

Sperner's Colorings and Optimal Partitioning of the Simplex*

Maryam Mirzakhani[†] Jan Vondrák[‡]

Abstract

We discuss coloring and partitioning questions related to Sperner's Lemma, originally motivated by an application in hardness of approximation. Informally, we call a partitioning of the $(k - 1)$ -dimensional simplex into k parts, or a labeling of a lattice inside the simplex by k colors, "Sperner-admissible" if color i avoids the face opposite to vertex i . The questions we study are of the following flavor: What is the Sperner-admissible labeling/partitioning that makes the total area of the boundary between different colors/parts as small as possible?

First, for a natural arrangement of "cells" in the simplex, we prove an optimal lower bound on the number of cells that must be non-monochromatic in any Sperner-admissible labeling. This lower bound is matched by a simple labeling where each vertex receives the minimum admissible color.

Second, we show for this arrangement that in contrast to Sperner's Lemma, there is a Sperner-admissible labeling such that every cell contains at most 4 colors.

Finally, we prove a geometric variant of the first result: For any Sperner-admissible partition of the regular simplex, the total surface area of the boundary shared by at least two different parts is minimized by the Voronoi partition (A_1^*, \dots, A_k^*) where A_i^* contains all the points whose closest vertex is \mathbf{e}_i . We also discuss possible extensions of this result to general polytopes and some open questions.

1 Introduction

Sperner's Lemma is a gem in combinatorics which was originally discovered by Emmanuel Sperner [12] as a tool to derive a simple proof of Brouwer's Fixed Point Theorem. Since then, Sperner's Lemma has seen numerous applications, notably in the proof of existence of mixed Nash equilibria [11], in fair division [13], and recently it played an important role in the study of computational complexity of finding a Nash equilibrium [3, 2]. At a high level, Sperner's Lemma states that for any coloring of a simplicial subdivision of a simplex satisfying certain boundary conditions, there must be a "rainbow cell" that receives all possible colors. We review Sperner's Lemma in Section 3.

*To appear in "*A Journey through Discrete Mathematics. A Tribute to Jiří Matoušek*", edited by Martin Loebl, Jaroslav Nešetřil and Robin Thomas, due to be published by Springer. The first two results of this paper have also appeared in [10].

[†]Dept. of Mathematics, Stanford University, Stanford, CA; mmirzakh@stanford.edu

[‡]Dept. of Mathematics, Stanford University, Stanford, CA; jvondrak@stanford.edu

The starting point of this work was a question that arises in the study of approximation algorithms for a certain hypergraph labeling problem [4]. The question posed by [4], while in some ways reminiscent of Sperner’s Lemma, is different in the following sense: Instead of asking whether there exists a rainbow cell for any admissible coloring, the question is what is the minimum possible number of cells that must be *non-monochromatic*. (Also, the question arises for a particular regular lattice inside the simplex rather than an arbitrary subdivision.) In this paper, we resolve this question and investigate some related problems.

Before we state our results, let us note the following connection. As the granularity of the subdivision tends to zero, Sperner’s Lemma becomes a statement about certain geometric partitions of the simplex: for any Sperner-admissible partition, where part i avoids the face opposite to vertex i , there must be a point where all parts meet. This result is known as the Knaster-Kuratowski-Mazurkiewicz Lemma [7]. In contrast, the questions we are studying are concerned with the measure of the boundary where at least two different parts meet: This can be viewed as a multi-colored isoperimetric inequality, where we try to partition the simplex in a certain way, so that the surface area of the union of all pairwise boundaries (what we call a *separating set*) is minimized. The way we measure the separating set also affects the problem; the discrete version of the question that is of primary interest to us is mandated by the application in [4]. In the geometric setting, a natural notion of surface area is the Minkowski content of the separating set (which coincides with other notions of volume for well-behaved sets). We give an optimal answer to this question for a regular simplex and discuss other related questions.

To state our results formally, we need some notation that we introduce in Section 2. We postpone our contributions to Sections 4–6, after a discussion of Sperner’s Lemma in Section 3.

2 Preliminaries

We denote vectors in boldface, such as $\mathbf{v} \in \mathbb{R}^k$. The coordinates of \mathbf{v} are written in italics, such as $\mathbf{v} = (v_1, \dots, v_k)$. By \mathbf{e}_i , we denote the canonical basis vectors $(0, \dots, 1, \dots, 0)$. By $\text{conv}(\mathbf{v}_1, \dots, \mathbf{v}_k)$, we denote the convex hull of the respective vectors.

2.1 Simplicial subdivisions of the simplex

Consider the $(k - 1)$ -dimensional simplex defined by

$$\Delta_k = \text{conv}(\mathbf{e}_1, \dots, \mathbf{e}_k) = \left\{ \mathbf{x} = (x_1, x_2, \dots, x_k) \in \mathbb{R}^k : \mathbf{x} \geq 0, \sum_{i=1}^k x_i = 1 \right\}.$$

Simplicial subdivision. A simplicial subdivision of Δ_k is a collection of simplices (“cells”) Σ such that

- The union of the cells in Σ is the simplex Δ_k .
- For any two cells $\sigma_1, \sigma_2 \in \Sigma$, their intersection is either empty or a full face of a certain dimension shared by σ_1, σ_2 .

The Simplex-Lattice Hypergraph. Next, we describe a specific configuration of cells in a simplex; this configuration is actually not a full subdivision since its cells do not cover the full volume of the simplex. It can be completed to a subdivision if desired.¹

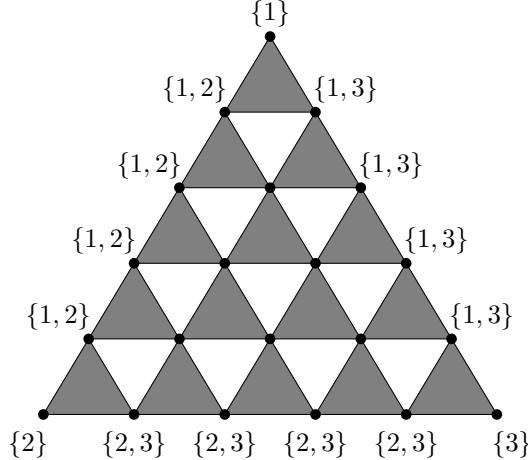
Let $q \geq 1$ be an integer and define

$$\Delta_{k,q} = \left\{ \mathbf{x} = (x_1, x_2, \dots, x_k) \in \mathbb{R}^k : \mathbf{x} \geq 0, \sum_{i=1}^k x_i = q \right\}.$$

We consider a vertex set of all the points in $\Delta_{k,q}$ with integer coordinates:

$$V_{k,q} = \left\{ \mathbf{a} = (a_1, a_2, \dots, a_k) \in \mathbb{Z}^k : \mathbf{a} \geq 0, \sum_{i=1}^k a_i = q \right\}.$$

Figure 1: The Simplex Lattice Hypergraph for $k = 3$ and $q = 5$, with hyperedges shaded in gray. The gray triangles together with the white triangles form a simplicial subdivision. The lists of admissible colors are given on the boundary; for internal vertices the lists are all $\{1, 2, 3\}$.



The *Simplex-Lattice Hypergraph* is a k -uniform hypergraph $H_{k,q} = (V_{k,q}, E_{k,q})$ whose hyperedges (which we also call *cells* due to their geometric interpretation) are indexed by $\mathbf{b} \in \mathbb{Z}_+^k$ such that $\sum_{i=1}^k b_i = q - 1$: we have

$$E_{k,q} = \left\{ e(\mathbf{b}) : \mathbf{b} \in \mathbb{Z}^k, \mathbf{b} \geq 0, \sum_{i=1}^k b_i = q - 1 \right\}$$

¹This specific configuration arises in [4] as an integrality gap example for a certain hypergraph labeling problem; see also [10] for more details.

where $e(\mathbf{b}) = \{\mathbf{b} + \mathbf{e}_1, \mathbf{b} + \mathbf{e}_2, \dots, \mathbf{b} + \mathbf{e}_k\} = \{(b_1 + 1, b_2, \dots, b_k), (b_1, b_2 + 1, \dots, b_k), \dots, (b_1, b_2, \dots, b_k + 1)\}$. For each vertex $\mathbf{a} \in V_{k,q}$, we have a list of admissible colors $L(\mathbf{a})$, which is

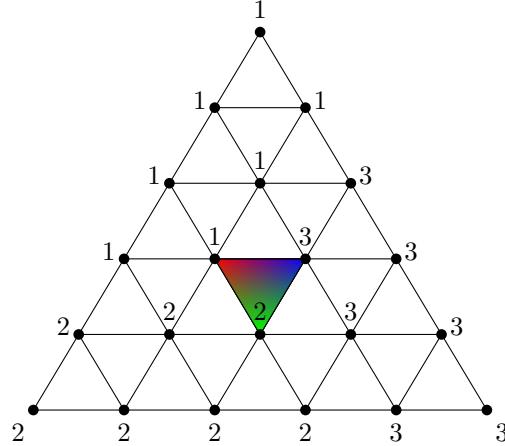
$$L(\mathbf{a}) = \{i \in [k] : a_i > 0\}.$$

3 Sperner's Lemma

First, let us recall the statement of Sperner's Lemma [12]. We consider labelings $\ell : V_{k,q} \rightarrow [k]$. We call a labeling ℓ Sperner-admissible if $\ell(\mathbf{a}) \in L(\mathbf{a})$ for each $\mathbf{a} \in V$; i.e., if $\ell(\mathbf{a}) = j$ then $a_j > 0$.

Lemma 1 (Sperner's Lemma). *For every Sperner-admissible labeling of the vertices of a simplicial subdivision of Δ_k , there is a cell whose vertices receive all k colors.*

Figure 2: A Sperner-admissible labeling for $k = 3$ and $q = 5$. At least one cell in the triangulation (not necessarily in $E_{k,q}$) must be k -colored (rainbow).



We remark that this does not say anything about the Simplex-Lattice Hypergraph: Even if the subdivision uses the point set $V_{k,q}$, the rainbow cell given by Sperner's Lemma might not be a member of $E_{k,q}$ since $E_{k,q}$ consists only of scaled copies of $\Delta_{k,q}$ without rotation; it is not a full subdivision of the simplex. (See Figure 2.)

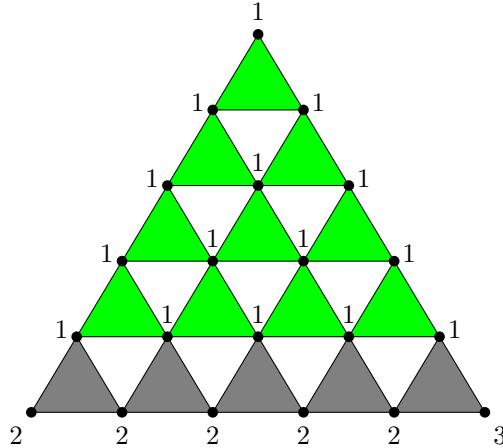
4 The Simplex-Lattice Coloring Lemma

Instead of rainbow cells, the statement proposed (and proved for $k = 3$) in [4] involves non-monochromatic cells.

Proposition 1 (Simplex-Lattice Coloring Lemma). *For any Sperner-admissible labeling $\ell : V_{k,q} \rightarrow [k]$, there are at least $\binom{q+k-3}{k-2}$ hyperedges $e \in E_{k,q}$ that are non-monochromatic under ℓ .*

The first-choice labeling. In particular, Proposition 1 is that a Sperner-admissible labeling minimizing the number of non-monochromatic cells is a “first-choice one” which labels each vertex \mathbf{a} by the smallest coordinate i such that $a_i > 0$. Under this labeling, all the hyperedges $e(\mathbf{b})$ such that $b_1 > 0$ are labeled monochromatically by 1. The only hyperedges that receive more than 1 color are those where $b_1 = 0$, and the number of such hyperedges is exactly $\binom{q+k-3}{k-2}$ (see [4]). Here we give a proof of Proposition 1.

Figure 3: The first-choice labeling.



Proof. Consider the set of hyperedges $E_{k,q}$: observe that it can be written naturally as

$$E_{k,q} = \{e(\mathbf{b}) : \mathbf{b} \in V_{k,q-1}\}.$$

I.e., the hyperedges can be identified one-to-one with the vertices in $V_{k,q-1}$. Recall that $e(\mathbf{b}) = \{\mathbf{b} + \mathbf{e}_1, \mathbf{b} + \mathbf{e}_2, \dots, \mathbf{b} + \mathbf{e}_k\}$. Two hyperedges $e(\mathbf{b}), e(\mathbf{b}')$ share a vertex if and only if $\mathbf{b}' + \mathbf{e}_j = \mathbf{b} + \mathbf{e}_i$ for some pair $i, j \in [k]$; or in other words if \mathbf{b}, \mathbf{b}' are nearest neighbors in $V_{k,q-1}$ (differ by ± 1 in exactly two coordinates).

Consider a labeling $\ell : V_{k,q} \rightarrow [k]$. For each $i \in [k]$, let C_i denote the set of points in $V_{k,q-1}$ representing the monochromatic hyperedges in color i ,

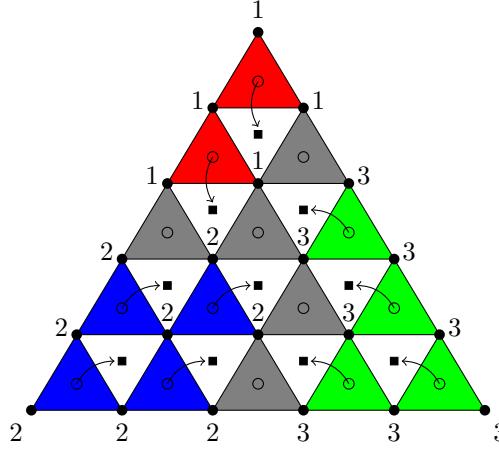
$$C_i = \{\mathbf{b} \in V_{k,q-1} : \forall \mathbf{v} \in e(\mathbf{b}); \ell(\mathbf{v}) = i\}.$$

Define an injective mapping $\phi_i : C_i \rightarrow V_{k,q-2}$ as follows:

$$\phi_i(\mathbf{b}) = \mathbf{b} - \mathbf{e}_i.$$

The image is indeed in $V_{k,q-2}$: if $\mathbf{b} \in C_i$, we have $b_i > 0$, or else $e(\mathbf{b})$ would contain a vertex \mathbf{a} such that $a_i = 0$ and hence $e(\mathbf{b})$ could not be monochromatic in color i . Therefore, $\mathbf{b} - \mathbf{e}_i \in \mathbb{Z}_+^k$ and $(\mathbf{b} - \mathbf{e}_i) \cdot \mathbf{1} = q - 2$ which means $\mathbf{b} - \mathbf{e}_i \in V_{k,q-2}$. (Here, $\mathbf{1}$ denotes the all-1's vector.)

Figure 4: The mappings $\phi_i : C_i \rightarrow V_{k,q-2}$. The hyperedges are represented by the empty circles; C_i is the subset of them monochromatic in color i . The black squares represent $V_{k,q-2}$; note that each point in $V_{k,q-2}$ is the image of at most one monochromatic hyperedge.



Further, we claim that $\phi_i[C_i] \cap \phi_j[C_j] = \emptyset$ for every $i \neq j$. If not, there would be $\mathbf{b} \in C_i$ and $\mathbf{b}' \in C_j$ such that $\mathbf{b} - \mathbf{e}_i = \mathbf{b}' - \mathbf{e}_j$. Then, the point $\mathbf{a} = \mathbf{b} + \mathbf{e}_j = \mathbf{b}' + \mathbf{e}_i$ would be an element of both the hyperedge $e(\mathbf{b})$ and the hyperedge $e(\mathbf{b}')$. This contradicts the assumption that $e(\mathbf{b})$ is monochromatic in color i and $e(\mathbf{b}')$ is monochromatic in color j . So the sets $\phi_i[C_i]$ are pairwise disjoint subsets of $V_{k,q-2}$. By the definition of ϕ_i , we clearly have $|\phi_i[C_i]| = |C_i|$. We conclude that the total number of monochromatic hyperedges is

$$\sum_{i=1}^k |C_i| = \sum_{i=1}^k |\phi_i[C_i]| \leq |V_{k,q-2}|.$$

The total number of hyperedges is $|E_{k,q}| = |V_{k,q-1}|$. Considering that $|V_{k,q}| = \binom{q+k-1}{k-1}$ (the number of partitions of q into a sum of k nonnegative integers), we obtain that the number of non-monochromatic hyperedges is

$$|E_{k,q}| - \sum_{i=1}^k |C_i| \geq |V_{k,q-1}| - |V_{k,q-2}| = \binom{q+k-2}{k-1} - \binom{q+k-3}{k-1} = \binom{q+k-3}{k-2}.$$

□

5 A labeling of $H_{k,q}$ with at most 4 colors on each hyperedge

We recall that Sperner's lemma states that any Sperner-admissible labeling of a subdivision of the simplex must contain a simplex with all k colors. The hypergraph $H_{k,q}$ defined in Section 2.1 is not a subdivision since it covers only a subset of the large simplex. It is easy to see that the conclusion of Sperner's lemma does not hold for $H_{k,q}$ — for example for $k = 3$, we can label a 2-dimensional triangulation so that exactly one triangle has 3 different colors, and this triangle is not in $E_{3,q}$. (See Figure 2.) Hence, each triangle in $E_{3,q}$ has at most 2 colors. By an extension of this argument, we can label $H_{k,q}$ so that each hyperedge in $E_{k,q}$ contains at most $k - 1$ colors. The question we ask in this section is, what is the minimum c such that there is a Sperner-admissible labeling with at most c different colors on each hyperedge in $E_{k,q}$? We prove the following result.

Proposition 2. *For any $k \geq 4$ and $q \geq k^2$, there is a Sperner-admissible labeling of $H_{k,q} = (V_{k,q}, E_{k,q})$ such that every hyperedge in $E_{k,q}$ contains at most 4 different colors.*

We note that this statement is not true for $q = 1$ and $k > 4$ (since $E_{k,1}$ consists of a single simplex which has k different colors). We have not identified the optimal lower bound on q that allows our statement to hold. Also, the statement could possibly hold with 2 or 3 colors instead of 4; the number 4 is just an artifact of our proof and we have no reason to believe that it is tight.

The intuition behind our construction is as follows: We want to label the vertices so that the number of different colors on each hyperedge is small. A natural choice is to label each vertex \mathbf{v} by its maximum-value coordinate. However, this does not work since a hyperedge in the center of the simplex may receive all k colors. The problem is that this labeling is possibly very sensitive to small changes in \mathbf{v} . A more “robust” labeling is one where we select a subset of “top coordinates” and choose one among them according to another rule. This rule should be such that incrementing the coordinates one at a time does not change the label too many times. One such rule that works well is described below.

Proof. We define a labeling $\ell : V_{k,q} \rightarrow [k]$ as follows:

- Given $\mathbf{a} \in V_{k,q}$, let $\pi : [k] \rightarrow [k]$ be a permutation such that $a_{\pi(1)} \geq a_{\pi(2)} \geq \dots \geq a_{\pi(k)}$ (and if $a_{\pi(i)} = a_{\pi(i+1)}$, we order π so that $\pi(i) < \pi(i+1)$).
- Define $t(\mathbf{a})$ to be the maximum $t \in [k]$ such that $\forall 1 \leq j \leq t, a_{\pi(j)} \geq k - j + 1$. We define the “Top coordinates” of \mathbf{a} to be $Top(\mathbf{a}) = (\pi(1), \dots, \pi(t(\mathbf{a})))$ (an ordered set).
- We define the label of \mathbf{a} to be $\ell(\mathbf{a}) = \pi(t(\mathbf{a}))$, the index of the “last Top coordinate”.

First, we verify that this is a well-defined Sperner-admissible labeling. Since $\sum_{i=1}^k a_i = q \geq k^2$, we have $a_{\pi(1)} = \max a_i \geq k$ and hence $1 \leq t(\mathbf{a}) \leq k$. For each $\mathbf{a} \in V_{k,q}$, we have: $a_{\ell(\mathbf{a})} = a_{\pi(t(\mathbf{a}))} \geq k - t(\mathbf{a}) + 1 > 0$, since $t(\mathbf{a}) \leq k$. Therefore, ℓ is Sperner-admissible.

Now, consider a hyperedge $e(\mathbf{b}) = (\mathbf{b} + \mathbf{e}_1, \mathbf{b} + \mathbf{e}_2, \dots, \mathbf{b} + \mathbf{e}_k)$ where $\mathbf{b} \geq 0, \sum_{i=1}^k b_i = q-1$. We claim that $\ell(\mathbf{b} + \mathbf{e}_i)$ attains at most 4 different values for $i = 1, \dots, k$. Without loss of generality, assume that $b_1 \geq b_2 \geq \dots \geq b_k$. Define ℓ^* to be the label assigned to \mathbf{b} by our construction (note that \mathbf{b} is not a vertex in $V_{k,q}$ but we can still apply our definition): ℓ^* is the maximum value in $[k]$ such that for all $1 \leq j \leq \ell^*$, $b_j \geq k - j + 1$. Hence, we have $Top(\mathbf{b}) = \{1, 2, \dots, \ell^*\}$.

Let $i \in [k]$, $\mathbf{a} = \mathbf{b} + \mathbf{e}_i$, and let π be the permutation such that $a_{\pi(1)} \geq \dots \geq a_{\pi(k)}$ as above. (Recall that for \mathbf{b} , we assumed that the respective permutation is the identity.) We consider the following cases:

- If $1 \leq i < \ell^*$, then we claim that $\ell(\mathbf{a}) = \ell(\mathbf{b} + \mathbf{e}_i) = \ell(\mathbf{b}) = \ell^*$. In the rule for selecting $t(\mathbf{a})$, one of the first $\ell^* - 1$ coordinates has been incremented compared to \mathbf{b} , which possibly pushes i forward in the ordering of the Top coordinates. However, the other coordinates remain unchanged, the condition $a_{\pi(j)} \geq k - j + 1$ is still satisfied for $1 \leq j \leq \ell^*$, and $Top(\mathbf{a}) = Top(\mathbf{b})$. In particular ℓ^* is still the last coordinate included in $Top(\mathbf{a})$ and hence $\ell(\mathbf{a}) = \ell^*$.
- If $i = \ell^*$, then $\ell(\mathbf{a}) = \ell(\mathbf{b} + \mathbf{e}_{\ell^*})$ is still one of the coordinates in $Top(\mathbf{b})$, possibly different from ℓ^* (due to a change in order, although we still have $Top(\mathbf{a}) = Top(\mathbf{b})$) — let us call this label ℓ_2^* .
- If $\ell^* < i \leq k$, then it is possible that in $\mathbf{a} = \mathbf{b} + \mathbf{e}_i$, we obtain additional Top coordinates ($Top(\mathbf{a}) \supset Top(\mathbf{b})$). It could be $a_i = b_i + 1$ itself which is now included among the Top coordinates, and possibly additional coordinates that already satisfied the condition $b_j \geq k - j + 1$ but were not selected due to the condition being false for b_{ℓ^*+1} . If this does not happen and we have $Top(\mathbf{a}) = Top(\mathbf{b})$, the label of \mathbf{a} is still $\ell(\mathbf{a}) = \ell^*$ (because the ordering of the Top coordinates remains the same).

Assume now that $Top(\mathbf{a})$ has additional coordinates beyond $Top(\mathbf{b})$. By the definition of ℓ^* , we have $b_{\ell^*} \geq k - \ell^* + 1$, and for each $j > \ell^*$, we have $b_j < k - \ell^*$; otherwise j would have been still chosen in $Top(\mathbf{b})$. For $Top(\mathbf{a}) = Top(\mathbf{b} + \mathbf{e}_i)$ to grow beyond $Top(\mathbf{b})$, a_i must become the $(\ell^* + 1)$ -largest coordinate and satisfy $a_i \geq k - \ell^*$. The only way this can happen is that $b_i = k - \ell^* - 1$ and hence $a_i = b_i + 1 = k - \ell^*$. In this case, a_i is the maximum coordinate among $\{a_j : j > \ell^*\}$, and still smaller than a_{ℓ^*} . Therefore, i will be included in $Top(\mathbf{a})$. Now, $Top(\mathbf{a})$ may grow further. However, note that the construction of $Top(\mathbf{a})$ will proceed in the same way for every $\mathbf{a} = \mathbf{b} + \mathbf{e}_i$ such that $b_i = k - \ell^* - 1$. This is because all the coordinates equal to $k - \ell^* - 1$ will be certainly included in $Top(\mathbf{a})$, and coordinates smaller than $k - \ell^* - 1$ remain the same in each of these cases (equal to the coordinates of \mathbf{b}). Therefore, the set $Top(\mathbf{a})$ will be the same in all these cases; let us call this set Top_+ .

The label assigned to $\mathbf{a} = \mathbf{b} + \mathbf{e}_i$ is the index of the last coordinate included in $Top_+ = Top(\mathbf{a})$. Since Top_+ is the same whenever $Top(\mathbf{a}) \neq Top(\mathbf{b})$, the label of \mathbf{a} will be the coordinate j^* minimizing b_j (and maximizing j to break ties) among all $j \in Top_+$, unless $j^* = i$ in which case the last included coordinate might be another one. This gives potentially two additional colors, let us call them ℓ_3^*, ℓ_4^* , that are assigned to $\mathbf{a} = \mathbf{b} + \mathbf{e}_i$ for all $i > \ell^*$ where $b_i = k - \ell^* - 1$. For other choices of $i > \ell^*$, we have $Top(\mathbf{b} + \mathbf{e}_i) = Top(\mathbf{b})$ and the label assigned to $\mathbf{b} + \mathbf{e}_i$ is $\ell(\mathbf{b} + \mathbf{e}_i) = \ell^*$.

To summarize, all the colors that appear in the labeling of $e(\mathbf{b})$ are included in $\{\ell^*, \ell_2^*, \ell_3^*, \ell_4^*\}$. \square

6 Boundary-minimizing partitioning of the simplex

Let us turn now to a geometric variant of Proposition 1. We recall that Sperner's Lemma has a geometric variant known as the Knaster-Kuratowski-Mazurkiewicz Lemma [7]:

Consider a covering of the simplex Δ_k by closed sets A_1, \dots, A_k such that each point $\mathbf{x} \in \Delta_k$ is contained in some set A_i such that $x_i > 0$. Then $\bigcap_{i=1}^k A_i \neq \emptyset$.

Here we consider a similar setup, but instead of the intersection of all sets, we are interested in the measure of the boundaries between pairs of adjacent sets. To avoid technicalities, let us assume that the A_i 's are closed, disjoint except on the boundary, and each A_i is disjoint from the face $\{\mathbf{x} \in \Delta_k : x_i = 0\}$.

Definition 1. A Sperner-admissible partition of Δ_k is a k -tuple of closed sets (A_1, \dots, A_k) such that

- $\bigcup_{i=1}^k A_i = \Delta_k$,
- A_1, \dots, A_k are disjoint except on their boundary,
- $x_i > 0$ for every $\mathbf{x} \in A_i$.

We call the union of pairwise boundaries $\bigcup_{i \neq j} (A_i \cap A_j)$ the separating set.

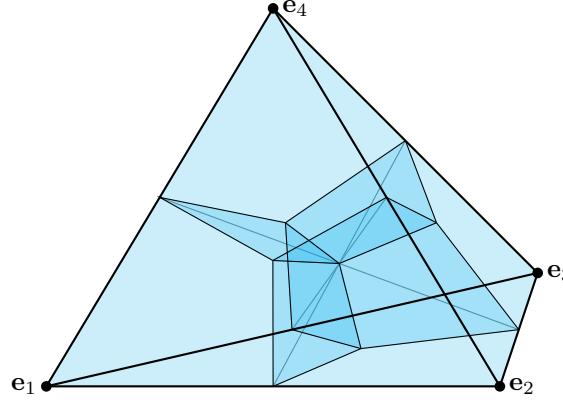
The question we ask here is, in analogy with Proposition 1, what is the Sperner-admissible partition with the separating set of minimum measure? A candidate partition is depicted in Figure 5, where A_i is the set of all points in Δ_k for whom \mathbf{e}_i is the closest vertex. We call this the *Voronoi partition*.

We prove that for the regular simplex Δ_k this is indeed the optimal partition (along with other, similar configurations). In the following, we denote by μ_k the usual Lebesgue measure on \mathbb{R}^k , and by μ_ℓ ($\ell < k$) the ℓ -dimensional Minkowski content.

Definition 2. For $A \subset \mathbb{R}^k$, the ℓ -dimensional Minkowski content is (if the limit exists)

$$\mu_\ell(A) = \lim_{\epsilon \rightarrow 0^+} \frac{\mu_k(A_\epsilon)}{\alpha_{k-\ell} \epsilon^{k-\ell}}$$

Figure 5: The Voronoi partition of a simplex.



where $A_\epsilon = \{\mathbf{y} \in \mathbb{R}^k : \exists \mathbf{x} \in A, \|\mathbf{x} - \mathbf{y}\| \leq \epsilon\}$ is the ϵ -neighborhood of A and $\alpha_{k-\ell}$ is the volume of a unit ball in $\mathbb{R}^{k-\ell}$. We also define $\mu_\ell^+(A)$ to be the upper limit and $\mu_\ell^-(A)$ the lower limit of the expression above.

We remark that for ℓ -rectifiable sets (polyhedral faces, smooth surfaces, etc.) the notion of Minkowski content coincides with that of Hausdorff measure (under suitable normalization).

Theorem 1. For every Sperner-admissible partition (A_1, \dots, A_k) of Δ_k ,

$$\mu_{k-2}^-\left(\bigcup_{i \neq j}(A_i \cap A_j)\right) \geq \frac{k-1}{\sqrt{2}} \mu_{k-1}(\Delta_k)$$

and the Voronoi partition achieves this with equality.

First, let us analyze the Voronoi partition and more generally the following kind of partition.

Lemma 2. For any \mathbf{z} in the interior of Δ_k , the partition (A_1^z, \dots, A_k^z) where

$$A_i^z = \{\mathbf{x} \in \Delta_k : x_i - z_i = \max_{1 \leq j \leq k} (x_j - z_j)\}$$

satisfies

$$\mu_{k-2}\left(\bigcup_{i \neq j}(A_i^z \cap A_j^z)\right) = \frac{k-1}{\sqrt{2}} \mu_{k-1}(\Delta_k) = \frac{1}{(k-2)!} \sqrt{\frac{k}{2}}.$$

We call this kind of partition ‘‘Voronoi-type’’.² We note that for $\mathbf{z} = (\frac{1}{k}, \frac{1}{k}, \dots, \frac{1}{k})$ we obtain the Voronoi partition in Figure 5. Other choices of

²We note that these partitions are also known as ‘‘power diagrams’’.

\mathbf{z} correspond to similar configurations where all the colors meet at the point \mathbf{z} . Note that \mathbf{z} is the “rainbow point” guaranteed by the Knaster-Kuratowski-Mazurkiewicz Lemma.

Proof. First let us compute some basic quantities that we will need. The sides of our simplex Δ_k have length $\sqrt{2}$. Denote by h_k the height of Δ_k , that is the distance of any vertex from the opposite facet. We have

$$h_k = \left\| (1, 0, \dots, 0) - (0, \frac{1}{k-1}, \dots, \frac{1}{k-1}) \right\| = \sqrt{1 + (k-1) \cdot \frac{1}{(k-1)^2}} = \sqrt{\frac{k}{k-1}}.$$

The volume of the simplex can be computed inductively as follows; we have $\mu_1(\Delta_2) = \sqrt{2}$, and $\mu_k(\Delta_{k+1}) = \frac{1}{k} h_{k+1} \cdot \mu_{k-1}(\Delta_k)$. This implies

$$\mu_{k-1}(\Delta_k) = \frac{\sqrt{k}}{(k-1)!}.$$

Now let us compute the measure of the separating set for the partition (A_1^z, \dots, A_k^z) defined above, by induction. The separating set can be described explicitly as

$$\bigcup_{i \neq j} (A_i^z \cap A_j^z) = \{ \mathbf{x} \in \Delta_k : \exists i \neq j, x_i - z_i = x_j - z_j = \max_{1 \leq \ell \leq k} x_\ell - z_\ell \}.$$

For $k = 2$, $A_1^z \cap A_2^z$ is just a single point, and $\mu_0(A_1^z \cap A_2^z) = 1$. For $k \geq 3$, denote by S the separating set for (A_1^z, \dots, A_k^z) and define $S_i = S \cap \text{conv}(\{\mathbf{e}_j : j \neq i\})$, the separating set restricted to the facet opposite vertex \mathbf{e}_i . Since S_i is a Voronoi-type separating set for Δ_{k-1} , by induction we assume that $\mu_{k-3}(S_i) = \frac{1}{(k-3)!} \sqrt{\frac{k-1}{2}}$. The separating set S can be written as $S = \bigcup_{i=1}^k \text{conv}(S_i \cup \{\mathbf{z}\})$, see Figure 5. Denote by h'_i the distance of \mathbf{z} from the facet containing S_i . By the pyramid formula in dimension $k-2$,

$$\mu_{k-2}(\text{conv}(S_i \cup \{\mathbf{z}\})) = \frac{1}{k-2} h'_i \mu_{k-3}(S_i) = \frac{h'_i}{(k-2)!} \sqrt{\frac{k-1}{2}}.$$

By a simple calculation, $\sum_{i=1}^k h'_i = h_k = \sqrt{\frac{k}{k-1}}$. The sets $\text{conv}(S_i \cup \{\mathbf{z}\})$ are disjoint except for lower-dimensional intersections. Hence,

$$\mu_{k-2}(S) = \sum_{i=1}^k \mu_{k-2}(\text{conv}(S_i \cup \{\mathbf{z}\})) = \sum_{i=1}^k \frac{h'_i}{(k-2)!} \sqrt{\frac{k-1}{2}} = \frac{1}{(k-2)!} \sqrt{\frac{k}{2}}.$$

□

Thus the proof of Theorem 1 will be complete if we prove the following bound.

Lemma 3. *For every Sperner-admissible partition (A_1, \dots, A_k) of Δ_k ,*

$$\mu_{k-2}^-(\bigcup_{i \neq j}(A_i \cap A_j)) \geq \frac{k-1}{\sqrt{2}}\mu_{k-1}(\Delta_k) = \frac{1}{(k-2)!}\sqrt{\frac{k}{2}}.$$

Proof. We pursue an approach similar to the proof of Proposition 1, with some additional technicalities. The high-level approach is to shrink the sets A_i somewhat, by excluding a small neighborhood of the separating set. This creates a buffer zone between the shrunk sets A'_i (yellow in Figure 6) whose measure corresponds to the measure of the separating set. Since we have this extra space, we are able to push the sets A'_i closer together and obtain sets A''_i that fit inside a slightly smaller simplex. The difference between the volume of this simplex and the original one gives a bound on the measure of the separating set.

First, let $\epsilon_0 = \inf_{i \in [k], \mathbf{x} \in A_i} x_i$. Recall that $x_i > 0$ for each $\mathbf{x} \in A_i$, and moreover each A_i is closed. Hence $\epsilon_0 > 0$.

Define $S = \bigcup_{i \neq j}(A_i \cap A_j)$, the separating set whose measure we are trying to lower-bound. Fix $\epsilon \in (0, \frac{1}{2}\epsilon_0)$ (eventually we will let $\epsilon \rightarrow 0$) and define S_ϵ as the ϵ -neighborhood of S ,

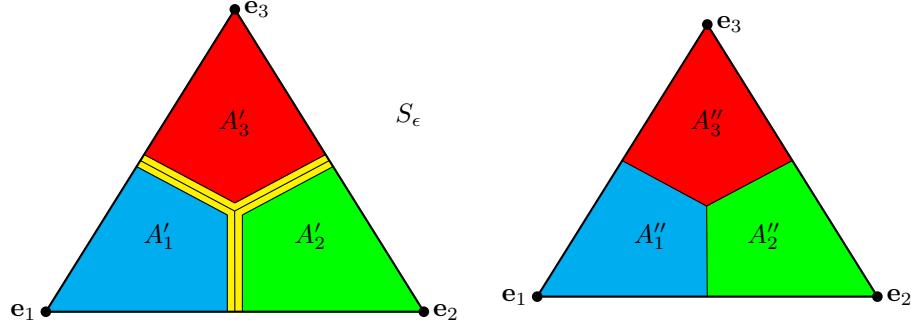
$$S_\epsilon = \{\mathbf{x} \in \Delta_k : \exists \mathbf{y} \in S, \|\mathbf{x} - \mathbf{y}\| \leq \epsilon\}.$$

We define subsets $A'_i \subset A_i$ as follows:

$$A'_i = A_i \setminus S_\epsilon.$$

Thus we have $\bigcup_{i=1}^k A'_i = \Delta_k \setminus S_\epsilon$. Also, the sets A'_i are clearly disjoint (see Figure 6).

Figure 6: The construction of A'_i and A''_i .



Next, we set $\epsilon' = \epsilon\sqrt{2}$ and define

$$A''_i = A'_i - \epsilon' \mathbf{e}_i = \{\mathbf{x} - \epsilon' \mathbf{e}_i : \mathbf{x} \in A'_i\}.$$

Thus A''_i is a shifted copy of A'_i , where we push A'_i slightly away from vertex \mathbf{e}_i . The sets A''_i live in the hyperplane $\sum_{i=1}^k x_i = 1 - \epsilon'$ rather than $\sum_{i=1}^k x_i = 1$.

We claim that the sets A''_i are still disjoint: Suppose that $A''_i \cap A''_j = (A'_i - \epsilon' \mathbf{e}_i) \cap (A'_j - \epsilon' \mathbf{e}_j) \neq \emptyset$. This would mean that there are points $\mathbf{x} \in A'_i, \mathbf{y} \in A'_j$ such that $\mathbf{x} - \epsilon' \mathbf{e}_i = \mathbf{y} - \epsilon' \mathbf{e}_j$. In other words, $\|\mathbf{x} - \mathbf{y}\| = \epsilon' \|\mathbf{e}_i - \mathbf{e}_j\| = \epsilon' \sqrt{2} = 2\epsilon$. Take the midpoint $\frac{1}{2}(\mathbf{x} + \mathbf{y})$: this point is in the simplex Δ_k (by convexity), and hence it is in some set A_ℓ , where either $\ell \neq i$ or $\ell \neq j$ (possibly both). Assume without loss of generality that $\ell \neq i$. Then by the closedness of A_i and A_ℓ , between \mathbf{x} and $\frac{1}{2}(\mathbf{x} + \mathbf{y})$ there exists a point $\mathbf{x}' \in A_i \cap A_\ell$. We get a contradiction, because $\|\mathbf{x}' - \mathbf{x}\| \leq \epsilon$ and so \mathbf{x} would not be included in A'_i .

We also observe that $A''_i \subseteq (1 - \epsilon') \cdot \Delta_k = \{\mathbf{x} \geq 0 : \sum_{i=1}^k x_i = 1 - \epsilon'\}$. This is because for every $\mathbf{x} \in A''_i$, we have $\mathbf{y} \in A'_i$ such that $\mathbf{x} = \mathbf{y} - \epsilon' \mathbf{e}_i$. By assumption, $y_i \geq \epsilon_0 > \epsilon'$, and $\mathbf{y} \in \Delta_k$. Therefore $x_i = y_i - \epsilon' > 0$ and $\sum_{i=1}^k x_i = 1 - \epsilon'$. We conclude that A''_1, \dots, A''_k are disjoint subsets of $(1 - \epsilon') \cdot \Delta_k$, obtained by an isometry from A'_1, \dots, A'_k and therefore

$$\sum_{i=1}^k \mu_{k-1}(A'_i) = \sum_{i=1}^k \mu_{k-1}(A''_i) \leq (1 - \epsilon')^{k-1} \mu_{k-1}(\Delta_k).$$

Recall that A'_1, \dots, A'_k are also disjoint and $\bigcup_{i=1}^k A'_i = \Delta_k \setminus S_\epsilon$. Therefore,

$$\mu_{k-1}(S_\epsilon) = \mu_{k-1}(\Delta_k) - \sum_{i=1}^k \mu_{k-1}(A'_i) \geq (1 - (1 - \epsilon')^{k-1}) \mu_{k-1}(\Delta_k).$$

By the definition of Minkowski content, we have

$$\begin{aligned} \mu_{k-2}^-(S) &= \liminf_{\epsilon \rightarrow 0^+} \frac{\mu_{k-1}(S_\epsilon)}{2\epsilon} \geq \liminf_{\epsilon \rightarrow 0^+} \frac{1 - (1 - \epsilon')^{k-1}}{2\epsilon} \mu_{k-1}(\Delta_k) \\ &= \lim_{\epsilon \rightarrow 0^+} \frac{1 - (1 - \epsilon\sqrt{2})^{k-1}}{2\epsilon} \mu_{k-1}(\Delta_k) = \frac{k-1}{\sqrt{2}} \mu_{k-1}(\Delta_k). \end{aligned}$$

□

Alternative proof of optimality. Here we give an alternative proof that the Voronoi partition has a separating set of minimum Minkowski content, avoiding an explicit computation of its volume.

Proof of Lemma 2. Let us consider the Voronoi partition (A_1^*, \dots, A_k^*) (the proof for a general \mathbf{z} is similar). We argue that the proof of Lemma 3 is tight for this partition. As in the proof of Lemma 3, we define $S^* = \bigcup_{i \neq j} (A_i^* \cap A_j^*)$, $S_\epsilon^* = \{\mathbf{x} \in \Delta_k : \exists \mathbf{y} \in S, \|\mathbf{x} - \mathbf{y}\| \leq \epsilon\}$, $A'_i = A_i^* \setminus S_\epsilon^*$ and $A''_i = A'_i - \epsilon' \mathbf{e}_i$, $\epsilon' = \epsilon\sqrt{2}$. In the case of the Voronoi partition, these sets are explicitly described as follows:

- $A_i^* = \{\mathbf{x} \in \Delta_k : x_i = \max_{\ell \in [k]} x_\ell\}$,
- $S^* = \{\mathbf{x} \in \Delta_k : \exists i \neq j, x_i = x_j = \max_{\ell \in [k]} x_\ell\}$,

- $A'_i = \{\mathbf{x} \in \Delta_k : x_i > \epsilon' + \max_{\ell \neq i} x_\ell\},$
- $A''_i = \{\mathbf{x} - \epsilon' \mathbf{e}_i : \mathbf{x} \in \Delta_k, x_i > \epsilon' + \max_{\ell \neq i} x_\ell\}.$

The description of A'_i is valid because for $\mathbf{x} \in A'_i$, it is possible to find a point in S^* within distance ϵ of \mathbf{x} if and only if the maximum coordinate x_i is within $\epsilon' = \epsilon\sqrt{2}$ of the second largest coordinate — then we can replace the two largest coordinates by their average and obtain a point in S^* . The description of A''_i follows by definition.

Consider now the scaled-down simplex $(1-\epsilon') \cdot \Delta_k$. By the proof of Lemma 3, the sets A''_i are disjoint subsets of $(1-\epsilon') \cdot \Delta_k$. We show that in this case, we actually have $\sum_{i=1}^k \mu_{k-1}(A''_i) = \mu_{k-1}((1-\epsilon')\Delta_k)$. This is because for any point $\mathbf{x}' \in (1-\epsilon') \cdot \Delta_k$, if the maximum coordinate x'_i of \mathbf{x}' is unique then $\mathbf{x} = \mathbf{x}' + \epsilon' \mathbf{e}_i$ is a point in Δ_k such that $x_i > \epsilon' + \max_{\ell \neq i} x_\ell$. Therefore, $\mathbf{x} \in A'_i$ which implies that $\mathbf{x}' \in A''_i$. The points $\mathbf{x}' \in (1-\epsilon') \cdot \Delta_k$ whose maximum coordinate is not unique form a set of $(k-1)$ -dimensional measure zero. Therefore, $(1-\epsilon')\Delta_k$ is covered by $\bigcup_{i=1}^k A''_i$ up to a set of measure zero, and $\sum_{i=1}^k \mu_{k-1}(A'_i) = \sum_{i=1}^k \mu_{k-1}(A''_i) = \mu_{k-1}((1-\epsilon')\Delta_k) = (1-\epsilon\sqrt{2})^{k-1} \mu_{k-1}(\Delta_k)$. We also have $S_\epsilon^* = \Delta_k \setminus \bigcup_{i=1}^k A'_i$. This shows that all the inequalities in the proof of Lemma 3 are tight and the Minkowski content of the separating set S^* is exactly

$$\mu_{k-2}(S^*) = \lim_{\epsilon \rightarrow 0^+} \frac{\mu_{k-1}(S_\epsilon^*)}{2\epsilon} = \lim_{\epsilon \rightarrow 0^+} \frac{1 - (1 - \epsilon\sqrt{2})^{k-1}}{2\epsilon} \mu_{k-1}(\Delta_k) = \frac{k-1}{\sqrt{2}} \mu_{k-1}(\Delta_k).$$

□

7 Discussion and open questions

Sperner's Lemma extends to general polytopes in the following sense [8]:

For any coloring of a triangulation of a d -dimensional polytope with n vertices by n colors, such that each point on a face $F = \text{conv}(\{\mathbf{v}_i : i \in A\})$ must be colored with a color in A , there are at least $n-d$ full-dimensional simplices with $d+1$ distinct colors.

It is natural ask whether our results also extend to general polytopes.

Possible extensions to polytopes. Consider the example of P being a square (Figure 7). The Voronoi partition (A_1, A_2, A_3, A_4) is not optimal with respect to the total length of the separating set. The separating set of the Voronoi partition has total length 2, whereas total length arbitrarily close to $\sqrt{2}$ is achieved by the partition (B_1, B_2, B_3, B_4) .

In general, we do not know what the partition minimizing $\mu(\bigcup_{i \neq j} (A_i \cap A_j))$ looks like, even in the case of a non-regular simplex. We believe that the separating set should still be polyhedral (piecewise linear) for an optimal Sperner-admissible partition of any polytope.

Figure 7: Two partitions of a square.

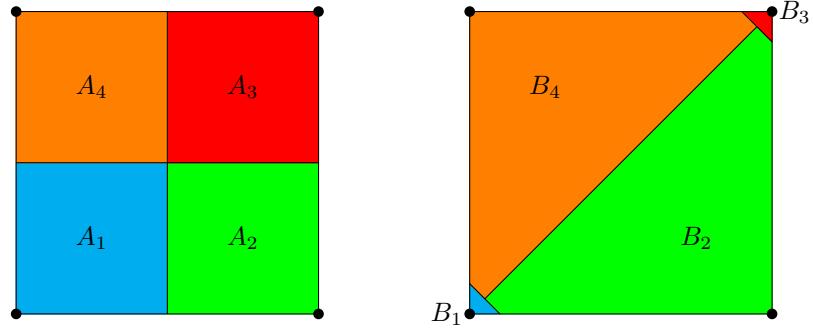
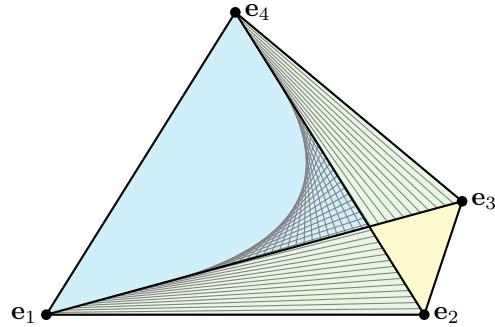


Figure 8: An optimal partition between two pairs of faces of the tetrahedron.



We remark that depending on the coloring conditions on the surface of the polyhedron, the optimal separating set may be non-linear: For a tetrahedron, the optimal partition that separates the pair of faces $\text{conv}(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3) \cup \text{conv}(\mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4)$ from $\text{conv}(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_4) \cup \text{conv}(\mathbf{e}_1, \mathbf{e}_3, \mathbf{e}_4)$, is the minimal surface whose boundary is the non-planar 4-gon $\mathbf{e}_1\mathbf{-}\mathbf{e}_2\mathbf{-}\mathbf{e}_4\mathbf{-}\mathbf{e}_3$. This is a saddle-shaped surface (see Figure 8).

Other open questions. We have proved several results about colorings of the simplex. Our first result (Proposition 1) can be viewed as being at the opposite end of the spectrum from Sperner’s Lemma: Instead of the existence of a rainbow cell, we proved a lower bound on the number of non-monochromatic cells. Due to the motivating application of [4], we considered a special hypergraph embedded in the simplex rather than a full subdivision. A natural question is whether an analogous statement holds for simplicial subdivisions.

More generally, we might “interpolate” between Sperner’s Lemma and our result, and ask: How many cells must contain at least j colors? It is clear that these questions depend on the structure of the subdivision, and some assumption of regularity would be needed to obtain a general result. Similarly, we may

ask, for Sperner-admissible geometric partitions of the simplex, what is the minimum possible volume of the set where at least j colors meet? Furthermore, as we discussed above, are there generalizations of these statements to other polytopes?

Another question is, what is the Sperner-admissible labeling of the Simplex-Lattice Hypergraph $H_{k,q}$ (defined in Section 2) minimizing the maximum number of colors on a hyperedge? We have proved that 4 colors suffice but it is possible that 2 colors are enough (see Proposition 2). Is there a Sperner-admissible labeling of the hypergraph $H_{k,q}$, for sufficiently large q , such that each hyperedge uses at most 2 colors?

Finally, we remark that Proposition 2 does not have a continuous counterpart for geometric partitions: As we discussed earlier, for any Sperner-admissible partition of a simplex there is a point where all the parts meet, by the Knaster-Kuratowski-Mazurkiewicz Lemma [7].

Acknowledgement. *The second author is indebted in many ways to Jirka Matoušek, who introduced him to Sperner’s Lemma in an undergraduate course at Charles University a long time ago.*

References

- [1] Chandra Chekuri and Alina Ene. Submodular cost allocation problem and applications. In *Proc. of ICALP*, 354–366, 2011.
- [2] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *J. ACM* 56:3, 2009.
- [3] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM J. Computing* 39:1, 195–259, 2009.
- [4] Alina Ene and Jan Vondrák. Hardness of submodular cost allocation: Lattice matching and a simplex coloring conjecture. In *Proc. of APPROX*, 144–159, 2014.
- [5] Alina Ene, Jan Vondrák, and Yi Wu. Local distribution and the symmetry gap: Approximability of multiway partitioning problems. In *Proc. of ACM-SIAM SODA*, 306–325, 2013.
- [6] Jon Kleinberg and Éva Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and Markov random fields. *Journal of the ACM*, 49:5, 616–639, 2002.
- [7] B. Knaster, C. Kuratowski, and S. Mazurkiewicz. Ein Beweis des Fixpunktsatzes für n -dimensionale Simplexe. *Fundamenta Mathematicae*, 14:1, 132–137, 1929.

- [8] Jesus de Loera, Elisha Peterson, and Francis Edward Su. A polytopal generalization of Sperner’s Lemma. *J. Combinatorial Theory A* 100, 1–26, 2002.
- [9] Rajsekhar Manokaran, Joseph Naor, Prasad Raghavendra, and Roy Schwartz. SDP gaps and UGC hardness for multiway cut, 0-extension, and metric labeling. *Proc. of ACM STOC*, 11–20, 2008.
- [10] Maryam Mirzakhani and Jan Vondrák. Sperner’s colorings, hypergraph labeling problems and fair division. *Proc. of ACM-SIAM SODA*, 873–886, 2015.
- [11] John Nash. Noncooperative games. *Annals of Mathematics* 54, 289–295, 1951.
- [12] Emanuel Sperner. Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes. *Math. Sem. Univ. Hamburg*, 6:265–272, 1928.
- [13] Francis Edward Su. Rental harmony: Sperner’s lemma in fair division. *Amer. Math. Monthly* 106, 930–942, 1999.