6.896 Topics in Algorithmic Game Theory	February 22, 2010

Lecture 6

Lecturer: Constantinos Daskalakis *Scribe: Jason Biddle, Debmalya Panigrahi*

NOTE: The content of these notes has not been formally reviewed by the lecturer. It is recommended that they are read critically.

In our last lecture, we proved Nash's Theorem using Broweur's Fixed Point Theorem. We also showed Brouwer's Theorem via Sperner's Lemma in 2-d using a limiting argument to go from the discrete combinatorial problem to the topological one. In this lecture, we present a multidimensional generalization of the proof from last time. Our proof differs from those typically found in the literature. In particular, we will insist that each step of the proof be constructive. Using constructive arguments, we shall be able to pin down the complexity-theoretic nature of the proof and make the steps algorithmic in subsequent lectures. In the first part of our lecture, we present a framework for the multidimensional generalization of Sperner's Lemma.

- A Canonical Triangulation of the Hypercube
- A Legal Coloring Rule

In the second part, we formally state Sperner's Lemma in n dimensions and prove it using the following constructive steps:

- Colored Envelope Construction
- Definition of the Walk
- Identification of the Starting Simplex
- Direction of the Walk

1 Framework

Recall that in the 2-dimensional case we had a square which was divided into triangles. We had also defined a legal coloring scheme for the triangle vertices lying on the boundary of the square. Now let us extend those concepts to higher dimensions.

1.1 Canonical Triangulation of the Hypercube

We begin by introducing the n-simplex as the n-dimensional analog of the triangle in 2 dimensions as shown in Figure 1. The n-simplex is simply the n-dimensional polytope formed by the convex hull of $n+1$ points in general position.

Next we introduce the hypercube $[0,1]^n$ as the n-dimensional analog of our original square in 2 dimensions. We divide the hypercube into cubelets of equal size. That is, we divide each dimension of the hypercube $[0,1]^n$ into integer multiples of 2^{-m} for some positive integer m as shown in Figure 2. This division into cubelets provides us with a set of vertices for the hypercube, namely all points in the hypercube whose coordinates are multiples of 2^{-m}.

Now we define a simplicization of the hypercube on these vertices by splitting each cubelet into simplices. We define our simplicization of the cubelets such that if two cubelets share a facet, the simplicization defined in the two cubelets coincides on that facet. We do this by ensuring all simplices of the cubelet use the vertices $[0]^n$ and $[1]^n$. An example simplicization for a cubelet in 3 dimensions where all tetrahedra share the vertices $(0,0,0)$ and $(1,1,1)$ is shown in Figure 3.

Formally, every simplex corresponds to a permutation of the coordinates. The points contained in each simplex are all of the points inside the hypercube which satisfy the following definition:

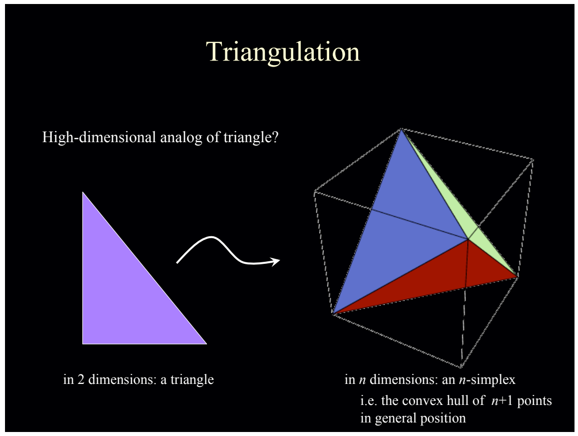

Figure 1: The n-simplex as the high-dimensional analog of the triangle.

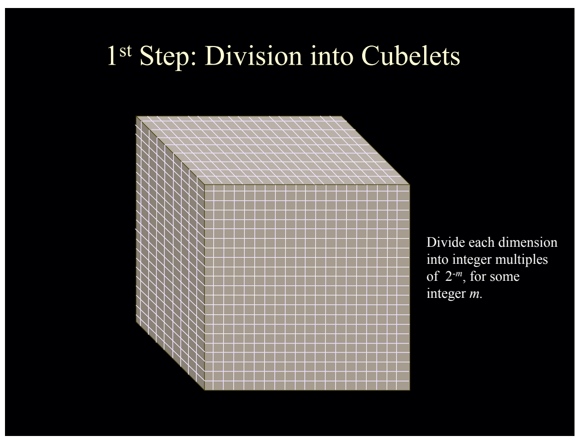

Figure 2: Dividing the hypercube into cubelets.

Definition 1. *For a permutation $\pi : [n] \to [n]$,*

$$\mathcal{T}_\pi := \{x \in [0,1]^n \mid x_{\pi(1)} \leq x_{\pi(2)} \leq ... \leq x_{\pi(n)}\}.$$

Claim 1. *The unique integral corners of $\mathcal{T}_\pi$ are the following $n+1$ points:*

$$\begin{array}{rcccccc}
 & x_{\pi(1)} & x_{\pi(2)} & \dots & x_{\pi(n-2)} & x_{\pi(n-1)} & x_{\pi(n)} \\
 & & & & & & \\
v_1^\pi = & 0 & 0 & \dots & 0 & 0 & 0 \\
v_2^\pi = & 0 & 0 & \dots & 0 & 0 & 1 \\
v_3^\pi = & 0 & 0 & \dots & 0 & 1 & 1 \\
v_4^\pi = & 0 & 0 & \dots & 1 & 1 & 1 \\
\vdots & & & & & & \\
v_n^\pi = & 0 & 1 & \dots & 1 & 1 & 1 \\
v_{n+1}^\pi = & 1 & 1 & \dots & 1 & 1 & 1
\end{array}$$

Proof: Any other integral point in the hypercube that does not respect the ordering must violate the inequality in Definition 1. Suppose vertex v of the hypercube $[0,1]^n$ belongs to set $\mathcal{T}_\pi$ and is not listed

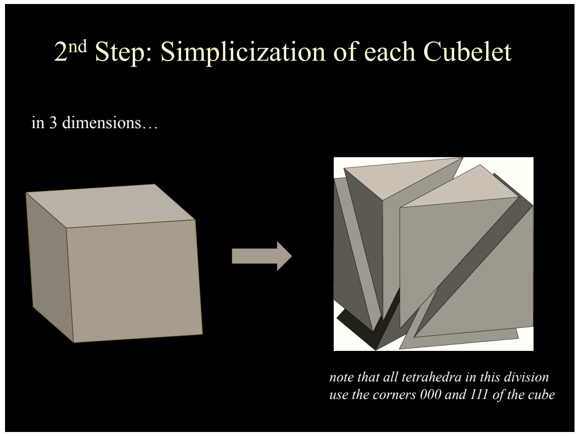

Figure 3: Simplicization of a 3-dimensional cubelet.

above. Vertex v must necessarily contain $v_{\pi(i)} = 1$ and $v_{\pi(i+1)} = 0$ for some $i \in \{1, ..., n-1\}$ which implies $v_{\pi(i)} \not\leq v_{\pi(i+1)}$, a contradiction. Therefore, vertex v cannot belong to set $\mathcal{T}_\pi$.

$\square$

Claim 2. *$\mathcal{T}_\pi$ is a simplex.*

Proof: We can easily express any point $x \in \mathcal{T}_\pi$ as a convex combination of the integral corners $v_1^\pi, v_1^\pi, ..., v_{n+1}^\pi$ of $\mathcal{T}_\pi$ via the following procedure:

$$
\begin{aligned}
\text{Let } y &= x \\
\alpha_{n+1} &= y_{i_{n+1}} \geq 0 \qquad \text{where} \quad i_{n+1} = \arg\min_i \; y_i \\
y &\leftarrow y - \alpha_{n+1} v_{n+1}^\pi \\
\\
\alpha_n &= y_{i_n} \geq 0 \qquad \text{where} \quad i_n = \arg\min_{i \notin \{i_{n+1}\}} \; y_i \\
y &\leftarrow y - \alpha_n v_n^\pi \\
&\vdots \\
\alpha_2 &= y_{i_2} \geq 0 \qquad \text{where} \quad i_2 = \arg\min_{i \notin \{i_3, i_4, ..., i_{n+1}\}} \; y_i \\
y &\leftarrow y - \alpha_2 v_2^\pi \\
\alpha_1 &= 1 - \sum_{j=2}^{n+1} \alpha_j \\
\Rightarrow \sum_{j=1}^{n+1} \alpha_j v_j^\pi &= x \quad \text{where} \quad \sum_{j=1}^{n+1} \alpha_j = 1 \quad \Rightarrow \quad \mathcal{T}_\pi \text{ is a simplex.}
\end{aligned}
$$

$\square$

Claim 3. $\bigcup_\pi \mathcal{T}_\pi = [0, 1]^n$.

Proof: Trivially, any point $x \in [0, 1]^n$ satisfies at least one of the permutations, each of which is a simplex from Claim 2; therefore, the union of simplices equals the hypercube.

$\square$

Theorem 1. *$\{\mathcal{T}_\pi\}_\pi$ is a triangulation of $[0,1]^n$.*

Proof: A triangulation of $[0,1]^n$ is a collection of simplices with disjoint interiors whose union equals the hypercube. We demonstrated the latter in Claim 3. To show our simplices have disjoint interiors, we must argue no simplices can intersect at internal point. Obviously, any internal point must strictly satisfy the inequalities from Definition 1. There is no way an internal point can satisfy more than one set of strict inequalities on the permutations of the coordinates. For example, in the 2-dimensional case if an internal point (x_1, x_2) satisfies $x_1 < x_2$ then it cannot also satisfy $x_2 < x_1$.

□

We apply this triangulation to every cubelet in our hypercube division. We can think of each cubelet as living in $[0,1]^n$, and we can scale and translate the cubelet to the corresponding location in the hypercube division. In order for our simplicization of cubelets to be a simplicization of the entire hypercube, the following property must hold:

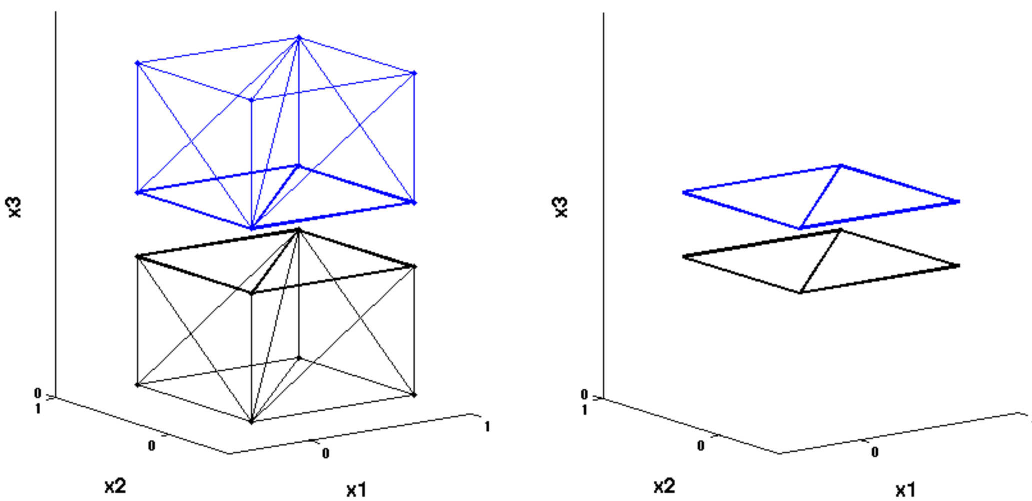

Figure 4: Projection of simplicies in the lower cubelet (black) onto its top face and projection of simplicies in the upper cubelet (blue) onto its bottom face.

Claim 4. *If two cubelets share a face, their simplicizations agree on a common simplicization of the face.*

Proof: Suppose cubelets C_1 and C_2 share a facet. WLOG we can think of each cubelet as living in $[0,1]^n$ and assume that their shared face is $x_n = 0$ for C_1 and $x_n = 1$ for C_2. The projections of the n-simplices of C_1 onto the face $x_n = 0$ produce the $(n-1)$–simplices

$$\mathcal{T}^1_{\pi'} = \{x \in [0,1]^n \mid x_n = 0,\ x_{\pi'(1)} \leq x_{\pi'(2)} \leq \dots \leq x_{\pi'(n-1)}\},$$

for all possible permutations $\pi' : [n-1] \to [n-1]$. Clearly, these simplices define the canonical simplicization of $[0,1]^{n-1}$. Observe that if we project all simplices of C_2 onto the face $x_n = 1$ we obtain the set of simplices

$$\mathcal{T}^2_{\pi'} = \{x \in [0,1]^n \mid x_i = 1,\ x_{\pi'(1)} \leq x_{\pi'(2)} \leq \dots \leq x_{\pi'(n-1)}\},$$

for all possible permutations $\pi' : [n-1] \to [n-1]$, which also define the canonical simplicization of $[0,1]^{n-1}$. Hence, the simplicization of C_1 and C_2 defines the same partition of their shared facet into $(n-1)$-dimensional simplices. □

For illustration consider the 3-dimensional case and suppose we have two cubelets stacked on top of one another as depicted in Figure 4. We can think of each cubelet as living in $[0,1]^3$. The faces of the two cubelets meet when $x_3 = 1$ for the bottom cubelet and when $x_3 = 0$ for the top cubelet. When this occurs, we have freedom in the permutation of the other coordinates in two cubelets. Each permutation defines a projection onto the face where the cubelets meet. Since we simplicize our cubelets on the long diagonal (i.e. all simplicies use the vertices $(0,0,0)$ and $(1,1,1)$), the projections of the top and bottom cubelet coincide at the shared face as shown in Figure 4. One can easily see this holds for the x_1 and x_2 dimensions of the 3-d case.

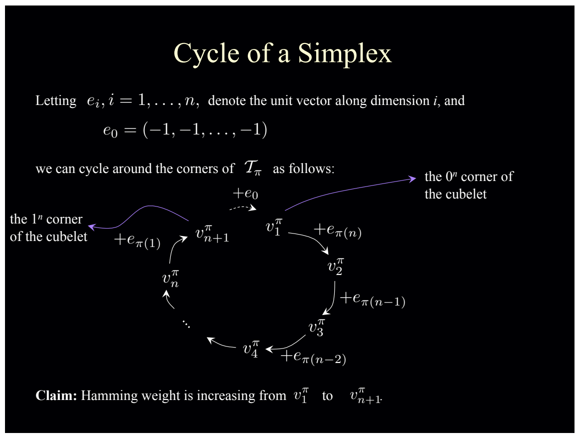

Figure 5: Cycle around the simplex vertices.

Next we introduce a natural way to traverse the vertices of each simplex in our simplicization. Let e_i, $\forall i \in \{1, ..., n\}$, denote the unit vector along dimension i, and $e_0 = (-1, -1, ..., -1)$. Starting at the all zero corner v_1^π, we can move around the $n+1$ corners of the simplex by adding unit vectors until we reach the all ones corner v_{n+1}^π, at which point we can add e_0 to return to the start.

$$
\begin{aligned}
v_1^\pi + e_{\pi(n)} &\rightarrow v_2^\pi \\
v_2^\pi + e_{\pi(n-1)} &\rightarrow v_3^\pi \\
&\vdots \\
v_n^\pi + e_{\pi(1)} &\rightarrow v_{n+1}^\pi \\
v_{n+1}^\pi + e_0 &\rightarrow v_1^\pi
\end{aligned}
$$

Claim 5. *The Hamming weight is increasing from v_1^π through v_{n+1}^π.*

Proof: It's trivial to see the Hamming weight increases by 1 at each step since we're adding a unit vector along dimension i at each step.

$$
\begin{aligned}
v_1^\pi + e_{\pi(n)} &= (0,0,...,0,0,0) + (0,0,...,0,0,1) = (0,0,...,0,0,1) = v_2^\pi \\
v_2^\pi + e_{\pi(n-1)} &= (0,0,...,0,0,1) + (0,0,...,0,1,0) = (0,0,...,0,1,1) = v_3^\pi \\
v_3^\pi + e_{\pi(n-2)} &= (0,0,...,0,1,1) + (0,0,...,1,0,0) = (0,0,...,1,1,1) = v_4^\pi \\
&\vdots \\
v_n^\pi + e_{\pi(1)} &= (0,1,...,1,1,1) + (1,0,...,0,0,0) = (1,1,...,1,1,1) = v_{n+1}^\pi
\end{aligned}
$$

□

1.2 Legal Coloring

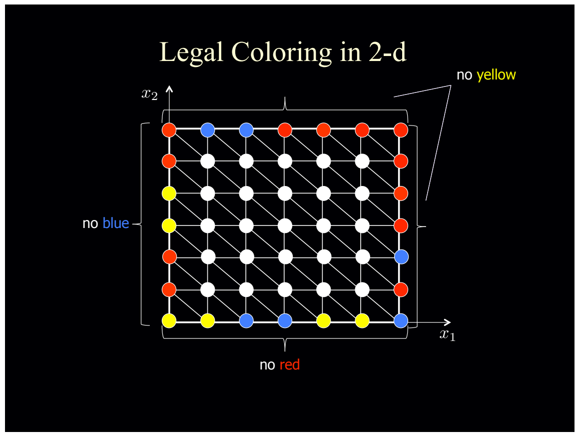

Figure 6: Sperner's coloring condition in 2 dimensions.

Recall for Sperner's coloring condition in 2 dimensions, we had 3 colors: blue (name this color 1), red (color 2), and yellow (color 0). As shown in Figure 6, none of the vertices along the left edge of the square could use blue, none of the vertices along the bottom edge could use red, and the remaining boundary vertices (i.e. the top and right edges of the square) could use yellow. We did not restrict the color of the internal vertices.

Definition 2. (P_2)*: None of the vertices on the left* ($x_1 = 0$) *side of the square uses blue, no vertex on the bottom side* ($x_2 = 0$) *uses red, and no vertex on the other two sides uses yellow.*

For Sperner's coloring condition in n dimensions, we have $n+1$ colors: $0, 1, 2, ..., n$. We color vertices on the boundary of the hypercube according to the following:

Definition 3. (P_n)*: For all $i \in \{1, ..., n\}$, none of the vertices on the face $x_i = 0$ of the hypercube uses color i; moreover, color 0 is not used by any vertex on a face $x_i = 1$, for some $i \in \{1, ..., n\}$.*

For example, in the 3-dimensional case (P_3) we have 4 colors: $0, 1, 2, 3$. Vertices on the left face of the cube (i.e. $x_1 = 0$) cannot use color 1, vertices on the front face (i.e. $x_2 = 0$) cannot use color 2, and vertices on the bottom face (i.e. $x_3 = 0$) cannot use color 3. As shown in Figure 7, vertices on the remaining three faces cannot use color 0.

2 N-Dimensional Sperner's Lemma

Theorem 2 (Sperner 1928 [1])**.** *Suppose that the vertices of the canonical simplicization of the hypercube $[0, 1]^n$ are colored with colors $0, 1, ..., n$ so that the following property is satisfied by the coloring on the boundary.*

(P_n): For all $i \in \{1, ..., n\}$, none of the vertices on the face $x_i = 0$ uses color i; moreover, color 0 is not used by any vertex on a face $x_i = 1$, for some $i \in \{1, ..., n\}$.

Then there exists a panchromatic simplex in the simplicization. In fact, there is an odd number of those.

Remark 1. *We need not restrict ourselves to the canonical simplicization of the hypercube (that is, divide the hypercube into cubelets and divide each cubelet into simplices in the canonical way shown*

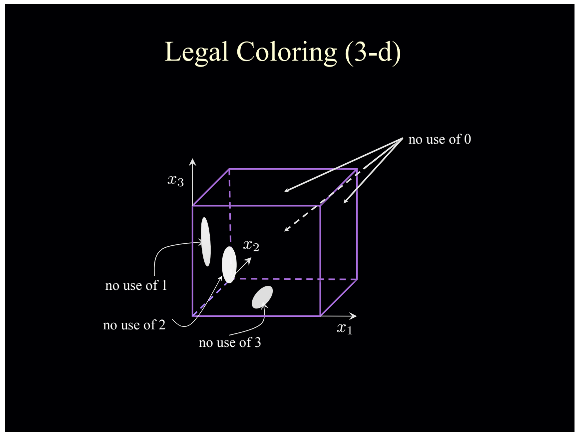

Figure 7: Sperner's coloring condition in 3 dimensions.

above). The conclusion of the theorem is true for any partition of the cube into n-simplices, as long as the coloring satisfies the property stated above.

The reason we state Sperner's lemma in terms of the canonical triangulation is in an effort to provide an algorithmically-friendly version of the computational problem related to Sperner, in which the triangulation and its simplices are easy to define, the neighbors of a simplex can be computed efficiently etc. We follow-up on this in the next lecture. Moreover, our set-up allows us to make all the steps in the proof of Sperner's lemma constructive (except for the length of the walk, see below).

Remark 2. *Sperner's Lemma was originally stated for a coloring of a simplicization of the n-simplex, as stated below:*

Theorem 3 (Original Statement of Sperner's Lemma)**.** *Color the vertices of any simplicization of the n-simplex (a convex combination of points $v_0, v_1, ..., v_n$) with colors $0, 1, ..., n$ so that the facet not containing vertex v_i does not use color i. Then there exists a panchromatic simplex in the simplicization.*

Our legal coloring of the hypercube $[0,1]^n$ corresponds essentially to a legal coloring of the n-simplex defined as in the original statement of Sperner's lemma. Here is why: let $v_0 = 0^n$, $v_i = e_i$, for all $i = 1, \ldots, n$, where e_i is the unit vector along dimension i. For all $i = 1, \ldots, n$, color i is disallowed in the facet of the cube that touches point 0^n and is opposite to point v_i. Moreover, all facets of the hypercube adjacent to 1^n are considered to be opposite to point v_0; hence color 0 is disallowed there.

We proceed to prove the n-dimensional version of Sperner's Lemma by generalizing every step of the 2-dimensional case.

3 Proof of Sperner's Lemma

3.1 Envelope Construction

Recall that in the 2-dimensional case, we introduced an outer boundary (envelope) containing our original square. We colored this envelope in a canonical way, called the *envelope coloring*, which does not violate the legal coloring rules. Figure 8 shows the square shown of Figure 6 enclosed inside an envelope. The envelope coloring is defined as follows (suppose that the envelope is identified with the facets of $[0,1]^2$):

- in the region of the boundary where yellow (color 0) is disallowed in a legal coloring, color all vertices with blue (color 1), except for the boundary of this region with $x_1 = 0$ (because blue is disallowed there);

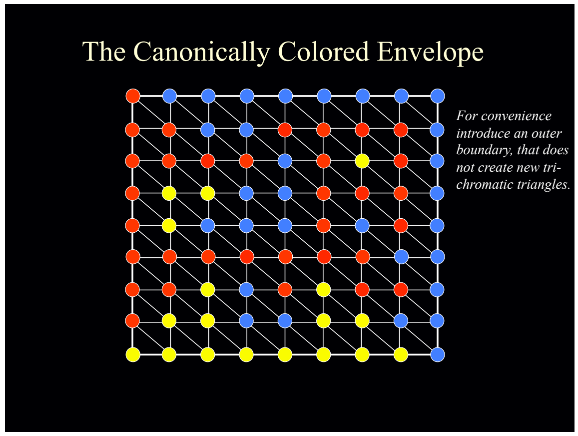

Figure 8: Envelope construction in 2 dimensions.

- given this, where blue (color 1) is disallowed, color with red (color 2), except for the boundary of this region with $x_2 = 0$;
- given this, where red (color 2) is disallowed, color with yellow (color 0).

In the last lecture, we proved that introducing the envelope maintains the legality of the coloring, and that the introduction of the envelope does not create any new trichromatic triangles that did not exist before the envelope was added.

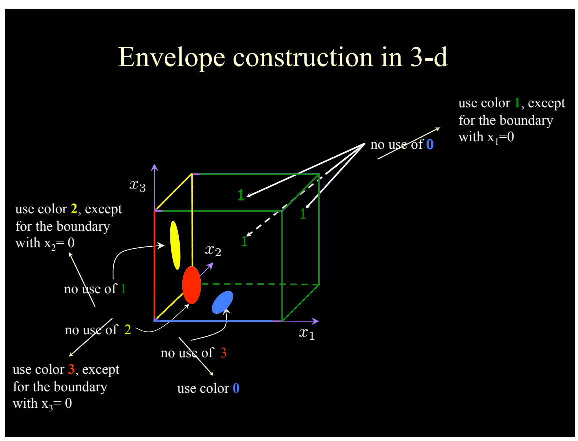

Figure 9: Envelope construction in 3 dimensions.

Definition 4 (Envelope Coloring)**.** *The* envelope coloring *is the coloring of the vertices of the simplicization of the hypercube that lie on the boundary of* $[0, 1]^n$ *in the following greedy manner (the steps are ordered):*

- *where 0 is disallowed, color with 1, except for the boundary of this region with* $x_1 = 0$*;*
- *where 1 is disallowed, color with 2, except for the boundary of this region with* $x_2 = 0$*;*

- *...*
- *where i is disallowed, color with $i+1$, except for the boundary of this region with $x_{i+1} = 0$;*
- *...*
- *where n is disallowed, color with 0.*

As an example, consider the 3-dimensional case shown in Figure 9. Color 0 is not allowed on the cap of the cube (i.e. the top, right, and back faces). We color the boundary there with color 1 except for the boundary with the left face of the cube (i.e. $x_1 = 0$), since color 1 is not allowed on the left face of the cube (i.e. $x_1 = 0$). We color the envelope there with color 2 except for the boundary with the front face of the cube (i.e. $x_2 = 0$). Color 2 was not allowed on the front face of the cube (i.e. $x_2 = 0$). We color the envelope there with color 3 except for the boundary with the bottom face of the cube (i.e. $x_3 = 0$). Finally, color 3 was not allowed on the bottom face of the cube (i.e. $x_3 = 0$). We color the boundary there with color 0.

Claim 6. *Suppose that we start off with a legal coloring of the hypercube and introduce an envelope around it, colored according to the envelope coloring rule defined above. No new panchromatic simplices are introduced by the envelope addition.*

Homework [2 points]
Prove Claim 6 for the n-dimensional case.

3.2 Definition of the Directed Walk

Like we did in the 2-d case, we show that a panchromatic simplex exists by defining a walk that jumps from simplex to simplex of our simplicization, starting at some fixed simplex (independent of the coloring) and is guaranteed to conclude at a panchromatic one. The simplices in our walk (except for the final one) will contain all the colors in the set $\{2, 3, \ldots, n, 0\}$, but will be missing color 1. Call such simplices *colorful*. In particular, every such simplex will have exactly one color repeated twice. So it will contain exactly two facets with colors $\{2, \ldots, n, 0\}$. Call these facets *colorful*. Our walk will be transitioning from simplex to simplex, by pivoting through a colorful facet.

The starting simplex of the walk belongs to the cubelet adjacent to the 0^n vertex of the hypercube, and corresponds to the permutation $\pi = (1, 2, \ldots, n-1, n)$ (refer to Figure 10). This has all the colors in $\{2, 3, \ldots, n, 0\}$ but not color 1, and hence is a colorful simplex. One of its colorful facets lies on $x_1 = 0$, while the other is shared with some neighboring simplex. The walk enters that neighboring simplex through the shared colorful facet. If the other vertex of that simplex has color 1 the walk is over, and the existence of a panchromatic simplex has been established. If the other vertex is not colored 1, that simplex has another colorful facet shared with one of its neighboring simplices. The walk traverses this colorful facet to enter the neighboring simplex. In general, the walk enters a colorful simplex through a colorful facet and exits via its other colorful facet if it exists; otherwise, the simplex is panchromatic and the walk terminates.

To enumerate the possible evolutions of the walk, we observe the following.

- The walk cannot loop into itself in a ρ-shape, since that would require a simplex with three colorful facets.
- The walk cannot exit the hypercube, since the only colorful facet on the boundary belongs to the starting simplex, and the walk cannot reach that simplex from the inside of the hypercube. This is because it would require a third colorful facet for the starting simplex or a violation of the above assertion somewhere else on the walk.
- The walk cannot get into a cycle by returning to the starting simplex since that would need it to enter the starting simplex from outside the hypercube.

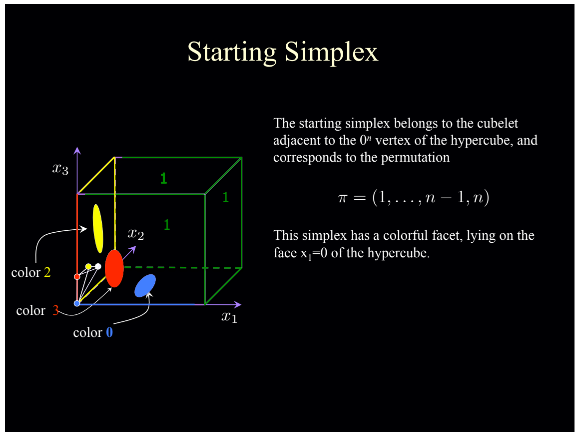

Figure 10: The starting simplex of the directed walk.

The single remaining possibility is that the walk traces a path (i.e. no simplex is repeated) inside the hypercube. Since there is a finite number of simplices, the walk must stop, and the only way this can happen is by encountering color 1 when entering into a simplex through a colorful facet. Then, the terminating simplex is panchromatic.

Lemma 1. *There exists an odd number of panchromatic simplices.*

Proof: After the original walk has terminated, we can start a new walk from some other simplex that is not part of the original walk. If the starting simplex of this walk has no colorful facet, we stop immediately. Otherwise, we start two simultaneous walks by crossing the two colorful facets of the simplex. For each walk, if S is a colorful simplex encountered, exit S from the colorful facet not used to enter S if such a facet exists; otherwise terminate the walk at the panchromatic simplex found. There are two possibilities:

- either the two walks meet and none of the simplices encountered by the walks is panchromatic, or
- the two walks stop at a pair of distinct panchromatic simplices.

In either case, each subsequent (pair of) walks (i.e. other than the first walk) yield either 0 or 2 panchromatic simplices. Thus, the total number of panchromatic simplices in the hypercube is odd. □

Abstractly, we can define a graph on the set of simplices where two simplices are neighbors iff they share a colorful facet (refer to Figure 11). In such a graph each vertex has degree at most 2; hence, it is a collection of cycles and paths. Each endpoint of a path is a panchromatic simplex, except the starting simplex of our first walk. This also shows that the number of panchromatic simplices is odd.

3.3 Directing the Walk

The above argument defines an undirected graph, whose vertex set is the set of simplices in the simplicization of the hypercube and which comprises of paths, cycles and isolated vertices. We will see in the next couple of lectures that in order to understand the precise computational complexity of Sperner's problem, we need to define a directed graph with the above structure (i.e. comprising of directed paths, directed cycles and isolated nodes). We now devise a convention (and an efficient method) for checking which of the two colorful facets of a colorful simplex corresponds to an incoming edge, and which facet corresponds to an outgoing edge.

Given a colorful facet f of some simplex S, we need to decide whether the facet corresponds to the inward or outward direction. To do this we define two permutations τ_f and σ_f as follows. Recall the

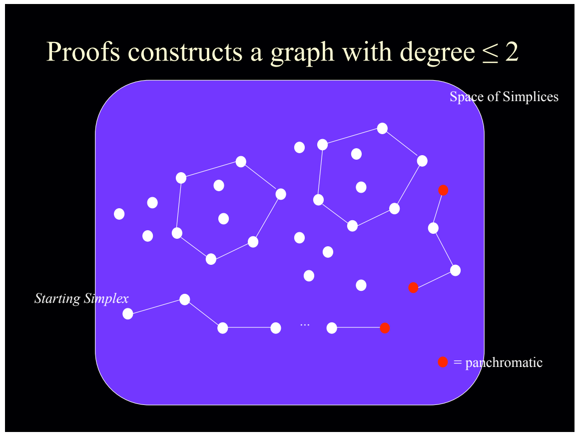

Figure 11: The graph on the set of simplices created by our proof of Sperner's lemma.

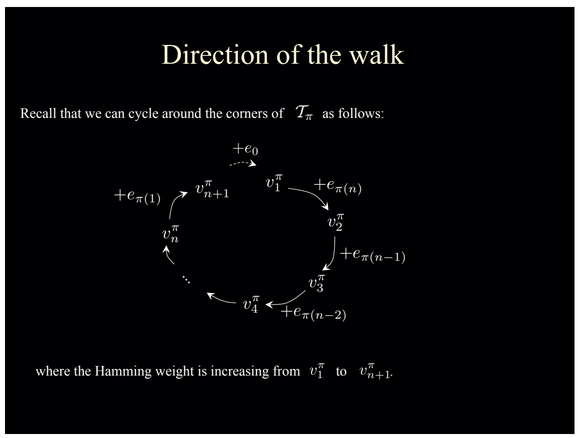

Figure 12: The canonical cycle through a simplex.

canonical cycle corresponding to a simplex (refer to Figure 12). Let w be the vertex not on the colorful facet; w falls somewhere on this cycle. If $w = v_k^\pi$, let τ_f be the following permutation of $0, 1 \ldots, n$: $\pi_{n-k+1}, \pi_{n-k}, \ldots, 0, \pi_n, \ldots, \pi_{n-k+2}$. In other words, suppose we start at w and travel around the cycle to get back to w. Then τ_f is the permutation of indices that we encounter on the arrows as subscripts of e. On the other hand, we define permutation $\sigma_f : \{2, 3, \ldots, n, 0\} \to \{2, 3, \ldots, n, 0\}$ as the order in which the colors $\{2, 3, \ldots, n, 0\}$ appear in the cycle, starting with the vertex of the cycle after w up until the vertex of the cycle before w.

The *sign* of a permutation is the parity of the number of pairwise inversions, i.e. $(-1)^{\#\text{pairwise inversions}}$. Define the sign of facet f in simplex S as

$$sign_S(f) = sign(\tau_f) \cdot sign(\sigma_f).$$

We now prove two properties about the sign function of facets defined above.

Lemma 2. *Suppose f is a colorful facet shared by simplices S and S'. Then*

$$sign_S(f) \cdot sign_{S'} f = -1,$$

i.e. simplices S and S' assign different signs to their shared colorful facet f.

Proof: There are two cases. First assume that S and S' belong to the same cubelet. Then their permutations π, π' are identical, except for a transposition of one adjacent pair of indices. Hence if w, w are the missing vertices from f in S and S' respectively, w, w' are located in the cycle of π, π' respectively between indices i and $i+1$, while all the other shared vertices appear in the same order. Hence, the color permutation σ_f is the same in S, S', while the permutation τ_f has the pair of indices $i, i+1$ transposed and hence has opposite sign in S, S'.

Now, assume that S and S' are in adjacent cubelets. Then f lies on a facet $x_i = 1$ of S and $x_i = 0$ of S'. The vertex not in f in S is 0^n, while the vertex of S' not in f is 1^n. Moreover, to obtain the vertices of f in S', we can replace coordinate i in the vertices of f in S with 0. In other words, permutations π, π' are identical, except that i is moved from the last position in π to the first position in π'. It follows that the color permutation σ_f is the same in S, S', while there is exactly one transposition in going from τ_f in S to τ_f in S'. □

Lemma 3. *Let S be a colorful simplex and f, f' be its two colorful facets. Then*

$$sign_S(f) \cdot sign_S(f') = -1.$$

Proof: Let w, w' be the vertices of S missing from f and f respectively. Without loss of generality, w appears before w' on the cycle, and they are separated by k arcs. Permutations τ_f and $\tau_{f'}$ differ by a cyclic shift of k positions. Thus,

$$sign(\tau_{f'}) \cdot sign(\tau_f) = (-1)^{k(n+1-k)}.$$

We now compare σ_f and $\sigma_{f'}$. Let $\sigma_f = i_2 i_3 \ldots i_1 \ldots i_{k+1} \ldots i_n$ and $\sigma_{f'} = i_{k+1} \ldots i_n i_1 i_2 \ldots i_k$. Thus,

$$sign(\sigma_{f'}) \cdot sign(\sigma_f) = (-1)^{n-1+(n-k)(k-1)}.$$

Hence,

$$sign_S(f) \cdot sign_S(f') = (-1)^{k(n-k+1)+n-1+(n-k)(k-1)} = (-1)^{2k(n-k+1)-1} = -1.$$

□

References

[1] E. Sperner. Neuer Beweis für die Invarianz der Dimensionzahl und des Gebietes. *Abh. Math. Sem.*, 6:265–272, 1928.