The Borsuk number of a graph*

José Cáceres†1, Delia Garijo‡2, Alberto Márquez§2, and Rodrigo I. Silveira¶3

[1]Dpto. de Matemáticas, Universidad de Almería, Spain
[2]Dpto. de Matemática Aplicada I, Universidad de Sevilla, Spain
[3]Dept. de Matemàtiques, Universitat Politècnica de Catalunya, Spain

Abstract

The *Borsuk problem* asks for the smallest number such that any bounded set in n-dimensional space can be cut into that many subsets with smaller diameter. It is a classical problem in combinatorial geometry that has been subject of much attention over the years, and research on variants of the problem continues nowadays in a plethora of directions. In this work, we propose a formulation of the problem in the context of graphs. Depending on how the graph is partitioned, we consider two different settings dealing either with the usual notion of diameter in abstract graphs, or with the so-called *continuous diameter* for the locus of plane geometric graphs. We present a complexity result, exact computations and upper bounds on the parameters associated to the problem.

1 Introduction

In 1933, Borsuk posed the question of whether every bounded set X in $\mathbb{R}^d$ could be partitioned into $d+1$ closed (sub)sets each with diameter smaller than that of X [1]. In this context, the diameter is defined as the maximum of the distances between two points in the set, under the Euclidean metric. This leads to the concept of *Borsuk number*. For a set $X \subset \mathbb{R}^d$, the Borsuk number $b(X)$ is the smallest number such that X can be partitioned into $b(X)$ subsets, each with diameter smaller than X. Borsuk's question can be thus stated as whether $b(X) \leq d+1$, for any bounded $X \subset \mathbb{R}^d$. The answer to this question was shown to be positive for $d = 2, 3$ [4, 10], and for general d for centrally symmetric convex bodies [11] and smooth convex bodies [6]. The general answer turned out to be negative, as shown in 1993 by Kahn and Kalai [8]. Since then, researchers have been trying to figure out the smallest dimension for which the partition does not exist, being $d = 64$ the currently best [7]. Many variants of the Borsuk problem have also been studied, see [12] for a recent survey.

We present a formulation of the problem in the context of graphs. Conceptually, we define the *Borsuk number* of a graph as the smallest number $b(G)$ such that G can be partitioned into $b(G)$ subgraphs, each with smaller diameter than the original graph. However, we need to define carefully how a graph can be partitioned. We propose two natural ways to do this, which lead to two variants of the problem: the *discrete* and the *continuous* Borsuk number of a graph. We define these formally in Section 2. Sections 3–5 contain our study on both parameters, encompassing a complexity result, exact computations and upper bounds. Proofs are omitted due to the page limit, although we very briefly explain the key ideas to prove our main results; they are based on an accurate analysis of how shortest paths and distances can change when modifying a graph.

*This research is supported by Grant PID2019-104129GB-I00 funded by MICIU/AEI/10.13039/501100011033.
†Email: jcaceres@ual.es.
‡Email: dgarijo@us.es.
§Email: almar@us.es.
¶Email: rodrigo.silveira@upc.edu.

1.1 Preliminaries

The *distance* between two vertices in an abstract graph G is the length of a shortest path connecting them. The *diameter* of G, denoted by $\text{diam}_d(G)$, is the maximum distance between any two vertices of G. A *plane geometric graph* is an undirected graph $G = (V(G), E(G))$ whose vertices are points in $\mathbb{R}^2$, and whose edges are straight-line segments, connecting pairs of points, that intersect only at their endpoints. Each edge e has a length, $|e|$, equal to the Euclidean distance between its endpoints. The *locus* $\mathcal{G}$ of a plane geometric graph G is the set of all points of the Euclidean plane that are on (the edges of) G. In contrast to (abstract) graphs, in $\mathcal{G}$, there can be an infinite number of pairs of points whose distance is equal to the diameter. Here, the distance between two points is again the length of a shortest path between the points, but now such a path will contain up to two fragments of edges if the points are not vertices. The *diameter* of $\mathcal{G}$ or *continuous diameter* of G, $\text{diam}_c(\mathcal{G})$, is the maximum distance between any two points in $\mathcal{G}$. Two points whose distance attains this value are called *diametral points*, and the shortest paths connecting diametral points are *diametral paths*. Problems dealing with the continuous diameter of a graph, also called *generalized* diameter [3], have received considerable attention recently, see [2, 5]. In the continuous case, we treat G and $\mathcal{G}$, interchangeably, as a closed point set, and assume that the distance between the endpoints of edge e is $|e|$.

2 Definitions of Borsuk number

2.1 Continuous Borsuk number

We consider a plane geometric graph G and partition its locus $\mathcal{G}$ by a sequence of cuts with straight lines. A line ℓ naturally partitions $\mathcal{G}$ into two geometric subgraphs (possibly, one empty). Moreover, to guarantee that the partition by ℓ does not produce a disconnected subgraph, we add to both subgraphs the longest segment in ℓ that has its endpoints in $\mathcal{G} \cup \ell$; this maximal segment is denoted by s. So, actually, the partition gives two subgraphs of $\mathcal{G} \cup \ell$, which are:

$$\mathcal{G}_1 = (\ell^+ \cap \mathcal{G}) \cup s \quad \text{and} \quad \mathcal{G}_2 = (\ell^- \cap \mathcal{G}) \cup s,$$

where ℓ^+ and ℓ^- are, respectively, the open half-planes above and below ℓ (right–left for vertical lines.)

We define the *continuous Borsuk number of G* or *Borsuk number of $\mathcal{G}$*, and denote it by $b_c(\mathcal{G})$, as the minimum cardinality of a partition of $\mathcal{G}$ by lines $\ell_1, \ldots, \ell_k$ into subgraphs $\mathcal{G}_1, \ldots, \mathcal{G}_{k+1}$ such that $\max\{\text{diam}_c(\mathcal{G}_1), \ldots, \text{diam}_c(\mathcal{G}_{k+1})\} < \text{diam}_c(\mathcal{G})$. In order to guarantee that the intersection with a line creates at most two subgraphs, each line ℓ_i is inserted only into one of the existing subgraphs.

Figure 1(a) illustrates this definition for a square. After partitioning the square with a vertical line ℓ (dashed) through its center point, we obtain two subgraphs: all points of $\mathcal{G}$ on each halfplane induced by ℓ, union the maximal segment in ℓ intersecting $\mathcal{G}$. Since this partitions the graph into two subgraphs (of $\mathcal{G} \cup \ell$), each with smaller diameter than that of $\mathcal{G}$, its continuous Borsuk number is two (best possible). However, sometimes more subgraphs are needed. The example in Figure 1(b) shows a 4-star graph, requiring at least two lines, giving continuous Borsuk number three. Note that the continuous diameter can increase when inserting a line, due to distances between points on the graph and new points on the line, see Figure 1(c).

One of the main open questions in this continuous setting is whether $b_c(\mathcal{G})$ can be upper-bounded by a constant. The following proposition gives a linear upper bound on the number of vertices of G.

Proposition 1. *Let $\mathcal{G}$ be the locus of a plane geometric graph with n vertices. Then, $b_c(\mathcal{G}) \leq 2n - 1$.*

Proof. (Brief sketch.) Consider a direction not parallel to any of the edges of $\mathcal{G}$; assume for simplicity that this is the vertical direction. For $\varepsilon > 0$, we split $\mathcal{G}$ by $2n$ vertical lines into $2n-1$ subgraphs; there are two lines associated to each vertex, one to the left and the other to the right, both at distance ε from the vertex. Thus, there are $2n-1$ vertical strips, each containing either only portions of edges of $\mathcal{G}$, or only one vertex and portions of edges. Each resulting graph, $\mathcal{G}_1, \ldots, \mathcal{G}_{2n-1}$, is in one of these

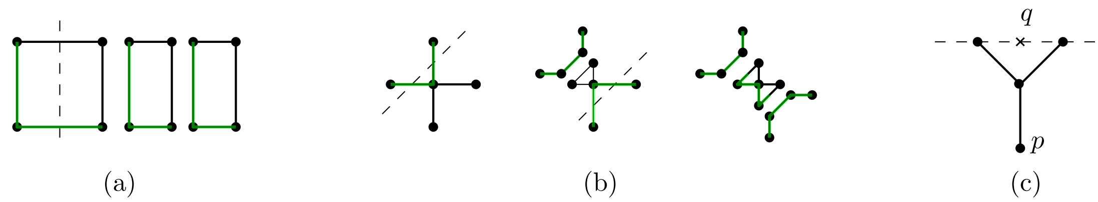

Figure 1: (a) A square with side length 1 and diameter 2 (given by green paths), and a partition with a line; (b) a 4-star partitioned into three subgraphs by inserting two lines; (c) the continuous diameter increases when inserting the dashed line into the tree (p, q is a diametral pair.)

strips. Analyzing the different types of diametral pairs of points that may have been generated in these graphs $\mathcal{G}_i$, we can prove that their diameter is smaller than $\mathrm{diam}_c(\mathcal{G})$. It is worth noting that this construction does not work using only n lines, since the width of the strips containing a vertex of $\mathcal{G}$ must tend to zero, in order to avoid diametral pairs of points located on the inserted lines whose distance is larger than the original diameter. □

2.2 Discrete Borsuk number

We now consider partitions of an abstract graph G by simply deleting edges; here all edges have the same length, equal to 1. The *discrete Borsuk number of* G, denoted by $b_d(G)$, is the minimum cardinality of a partition of G by deleting edges into subgraphs $G_1, \ldots, G_k$ (of G) such that $\max\{\mathrm{diam}_d(G_1), \ldots, \mathrm{diam}_d(G_k)\} < \mathrm{diam}_d(G)$. The following observation gives some simple examples.

Observation 2. *(i) If G is a path or a cycle of even length, $b_d(G) = 2$.*

(ii) If G is a cycle of odd length, $b_d(G) = 3$.

(iii) If G is a star graph on $k + 1$ vertices, $b_d(G) = k$.

In Section 5, we study the Borsuk number of an arbitrary tree T, in both, the discrete and the continuous setting. We show that while $b_c(\mathcal{T})$ is bounded by a constant, $b_d(T)$ can be linear with the number of vertices (as happens for the star). This linearity of the discrete Borsuk number also occurs in other families of graphs, such as unicycle graphs and maximal outerplanar graphs that are not trees.

3 Computational complexity

The problem of deciding whether the discrete Borsuk number of a graph G is below a given threshold is related to the *minimum clique cover problem.* A *clique cover* of a graph G is a partition of its vertex set into cliques. The *clique cover number* of G is the minimum size of a clique cover. The *minimum clique cover problem* seeks for a minimum clique cover.

Lemma 3. *The clique cover number of a non-complete graph G is an upper bound of $b_d(G)$, and both numbers coincide when* $\mathrm{diam}_d(G) = 2$.

Theorem 4. *Let G be a graph, and let k be a positive integer number. The problem of deciding whether $b_d(G) < k$ is NP–complete.*

Proof. Let G be a graph such that $\mathrm{diam}_d(G) > 1$ (otherwise, G is a complete graph, and $b_d(G)$ is simply the number of vertices of G, since all edges need to be removed to have each connected component with diameter zero.) The *cone* C_G of G is the graph obtained from G by adding a new vertex adjacent to all the vertices in G. The graph G has a clique cover of size k if and only if C_G has a clique cover of size k. Since $\mathrm{diam}_d(C_G) = 2$, by Lemma 3, the clique cover number of C_G is precisely $b_d(C_G)$. Thus, the result follows from the fact that deciding whether the clique cover number of an arbitrary graph is below a given threshold is an NP–complete problem [9]. □

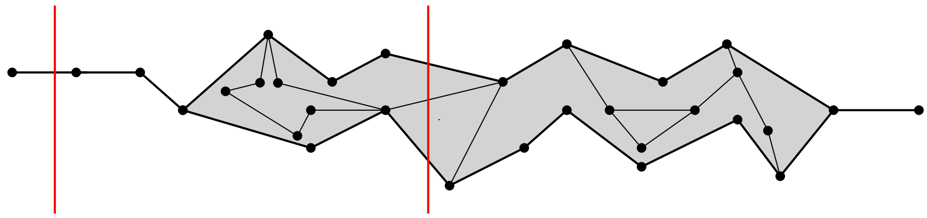

Figure 2: A graph that is monotone with respect to the x-axis; $G \cup \mathcal{F}_I$ consists of the graph (in black) and the gray region. Red vertical lines either intersect $G \cup \mathcal{F}_I$ at a single point or at a segment.

We conjecture that the problem in the continuous setting is also NP-hard, but, at the moment, a proof remains as future work.

4 Continuous Borsuk number of monotone graphs

Let G be a (plane geometric) graph, and let $G \cup \mathcal{F}_I$ be the part of the plane formed by the graph itself and all its interior faces. The graph G is said to be *ℓ-monotone* if the intersection of any line perpendicular to ℓ with $G \cup \mathcal{F}_I$ is either a single point or a segment; see Figure 2. We extend naturally this concept to the locus $\mathcal{G}$. For an ℓ-monotone graph $\mathcal{G}$, and a line ℓ' perpendicular to ℓ that is moving from left to right (parameterized by $\ell \cap \ell'$), we define the functions $d^+(\ell') = \mathrm{diam}_c((\ell'^+ \cap \mathcal{G}) \cup s')$ and $d^-(\ell') = \mathrm{diam}_c((\ell'^- \cap \mathcal{G}) \cup s')$, where s' is the maximal segment of ℓ' intersecting $\mathcal{G}$.

Lemma 5. *The functions $d^+(\ell')$ and $d^-(\ell')$ are monotone, respectively, decreasing and increasing.*

The continuous diameter can increase when partitioning a graph (see Figure 1b) but, as a straightforward consequence of the preceding lemma, we obtain that this is not true for monotone graphs.

Corollary 6. *The functions $d^+(\ell')$ and $d^-(\ell')$ associated to an ℓ-monotone graph $\mathcal{G}$ are upper-bounded by $\mathrm{diam}_c(\mathcal{G})$.*

In order to bound the continuous Borsuk number of a monotone graph, we introduce the concept of *diametral set*. The diametral set $D(p,q) \subseteq \mathcal{G}$ of a diametral pair p,q is defined as the union of all the shortest paths connecting p and q. Note that, for example, a cycle has an infinite number of diametral pairs of points, but only one distinct diametral set, which is the whole cycle (the union of the two diametral paths for each diametral pair is the same, the cycle). Thus, while a graph can have an infinite number of diametral pairs of points, we next state that this is not the case for diametral sets, which is key to prove Theorem 8 below.

Lemma 7. *Let $\mathcal{G}$ be the locus of any plane geometric graph with n vertices. The number of distinct diametral sets of $\mathcal{G}$ is in $O(n^2)$.*

Theorem 8. *Let $\mathcal{G}$ be an ℓ-monotone graph such that there are no $k+1$ disjoint diametral sets. Then, $b_c(\mathcal{G}) \leq k+2$.*

Proof. (Brief sketch.) By Corollary 6, in order to reduce the original diameter when cutting by lines, it suffices to intersect the $O(n^2)$ diametral sets of Lemma 7, with lines perpendicular to ℓ, since the new points on the cutting lines cannot cause an increase of $\mathrm{diam}_c(\mathcal{G})$. If we shorten one of the shortest paths connecting two diametral points, their distance decreases, and so each of the $O(n^2)$ diametral set only needs to be intersected once. We can prove that all the sets can be intersected using $k+1$ lines. The idea is to project each diametral set onto the line ℓ so that each set determines an interval on the line. Then, for each interval, we define a line that intersects it, and also crosses all the intervals overlapping with that one. This produces a sequence of subsets of diametral sets $\mathcal{D}_0 \supset \mathcal{D}_1 \supset \mathcal{D}_2 \ldots$, where $\mathcal{D}_0$ is the set of all diametral sets of $\mathcal{G}$, satisfying that the diametral sets in $\mathcal{D}_{i-2} \backslash \mathcal{D}_{i-1}$ do not intersect those in $\mathcal{D}_{i-1} \backslash \mathcal{D}_i$. Hence, we can find at most $k+1$ of the lines defined above, otherwise we would have $k+1$ disjoint diametral sets. □

We note that the previous bound can be attained, at least for $k = 1$. Consider, for example, the wheel graph on 33 vertices, $\mathcal{W}_{33}$, embedded in the plane such that its outer boundary is a regular 32-sided polygon, and the distance from the wheel center to each polygon vertex is one. This implies that each side has length $s = 2\sin(\pi/32) \approx 0.19$. Any two vertices of the polygon are connected by a path of length two, through the wheel center. This path is shorter than going along the boundary as soon as the other vertex is more than ten vertices away along the boundary (since $11s > 2$). It follows that the diametral pairs of this graph are given by pairs of midpoints of polygon sides that are at distance $2 + s \approx 2.19$. In fact, each midpoint has nine points at exactly that distance, corresponding to the midpoint exactly opposite, plus those of the first four sides neighboring the opposite side, in each direction. See side figure for the nine diametral pairs involving p.

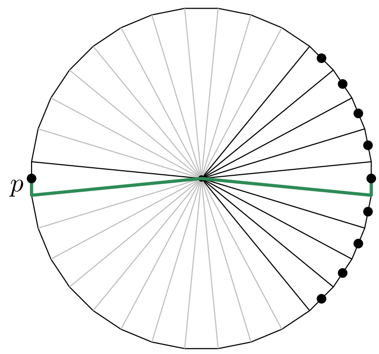

Next we argue that subdividing by one line is not enough to decrease the diameter of $\mathcal{W}_{33}$. Any line intersecting the wheel will leave at least 15 complete triangles of the wheel on one side. These triangles are contiguous, and form a fan. The diameter of any such a fan with 13 or more triangles remains the same as the original one, $2 + s$. It follows that two lines are necessary. Moreover, they are also sufficient, since two parallel lines at a very small distance that enclose the center will result in three subgraphs with smaller diameter. Therefore, $b_c(\mathcal{W}_{33}) = 3 = k + 2$.

5 Borsuk number of trees

In this section, we first compute $b_d(T)$ for an arbitrary tree T, and then we move to the continuous version of the problem, which behaves differently.

Proposition 9. *The discrete Borsuk number of any tree T with n vertices can be computed in $O(n)$ time. Furthermore,*

(i) If the center of T is not a unique vertex, then $b_d(T) = 2$.

(ii) If the center of T is a vertex v, then $b_d(T) = b_d(T') = \delta_{T'}(v)$, where T' is the subtree of T induced by the vertices of all diametral paths, and $\delta_{T'}(v)$ is the degree of v in T'.

While $b_d(T)$ depends on the center of T, we next show that the continuous Borsuk number is upper-bounded by a constant. We apply the following lemma that states that lines intersecting a tree at its center cannot cause an increase of the diameter of the tree.

Lemma 10. *Let $\mathcal{T}$ be the locus of a tree with center point $\mathcal{C}$, and let ℓ be a line that passes through $\mathcal{C}$. Then,* $\max\{\mathrm{diam}_c((\ell^+ \cap \mathcal{T}) \cup s), \mathrm{diam}_c((\ell^- \cap \mathcal{T}) \cup s)\} \leq \mathrm{diam}_c(\mathcal{T})$*, where s is the longest segment in ℓ that has its endpoints in $\mathcal{T} \cup \ell$.*

Lemma 10 also holds for lines that intersect the tree, not at the center, but infinitely close to it. Further, with lines that go exactly through the center $\mathcal{C}$, we cannot guarantee that the diameters obtained after cutting are strictly smaller than $\mathrm{diam}_c(\mathcal{T})$ (for example, the star graph with three edges of the same length and not contained in the same half-plane through the center). However, Proposition 11 below states that when a tree has Borsuk number 2, we can always find a line intersecting $\mathcal{T}$ at a point infinitely close to the center giving a correct partition (that is, the diameter decreases with respect to the original). This is an important step in order to design an algorithm for deciding whether the continuous Borsuk number of a tree is 2 or 3 which, by Theorem 12, are its possible values.

Proposition 11. *Let $\mathcal{T}$ be the locus of a tree with center point $\mathcal{C}$. If $b_c(\mathcal{T}) = 2$ then there exists a sequence of lines $\{\ell_i\}_{i\geq 0}$ satisfying that:*

(i) $\{d_T(t_i, \mathcal{C})\}_{i\geq 0}$ *approaches zero, where t_i is the closest point in $\mathcal{T} \cap \ell_i$ to $\mathcal{C}$.*

(ii) there exists $j \geq 0$ such that for every $i \geq j$, $\max\{\mathrm{diam}_c((\ell_i^+ \cap \mathcal{T}) \cup s_i), \mathrm{diam}_c((\ell_i^- \cap \mathcal{T}) \cup s_i)\} < \mathrm{diam}_c(\mathcal{T})$*, where s_i is the longest segment in ℓ_i that has its endpoints in $T \cup \ell_i$.*

Theorem 12. *Let $\mathcal{T}$ be the locus of a tree. Then, $b_c(\mathcal{T}) \leq 3$.*

Proof. (Brief sketch.) Consider a line ℓ that passes through the center point $\mathcal{C}$ of $\mathcal{T}$, and splits the tree into two graphs $\mathcal{T}_1$ and $\mathcal{T}_2$. We can assume that ℓ does not contain any edge incident or containing $\mathcal{C}$ as an interior point. By Lemma 10, the diameters of $\mathcal{T}_1$ and $\mathcal{T}_2$ are at most $\mathrm{diam}_c(\mathcal{T})$. A case analysis of how distances change after inserting the line, lets us conclude that $d_{T_1}(p, \mathcal{C}) \leq \mathrm{diam}_c(\mathcal{T})/2$ for every point p on $\mathcal{T}_1$ (analogous for $\mathcal{T}_2$). This fact is the key tool to prove that the diameters of $\mathcal{T}_1$ and $\mathcal{T}_2$ are strictly smaller than $\mathrm{diam}_c(\mathcal{T})$ when $\mathcal{C}$ is not a vertex of T, which leads to $b_c(\mathcal{T}) = 2$.

If $\mathcal{C}$ is a vertex of T, a diametral pair of $\mathcal{T}_1$ or $\mathcal{T}_2$ may consists of two leaves at distance $\mathrm{diam}_c(\mathcal{T})$. Then, we may need two lines to decrease the diameter; for example, this is the case in the star graph with all edges of the same length and such that no half-plane through the center contains all of them. It suffices to take two parallel lines to ℓ, one slightly above and the other below. This gives $b_c(\mathcal{T}) \leq 3$. □

6 Conclusions

We have introduced the concept of Borsuk number of a graph in a discrete and a continuous setting. Let us mention that this is ongoing research. In the continuous setting, we are currently focusing on proving the NP-hardness of computing $b_c(\mathcal{G})$, and whether there is a polynomial time algorithm to decide whether $b_c(G) = 2$. In addition, we are trying to answer the question of whether $b_c(\mathcal{G})$ can be upper-bounded by a constant, and designing an algorithm for trees as mentioned above. We are also delving deeper into the discrete version, currently studying the Borsuk number of unicycle graphs to better understand the behavior of this parameter.

References

[1] K. Borsuk. Three theorems on the n-dimensional sphere. *Fund. Math*, 20:177–190, 1933.

[2] J. L. De Carufel, C. Grimm, S. Schirra, and M. Smid. Minimizing the continuous diameter when augmenting a tree with a shortcut. In *Algorithms and Data Structures. WADS 2017. LNCS 10389*, pages 301– 312, 2017.

[3] C. E. Chen and R. S. Garfinkel. The generalized diameter of a graph. *Networks*, 12(3):335–340, 1982.

[4] H. G. Eggleston. Covering a three-dimensional set with sets of smaller diameter. *J. London Math.*, 30:11—24, 1955.

[5] D. Garijo, A. Márquez, N. Rodríguez, and R. I. Silveira. Computing optimal shortcuts for networks. *Eur. J. Oper. Res.*, 279(1):26–37, 2019.

[6] H Hadwiger. Mitteilung betreffend meine Note: Überdeckung einer Menge durch Mengen kleineren Durchmessers. *Commentarii Mathematici Helvetici*, 19(1):72–73, 1946.

[7] T. Jenrich and A. E. Brouwer. A 64-dimensional counterexample to Borsuk's conjecture. *Electron. J. Comb.*, 21(4):4, 2014.

[8] J. Kahn and G. Kalai. A counterexample to Borsuk's conjecture. *B. Am. Math. Soc.*, 29(1):60–62, 1993.

[9] R. M. Karp. Reducibility among combinatorial problems. In R.E. Miller, J.W. Thatcher, and J.D. Bohlinger, editors, *Complexity of Computer Computations*, Plenum Press, New York, pages 85–103. 1972.

[10] J. Perkal. Sur la subdivision des ensembles en parties de diamètre inférieur. *Colloq. Math.*, 1:45, 1947.

[11] A. S. Rissling. Das Borsuksche Problem in dreidimensionalen Räumen konstanter Krümmung. *Ukr. Geom. Sb.*, 11:78–83, 1971.

[12] C. Zong. Borsuk's partition conjecture. *Jpn. J. Math.*, 16:185–201, 2021.