bar{G})$. Since $\chi(G, F)$ is also independent of facet, we can define $\rho(G)$ similarly.

Since $\chi(G, F)$ is the sum of $\chi(\bar{G}, F)$ 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 the usual Sperner's lemma applied to (any) simplicial facet F shows that there are an odd number of full facets of T in the facet F.

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

Theorem 19. *Any Sperner-labelled triangulation of a simplicial* (n, d)*-polytope must contain at least* $(n-d)$ *full cells.*

Proof. This follows immediately from Theorem 18 and Theorem 17. □

To extend this proof for non-simplicial polytopes requires some new ideas but follows the basic pattern: (1) find a function χ that counts the number of times a component $\bar{G}$ of G meets a certain facet in a certain way, and show that this function only depends on $\bar{G}$, and (2) appeal to the usual Sperner's lemma for simplices in a lower dimension to constructively show that the parity of χ summed over all components $\bar{G}$ must be odd. For the non-simplicial case, we cannot guarantee that any faces of P except those in dimension 1 are simplicial. How to connect dimension 1 to dimension d is tackled in the next section, and the *flag graph* introduced there gives a constructive procedure for finding certain full cells. Then we construct a counting function χ to show that there are at least $(n-d)$ full cells for an (n,d)-polytope.

4. The Flag Graph and Arbitrary Polytopes.

Throughout this section, let the symbol $\equiv$ denote equivalence mod 2. Recall that $L(F)$ denotes the label set of a face F. Let $\mathcal{F}$ denote a *flag* of the polytope P, i.e., a choice of faces $F_1 \subset F_2 \subset \ldots \subset F_d$ where F_i is an i-face of P. When the choice of F_i is not understood by context, we refer to the i-face of a particular flag $\mathcal{F}$ by writing $F_i(\mathcal{F})$.

Given a flag $\mathcal{F}$, it will be extremely useful to construct "super-paths" containing simplices of P of various dimensions whose endpoints are either on a 1-dimensional edge or a d-dimensional full cell.

Definition. Let P be an (n,d)-polytope with a Sperner-labelled triangulation T. Let $\mathcal{F}$ be a flag of P. We define the *flag graph* $G_{\mathcal{F}}$ in the following way. For $1 \le k \le d$, a k-simplex $\sigma \in T$ is a node in the graph $G_{\mathcal{F}}$ if and only if σ is one of four types:

(I). the k-simplex σ is carried by the k-face F_k and
$|L(\sigma) \cap L(F_i)| = i+1$ for all $1 \le i \le k$.

(II). the k-simplex σ is carried by the $(k+1)$-face F_{k+1} and
$|L(\sigma) \cap L(F_i)| = i+1$ for all $1 \le i \le k$.

(III). the k-simplex σ is carried by the k-face F_k and
$|L(\sigma) \cap L(F_k)| = k$ and
$|L(\sigma) \cap L(F_i)| = i+1$ for all $1 \le i \le k-1$.

(IV). the k-simplex σ is carried by the k-face F_k and
there is an I such that
$|L(\sigma) \cap L(F_k)| = k+1$,
$|L(\sigma) \cap L(F_i)| = i+2$ for all $I \le i \le k-1$, and
$|L(\sigma) \cap L(F_i)| = i+1$ for all $1 \le i < I$.

Two nodes are adjacent in $G_{\mathcal{F}}$ if (as simplices) one is a facet of the other and at least one of the pair is of type (I) or (II).

Note that if σ is a type (I) simplex in $G_{\mathcal{F}}$, then it is a "non-degenerate" full cell of the k-face that it is carried in, i.e., the vertices of P corresponding to the labels in $L(\sigma)$ span a k-dimensional simplex. A type (II) simplex is a non-degenerate full facet in the $(k+1)$-face that it is carried in. A type (III) simplex has just one repeated label and satisfies a certain kind of non-degeneracy (that ensures its two full facets are non-degenerate). A type (IV) simplex is one kind of degenerate full cell in the k-face that it is carried in (but such that it has exactly two facets which are non-degenerate).

Example. Let P be a $(7,3)$-polytope P, i.e., a 3-dimensional polytope with 7 vertices, and suppose T is a Sperner-labelled triangulation of P. Let $F_1 \subset F_2 \subset F_3$ be a flag $\mathcal{F}$ of P with label sets $\{1,2\} \subset \{1,2,3,4,5\} \subset \{1,2,\ldots,7\}$, respectively.

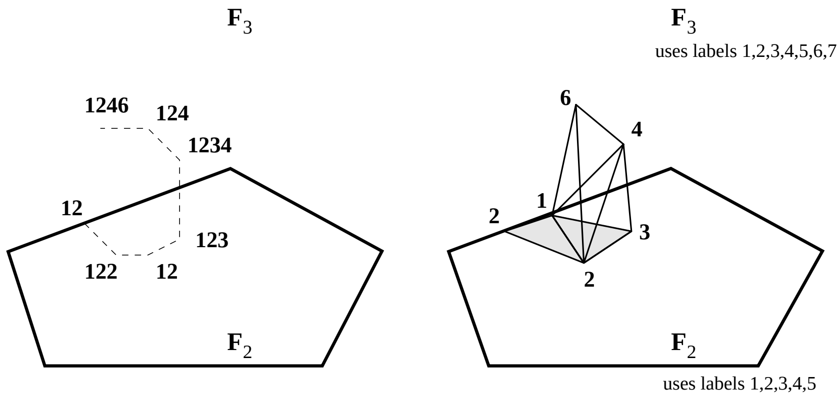

FIGURE 5. A path through a $(7,3)$-polytope. The figure at left shows the label sets along a path in the flag-graph. The figure at right shows the simplices of the same path in the triangulation.

Consider the following collection of simplices shown in Figure 5. Let $\sigma_1, ..., \sigma_7$ be simplices with label sets: $L(\sigma_1) = \{1,2\}$, $L(\sigma_2) = \{1,2\}$ (with repeated label 2), $L(\sigma_3) = \{1,2\}$, $L(\sigma_4) = \{1,2,3\}$, $L(\sigma_5) = \{1,2,3,4\}$, $L(\sigma_6) = \{1,2,4\}$. $L(\sigma_7) = \{1,2,4,6\}$ such that in each pair $\{\sigma_i, \sigma_{i+1}\}$, one is a facet of the other. The simplices $\sigma_1, ..., \sigma_4$ are carried by the face F_2 and other are carried by the face F_3. Each of these simplices is a node in the graph $G_{\mathcal{F}}$: simplices $\sigma_1, \sigma_4, \sigma_7$ are type (I), σ_3, σ_6 are of type (II), σ_2 is of type (III), and σ_5 is of type (IV). Furthermore, each pair σ_i and σ_{i+1} are adjacent in $G_{\mathcal{F}}$. Except for σ_1 and σ_7, each of these simplices has exactly 2 neighbors in $G_{\mathcal{F}}$ so the above sequence traces out a path.

The following result shows that $G_{\mathcal{F}}$ does, in fact, consist of a collection of loops or paths whose endpoints are either 1-dimensional and d-dimensional.

Lemma 20. *Every node σ of $G_{\mathcal{F}}$ has degree 1 or 2, and has degree 1 only when σ is a 1-simplex or a d-simplex in $G_{\mathcal{F}}$.*

Proof. Consider a k-simplex σ of type (I). If $k \geq 2$, then σ has a facet determined by the k labels in $L(\sigma) \cap L(F_{k-1})$, and this facet is a $(k-1)$-simplex of type (I) or (II) so it is adjacent to σ. No other facets of σ are types (I)-(IV). If $k \leq d-1$, then σ is a facet of exactly one $(k+1)$-simplex τ (carried by F_{k+1}) that must be of type (I) or (III) or (IV). Thus a type (I) simplex has degree 2 unless $k = 1$ or $k = d$, in which case it has degree 1.

In case (II), the k-simplex σ is the facet of exactly two $(k+1)$-simplices in F_{k+1}; these are either of type (I) or (III) or (IV) and are thus neighbors of σ in $G_{\mathcal{F}}$. Facets of σ are not types (I)-(IV) because they are co-dimension 2 in the face F_{k+1}. Thus type (II) vertices have degree 2.

In case (III), the k-simplex σ has exactly two facets determined by the k labels in $L(\sigma)$; each of these is a $(k-1)$-simplex adjacent to σ in $G_{\mathcal{F}}$ because it is either of type (I) in F_{k-1} or of type (II) in F_k. No other facets of σ are types (I)-(IV). Thus type (III) vertices have degree 2.

In case (IV), the labelling rules show that the set $L(\sigma) \cap (L(F_I) \setminus L(F_{I-1}))$ is of size two. Call these labels a and b. There is exactly one facet of σ that omits the label a and one facet which omits the label b; each of these is a $(k-1)$-simplex of type (I) in F_{k-1} or of type (II) in F_k, so is adjacent to σ in $G_{\mathcal{F}}$. No other facets of σ are non-degenerate; therefore σ cannot be types (I) or (II). Thus type (IV) vertices have degree 2. □

We remark that in the definition of the flag graph, we require at least one of an adjacent pair to be of type (I) or (II) because without this restriction, some type (IV) vertices could have degree greater than 2. For instance, in a Sperner-labelled, triangulated $(9,4)$-polytope, suppose $F_1 \subset F_2 \subset F_3 \subset F_4$ is a flag $\mathcal{F}$ of P with label sets $\{1,2\} \subset \{1,2,3,4,5\} \subset \{1,2,...,7\} \subset \{1,2,...,9\}$, respectively. If σ is a 4-simplex in F_4 with label set $\{1,2,3,4,6\}$ such that its face τ with labels $\{1,2,3,4\}$ is carried in F_3, then both σ and τ are of type (IV). They each already have two facets in $G_{\mathcal{F}}$ of type (II), so we would not want to define them to be adjacent to each other.

Theorem 21. *A Sperner-labelled triangulation of an (n,d)-polytope contains, for each edge F_1 of P, a non-degenerate full cell whose labels contain $L(F_1)$.*

Proof. For any flag $\mathcal{F}$ containing the edge F_1, Lemma 20 shows that $G_{\mathcal{F}}$ consist of loops or paths whose endpoints are non-degenerate full cells in F_1 or F_d; thus the total number of such end points full cells in F_1 and in F_d must be of the same parity. On the other hand, the 1-dimensional Sperner's lemma shows that the number of full cells in F_1 is odd. So the number of non-degenerate full cells in F_d in $G_{\mathcal{F}}$ must be odd. In particular there is at least one non-degenerate full cell in F_d whose label set contains $L(F_1)$.

□

Notice that the above proof is constructive; the graph $G_{\mathcal{F}}$ yields a method for locating a non-degenerate full cell for any choice of flag $\mathcal{F}$, by starting at one of the full cells on the edge F_1 (an odd number of them are available). At most an even number of them are matched by paths in $G_{\mathcal{F}}$, so at least one of them is matched by a path to a non-degenerate full cell in F_d.

However, as we show now, more can be said about the location of full cells. Rather than locating all of them at the endpoints of paths in a flag-graph, we can show that there is some component $\bar{G}$ of the nerve graph G that contains at least $(n-d)$ full cells. One can trace paths through this component to find them. We find a component $\bar{G}$ of G that carries all labels. Theorem 16 will imply that the component must have $(n-d)$ full cells. As in the case for simplicial polytopes, the key rests on defining a function χ that counts the number of times that $\bar{G}$ meets a facet in a certain way, and then showing that the parity of χ exhibits a certain kind of invariance— it really only depends on $\bar{G}$. Any component with non-zero parity will be the desired component.

Definition. Suppose F is a facet of P and R is a ridge of P that is a facet of F. Let $\chi(G, F, R)$ denote the number of nodes σ of the nerve graph G in the facet F such that $|L(\sigma) \cap L(R)| = d - 1$.

Similarly if K is any $(d-1)$-subset of $L(F)$, let $\chi(G, F, K)$ denote the number of nodes σ of the graph G in the facet F such that $|L(\sigma) \cap K| = d - 1$.

If $\bar{G}$ is a connected component of G, define $\chi(\bar{G}, F, R)$ and $\chi(\bar{G}, F, K)$ similarly using $\bar{G}$ instead of G.

Thus $\chi(G, F, R)$ (resp. $\chi(\bar{G}, F, R)$) counts non-degenerate full cells of type (I) from $G_{\mathcal{F}}$ (resp. $\bar{G}_{\mathcal{F}}$) in the facet F, for all flags $\mathcal{F}$ of P such that $F = F_{d-1}(\mathcal{F})$ and $R = F_{d-2}(\mathcal{F})$. It is easy to show that the parity of $\chi(G, F, R)$ is independent of both F and R:

Theorem 22. *Given any flag $\mathcal{F}$ of P, suppose $F = F_{d-1}(\mathcal{F})$ and $R = F_{d-2}(\mathcal{F})$. Then*

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

Proof. Consider the subgraph contains simplices of dimension $(d-1)$ or lower. This subgraph must be a collection of loops or paths (since $G_{\mathcal{F}}$ is) whose endpoints (even number of them) are non-degenerate full cells in either F_1 or F_{d-1}. But Sperner's lemma in 1-dimension (or simple inspection) shows that the number of full cells in F_1 must be odd. Hence the number of non-degenerate full cells of $G_{\mathcal{F}}$ that meet F_{d-1} must be odd as well. □

The next two theorems show that for *connected* components $\bar{G}$ of G, the parity of $\chi(\bar{G}, F, R)$ is also independent of F and R. This fact does not follow directly from Theorem 22 since we do not know that endpoints of $G_{\mathcal{F}}$ for different flags are connected in $\bar{G}$. To establish this we need to trace connected paths in the nerve graph G rather than the flag graph $G_{\mathcal{F}}$.

Lemma 23. *Let $\bar{G}$ be a connected component of G. Suppose that R, R' are ridges of P and F a facet of P such that R, R' are both facets of F. Then*

$$\chi(\bar{G}, F, R) \equiv \chi(\bar{G}, F, R').$$

Proof. First assume that R and R' are "adjacent" ridges sharing a common facet C (it has dimension $d-3$, if you are keeping track). We claim that

$$\sum_{A,x} \chi(\bar{G}, F, A \cup x) \equiv 0,$$

where x runs over all labels in $L(F)$ that are not in $L(C)$ and A runs over all $(d-2)$-subsets of $L(C)$. (Here we write $A \cup x$ instead of $A \cup \{x\}$ to reduce notation.) The above sum holds because it only counts fully-labelled simplices σ from $\bar{G}$ in F that contain a $(d-2)$-subset A of $L(C)$, and every such σ appears exactly *twice* in this sum; if $L(\sigma) = A \cup a \cup b$, then σ is counted once each in $\chi(\bar{G}, F, A \cup a)$ and in $\chi(\bar{G}, F, A \cup b)$.

On the other hand, if $K = A \cup x$ is not the label set of any ridge of P, then any fully-labelled simplex on the boundary of P that contains K must be contained in the facet F. Since K is of size $d-1$, by Lemma 11, we see that $\chi(\bar{G}, F, K) \equiv 0$ because there is an even number of endpoints of paths in G_K, and such paths are connected subgraphs of the connected graph $\bar{G}$.

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.,

$$\chi(\bar{G}, F, R) + \chi(\bar{G}, F, R') \equiv 0,$$

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. □

Lemma 24. *Let $\bar{G}$ be a connected component of G. Let F, F' be adjacent facets of the polytope P bordering on a common ridge R. Then*

$$\chi(\bar{G}, F, R) \equiv \chi(\bar{G}, F', R).$$

Proof. Let R denote the ridge common to both F and F'. Let K be any non-degenerate subset of $L(R)$ of size $(d-1)$, i.e., K is not a subset of F_i for $i < (d-2)$. Consider the derived graph G_K. By Lemma 11, this graph consists of paths connecting full cells from $\bar{G}$ on the boundary of P that contain the label set K. Since these paths are connected subgraphs of $\bar{G}$, there is an even number of endpoints of these paths in $\bar{G}$.

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

$$\chi(\bar{G}, F, R) + \chi(\bar{G}, F', R) \equiv 0,$$

which produces the desired conclusion. □

Theorem 25. *Let $\bar{G}$ be a connected component of G. The parity of $\chi(\bar{G}, F, R)$ is independent of F and R.*

Proof. Since all facet-ridge pairs (F, R) are connected by a sequence of adjacent facets and ridges, the statement follows from Lemmas 23 and 24. □

Hence we may define the *parity* of $\bar{G}$ to be $\rho(\bar{G}) \equiv \chi(\bar{G}, F, R)$ for any facet-ridge pair (F, R). Similarly, define the *parity* of G to be $\rho(G) \equiv \chi(G, F, R)$ for any facet-ridge pair (F, R), which is well-defined and equal to 1 in light of Theorem 22. Now we may prove

Theorem 26. *If P is an (n, d)-polytope, there is some component $\bar{G}$ of G which carries all labels of P.*

Proof. Fix some flag $\mathcal{F}$ of P, and let $F = F_{d-1}(\mathcal{F})$ and $R = F_{d-2}(\mathcal{F})$. Since $\chi(G, F, R)$ is the sum of $\chi(\bar{G}, F, R)$ 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 22 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 in $L(R)$. Since the flag $\mathcal{F}$ was arbitrary, $\bar{G}$ must carry all labels of P. ☐

This concludes our alternate "path-following" proof of Theorem 1, because the $(n-d)$ count follows immediately from Theorems 26 and 17, while the covering property follows (as before) from Proposition 4.

5. CONCLUSION

As applications of Theorem 1, we prove the corollaries mentioned in the introduction, and make some further remarks.

Corollary 2 follows directly from the degree and covering arguments in the non-constructive proof of Section 2 and the fact that stacked polytopes have triangulations of size $(n-d)$ [11].

Corollary 3, proved below, is stronger than the version stated in [10], which does not mention the covering property of the full cells, nor their cardinality. That weaker version follows also from the combinatorial results of Freund in [6] (in particular his Theorem 4). In fact we should remark the similar flavor of Freund's results to our theorem although he considers triangulations of polytopes with the number of labels equal to the number of facets (not vertices). His results seem to imply a "dual" version of our polytopal Sperner but without an estimate of how many full cells exist.

Proof of Corollary 3. Let C^j, $j = 1, \dots, n$ be the closed sets in the statement. Consider an infinite sequence of triangulations T_k of the polytope P with the property that the maximal diameter of their simplices tends to zero as k goes to infinity. For each triangulation, we label a vertex y of T_k with $i = \min\{j \in \{1, 2, \dots, n\} \mid y \in C^j\}$. This is clearly a Sperner labelling.

By Theorem 1, each triangulation T_k targets a collection of simplices of P corresponding to full cells in T_k. There are only finitely many possible collections (since they are subsets of the set of all simplices of P), and because there are infinitely many T_k, some collection C of simplices must be targeted infinitely many times by a subsequence T_{k_i} of T_k. By Theorem 1, this collection C is a cover of P and therefore has at least $c(P)$ elements.

For each simplex σ in C, choose one full cell σ_i in T_{k_i} that shares the same label set. The σ_i form a sequence of triangles decreasing in size. By the compactness of P, some subsequence of these triangles converges to a point, which (by the labelling rule) must be in the intersection of the closed sets C^j with $j \in L(\sigma)$.

Thus given a point $p \in P$, choose any simplex σ of C that contains P (since C is a cover of P), and let $J_p = L(\sigma)$. Then the above remarks show that J_p satisfies the conditions in the conclusion of the theorem. Moreover, there are at least $c(P)$ different such subsets, one for each σ in C. ☐

We remark that in statement of Theorem 1, the full cells not only correspond to a cover of the polytope P but, in fact, a face-to-face cover. While this is not apparent from the first proof of Theorem 1, it may be seen from the second proof; in particular, the adjacencies in full cell graph G' indicate how the simplices of the cover meet face-to-face. Thus the subsets in Corollary 1 also satisfy this property.

We close with a couple of questions. For a specific polytope P, define the *pebble number* $p(P)$ to be the size of its largest pebble set. The $(n-d)$ lower bound of Theorem 1 is tight, acheived by stacked polytopes whose vertices are assigned different labels. But for a specific polytope P, the arguments of Section 2 show that the lower bound $(n-d)$ can be improved to $p(P)$. What can be said about the value of $p(P)$?

We can provide at least two upper bounds for this number. On one hand $p(P) \leq c(P)$, because for a maximal pebble set, at most one pebble lies in each simplex of a minimal size cover. On the other hand, consider the simplex-chamber incidence 0/1 matrix M introduced in [1]. As the columns correspond to chambers a pebble selection is essentially a selection of a "row-echelon" submatrix; therefore the rank of M is an upper bound on the size of pebble sets. Our pebble construction gives an algorithm for selecting an explicit independent set of columns of M (although this may not always be a basis).

A related question is: for a specific polytope P, how can one determine the minimal cover size $c(P)$? Although Corollary 2 gives a general sharp lower bound for all polytopes, we know that sometimes minimal covers are much larger for specific polytopes, such as for cubes (as the volume arguments in [15] show). Also note that the minimal cover may be strictly smaller than the minimal triangulation (an example is contained in [3]).

Finding other explicit constructions of pebble sets (besides our "facet-pivoting" construction of Section 2) that work for specific polytopes may shed some light on these questions.

References

1. Alekseyevskaya, T. V. "Combinatorial bases in systems of simplices and chambers" *Discrete Math.* 157 (1996) 15-37.
2. Atanassov, K. T., "On Sperner's Lemma," *Stud. Sci. Math. Hungarica* **32** (1996), 71-74.
3. Below, A., Brehm, U., De Loera, J.A., and Richter-Gebert, J., "Minimal simplicial dissections and triangulations of convex 3-polytopes", *Discrete and computational Geometry* **24** (2000), 35-48.
4. Cohen, D. I. A., "On the Sperner Lemma," *J. Comb. Theory* **2** (1967), 585-587.
5. Cronin, J., *Fixed Points and Topological Degree in Nonlinear Analysis*, Amer. Math. Soc., Providence, RI, 1964.
6. Freund, R. M., "Combinatorial Analogs of Brouwer's Fixed-Point Theorem on a Bounded Polyhedron," *J. Comb. Theory, Series B* **47** (1989), 192-219.
7. Guillemin, V. and Pollack, A. "Differential Topology". Prentice-Hall, Englewood Cliffs, NJ, 1974.
8. Knaster, B., Kuratowski, C., and Mazurkiewicz, S., Ein Beweis des Fixpunktsatzes für n-dimensionale Simplexe, *Fund. Math.* **14** (1929), 132-137.

9. Kuhn, H. W., "Simplicial Approximation of Fixed Points," *Proc. Natl. Acad. Sci.* **61** (1968), 1238-1242.
10. van der Laan, G., Talman, D., and Yang, Z., "Intersection theorems on polytopes", *Math. Program.* **84** (1999) 25–38.
11. Rothschild B. L. and Straus, E.G, " On triangulations of the convex hull of n points", *Combinatorica* **5** (1985), 167–179.
12. Scarf, H., "The Approximation of Fixed Points of a Continuous Mapping," *SIAM J. Appl. Math.* **15** (1967), 1328-1343.
13. Seifert, H., and Threlfall, W., *Seifert and Threlfall: A Textbook of Topology*, Academic Press, New York, 1980.
14. Simmons, F. W. and F. E. Su, "Consensus-halving via theorems of Borsuk-Ulam and Tucker," preprint.
15. Smith, W.D., "A lower bound for the simplexity of the n-cube via hyperbolic volumes" special issue Combinatorics of polytopes. European J. Combin. 21, 2000, no. 1, 131–137.
16. Sperner, E., Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes, *Abh. Math. Sem. Hamburg* **6** (1928), 265-272.
17. Su, F. E., "Rental Harmony: Sperner's Lemma in Fair Division," *Amer. Math. Monthly* **106** (1999), 930-942.
18. Todd, M. J., *The Computation of Fixed Points and Applications*, Lecture Notes in Econ. and Math. Syst., Springer-Verlang, New York (1976).
19. Yang, Z., *Computing Equilibria and Fixed Points*, Kluwer Academic Publishers, Boston (1999).
20. Ziegler, G. M., *Lectures on Polytopes*, Springer-Verlag, New York (1995).