

The Complexity of Sharing a Pizza

Patrick Schneider   

Department of Mathematical Sciences, University of Copenhagen, Denmark

Abstract

Assume you have a 2-dimensional pizza with $2n$ ingredients that you want to share with your friend. For this you are allowed to cut the pizza using several straight cuts, and then give every second piece to your friend. You want to do this fairly, that is, your friend and you should each get exactly half of each ingredient. How many cuts do you need?

It was recently shown using topological methods that n cuts always suffice. In this work, we study the computational complexity of finding such n cuts. Our main result is that this problem is PPA-complete when the ingredients are represented as point sets. For this, we give a new proof that for point sets n cuts suffice, which does not use any topological methods.

We further prove several hardness results as well as a higher-dimensional variant for the case where the ingredients are well-separated.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases pizza sharing, Ham-Sandwich theorem, PPA, computational geometry, computational complexity

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2021.13

Related Version *Full Version*: <https://arxiv.org/abs/2109.06752>

Funding The author has received funding from the European Research Council under the European Unions Seventh Framework Programme ERC Grant agreement ERC StG 716424 – CASe.

1 Introduction

1.1 Mass partitions

The study of *mass partitions* is a large and rapidly growing area of research in discrete and computational geometry. It has its origins in the classic *Ham-Sandwich theorem* [44]. This theorem states that any d mass distributions in \mathbb{R}^d can be simultaneously bisected by a hyperplane. A mass distribution μ in \mathbb{R}^d is a measure on \mathbb{R}^d such that all open subsets of \mathbb{R}^d are measurable, $0 < \mu(\mathbb{R}^d) < \infty$ and $\mu(S) = 0$ for every lower-dimensional subset S of \mathbb{R}^d . A vivid example of this result is, that it is possible to share a 3-dimensional sandwich, consisting of bread, ham and cheese, with a friend by cutting it with one straight cut such that both will get exactly half of each ingredient. This works no matter how the ingredients lie. In fact, as Edelsbrunner puts it, this even works if the cheese is still in the fridge [19].

But what if there are more ingredients, for example on a pizza? One way to bisect more than d masses is to use more complicated cuts, such as algebraic surfaces of fixed degree [44] or piece-wise linear cuts with a fixed number of turns [28, 40]. Another option is to use several straight cuts, as introduced by Bereg et al. [8]: Consider some arrangement of n hyperplanes in \mathbb{R}^d . The cells of this arrangement allow a natural 2-coloring, where two cells get a different color whenever they share a $(d - 1)$ -dimensional face. We say that an arrangement bisects a mass distribution μ if the cells of each color contain exactly half of μ . See Figure 1 for an illustration. It was conjectured by Langerman that any nd mass distributions in \mathbb{R}^d can be simultaneously bisected by an arrangement of n hyperplanes ([30], see also [5]). In a series of papers, this conjecture has been resolved for 4 masses in the



© Patrick Schneider;

licensed under Creative Commons License CC-BY 4.0

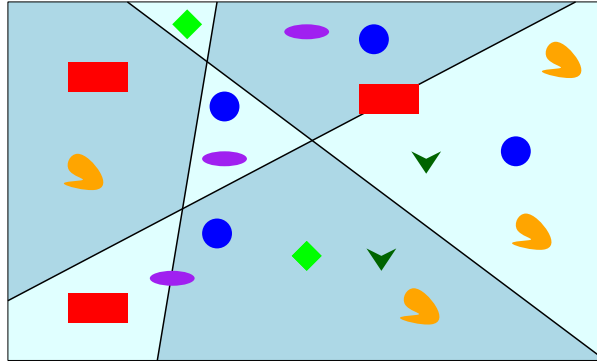
32nd International Symposium on Algorithms and Computation (ISAAC 2021).

Editors: Hee-Kap Ahn and Kunihiko Sadakane; Article No. 13; pp. 13:1–13:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** A bisection of 6 masses with 3 cuts.

plane [5], for any number of masses in any dimension that is a power of 2 (and thus in particular also in the plane) [26] and in a relaxed setting for any number of masses in any dimension [41]. However, the general conjecture remains open.

There are many other variants of mass partitions that have been studied, see e.g. [27, 38] for recent surveys. In this work, we focus on the algorithmic aspects of the 2-dimensional variant of bisections with hyperplane arrangements, that is, bisections with line arrangements. For this, let us formally define the involved objects.

► **Definition 1** (Partition induced by an arrangement of oriented lines). *Let $A = (\ell_1, \dots, \ell_n)$ be a set of oriented lines in the plane. For each ℓ_i , define R_i^+ and R_i^- the part of the plane on the positive and negative side of ℓ_i , respectively. Define $R^+(A)$ as the part of the plane lying in an even number of R_i^+ and not on any of the ℓ_i . Similarly, define $R^-(A)$ as the part of the plane lying in an odd number of R_i^+ and not on any of the ℓ_i . Now, $R^+(A)$ and $R^-(A)$ are disjoint, and they partition $\mathbb{R}^2 \setminus \{\ell_1, \dots, \ell_n\}$ into two parts.*

Note that reorienting one line just swaps $R^+(A)$ and $R^-(A)$, so up to symmetry, the two sides are already determined by the underlying unoriented line arrangement. We will thus often forget about the orientations and just say that a mass is bisected by a line arrangement. Hubbard and Karasev [26] have shown the following:

► **Theorem 2** (Planar pizza cutting theorem [26]). *Any $2n$ mass distributions in the plane can be simultaneously bisected by an arrangement of n lines.*

From an algorithmic point of view, we want to restrict our attention to efficiently computable mass distributions.

► **Definition 3** (Computable mass distribution). *A computable mass distribution is a continuous function μ which assigns to each arrangement of n oriented lines two values $\mu(R^+(A))$ and $\mu(R^-(A))$, such that $\mu(R^+(A)) + \mu(R^-(A)) = \mu(R^+(A')) + \mu(R^-(A'))$ for any two arrangements A and A' . We further assume that μ can be computed in time polynomial in the description of the input arrangement.*

We now have that the following problem always has a solution.

► **Definition 4** (PIZZACUTTING). *The problem PIZZACUTTING takes as input $2n$ computable mass distributions μ_1, \dots, μ_{2n} and returns an arrangement A of oriented lines in the plane such that for each i we have $\mu_i(R^+(A)) = \mu_i(R^-(A))$.*

An important special case of masses are point sets with the counting measure. They do not quite fit the above framework of mass distributions, as the number of points on a line can be non-zero. This can however be resolved by rounding: we say that a line arrangement bisects a point set P , if there are at most $\lfloor \frac{|P|}{2} \rfloor$ many points of P in both parts. Note that if the number of points in P is odd, this implies that at least one point needs to lie on some line. In the following, we will assume that all point sets are in general position, that is, no three points lie on a common line. With this, we can assume that for each point set at most one point lies on a line. Standard arguments (see e.g. [33]) show that the existence of partitions for mass distributions imply the analogous result for point sets with this definition of bisection. Alternatively, we give a direct proof of the following in Section 4.

► **Corollary 5** (Discrete planar pizza cutting theorem). *Any $2n$ point sets in the plane can be simultaneously bisected by an arrangement of n lines.*

Thus, also the following discrete version always has a solution.

► **Definition 6** (DISCRETEPIZZACUTTING). *The problem DISCRETEPIZZACUTTING takes as input $2n$ point sets P_1, \dots, P_{2n} in general position in the plane and returns an arrangement A of oriented lines in the plane such that for each i we have $|P_i \cap R^+(A)| = |P_i \cap R^-(A)| = \lfloor \frac{|P_i|}{2} \rfloor$.*

The pizza cutting problem can be viewed as a higher-dimensional generalization of the consensus halving and necklace splitting problems.

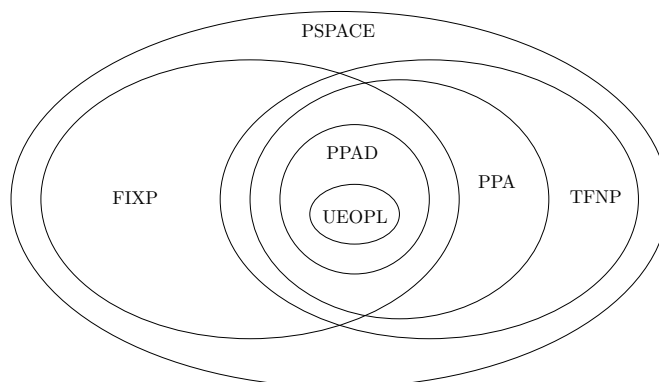
► **Definition 7** (CONSENSUSHALVING/NECKLACESPLITTING). *The problem CONSENSUSHALVING takes as input n valuation functions v_1, \dots, v_n on the interval $[0, 1]$ and returns a partition of $[0, 1]$ into $n + 1$ intervals (that is, using n cuts), each labeled “+” or “−”, such that for each valuation function we have $v_i(\mathcal{I}^+) = v_i(\mathcal{I}^-)$ (where \mathcal{I}^x denotes the union of intervals labeled “ x ”). The problem NECKLACESPLITTING is the same, but taking as input n point sets, again using the above definition of bisections of point sets.*

Again, a solution to the problems is always guaranteed to exist. In the case of mass distributions, this result is known as the *Hobby-Rice theorem* [25]. For necklaces, the statement holds even for the generalized problem of sharing with more than two people [2, 3]. In this work, whenever we refer to the *Necklace splitting theorem*, we mean the version for two people.

Finally, for all above problems, we can also consider the decision version, where we are given one more measure or point set than the number that can always be bisected, and we need to decide whether there still is a bisection. We denote these problems by adding “DECISION” to their name.

1.2 Algorithms and complexity

Most proofs of existence of certain mass partitions use topological methods, which, by their nature, are not algorithmic. Thus, there has been quite some effort in developing algorithms that find these promised partitions, ideally efficiently. Arguably the most famous result in this direction are the algorithms for Ham-Sandwich cuts by Lo, Steiger and Matoušek [32, 31]. While in the plane, their algorithm runs in linear time, in general the runtime shows an exponential dependency on the dimension. This curse of dimensionality seems to be a common issue for many algorithmic version of mass partition problems, and most problems have only been studied from an algorithmic point of view in low dimensions, where the constructed algorithms either rely on a relatively small space of solutions or a simplified proof which allows for an algorithmic formulation, see e.g. [1, 6, 37].



■ **Figure 2** Containment relations of some complexity classes for total problems.

The curse of dimensionality was made explicit for the first time by Knauer, Tiwary and Werner, who showed that deciding whether there is a Ham-Sandwich cut through a given point in arbitrary dimensions is $W[1]$ -hard (and thus also NP-hard) [29]. More recently, in several breakthrough papers, Filos-Ratsikas and Goldberg have shown that computing Ham-Sandwich cuts in arbitrary dimensions is PPA-complete, and so are NECKLACESPLITTING and CONSENSUSHALVING, the latter even in an approximation version [23, 24].

The class PPA was introduced in 1994 by Papadimitriou [36]. It captures search problems, where the existence of a solution is guaranteed by a parity argument in a graph. More specifically, the defining problem is the following search problem in a (potentially exponentially sized) graph G : given a vertex of odd degree in G , where G is represented via a polynomially-sized circuit which takes as input a vertex and outputs its neighbors, find another vertex of odd degree. The class PPA is a subclass of TFNP (Total Function NP), which are total search problems where solutions can be verified efficiently.

A subclass of PPA that is of importance in this work is UEOPL [22], which is a subclass of PPAD [36]. PPAD is similar to PPA, but instead of an undirected graph, we are given a directed graph in which each vertex has at most one predecessor and at most one successor. We are given a vertex without predecessor, and our goal is to find another vertex without predecessor or successor. If we are further given a potential function which strictly increases on a directed path such that there is a unique vertex with maximal potential, finding this vertex is the defining problem for the class UEOPL.

The class UEOPL is related to mass partitions through the fact that finding the unique discrete Ham-Sandwich cut in the case that the point sets are *well-separated* is in UEOPL [14]. We say that k point sets P_1, \dots, P_k in \mathbb{R}^d are well-separated, if for no $d-2$ -dimensional affine subspace their convex hulls can be intersected with a $(d-2)$ -dimensional affine subspace.¹ This definition extends to masses by forbidding intersections of affine subspaces with the convex hulls of their supports. In fact, for well-separated masses and point sets, the α -Ham-Sandwich theorem states that it is always possible to simultaneously cut off an arbitrary given fraction from each mass or point set with a single hyperplane [4, 43].

The class PPAD has so far mostly been related to the computation of Nash and market equilibria [10, 11, 12, 13, 15, 16, 20, 34, 39, 42, 45].

¹ Admittedly, this would be a very weird pizza.

Finally, the two last classes that are relevant for this work are FIXP and $\exists\mathbb{R}$. The class FIXP is the class of problems that can be reduced to finding a Brouwer fixed point [21], whereas $\exists\mathbb{R}$ is the class of decision problems which can be written in the existential theory of the reals.

In the context of mass partitions, apart from the above mentioned results on Ham-Sandwich cuts, Consensus halvings and Necklace splittings, some of the above classes also appear in the complexity of *Square-Cut Pizza sharing*. In this variant, introduced by Karasev, Roldán-Pensado and Soberón, n masses in the plane are bisected by a cut which is a union of at most n axis-parallel segments, or in other words, a piecewise linear cut with at most $n - 1$ 90° -turns [28]. It was recently shown by Deligkas, Fearnley and Melissourgos that finding such a cut even for restricted inputs is PPA-complete, finding a cut where a constant number of additional turns are allowed is PPAD-hard, and that the corresponding decision version is NP-hard [17]. For more general masses, they show the problem to be FIXP-hard and the decision version to be $\exists\mathbb{R}$ -complete. This work was heavily influenced and motivated by their paper: apart from the different setting, their results are very similar from the results in this work, showing the relation between those two pizza cutting variants. Indeed, in a new version of their paper, the authors of [17] prove some of the same hardness results as this work, and also some stronger hardness results for approximate bisections.

1.3 Our contributions

Our main contribution is that DISCRETEPIZZACUTTING is PPA-complete. While the hardness is rather straight-forward, the containment requires some more work. We give a new proof for the existence of pizza cuttings for point sets which, while inspired by the ideas of Hubard and Karasev, uses only elementary geometric techniques which allow us to place the problem in PPA.

We further prove that PIZZACUTTING is FIXP-hard and that PIZZACUTTINGDECISION is $\exists\mathbb{R}$ -hard and that finding the minimum number of cuts required to bisect an instance of DISCRETEPIZZACUTTING is NP-hard.

Finally, for well-separated masses, we show that the α -Ham-Sandwich theorem generalizes to pizza cuttings. For point sets in fixed dimensions, we give a linear time algorithm to find such a cut, whereas for arbitrary dimensions we place the problem in UEOPL.

Many of our results are more or less direct applications of known results. In particular, all hardness results follow from a proof of the existence of a Consensus Halving from the existence of a pizza cut. We present this proof and the hardness results in Section 2. In Section 3, we consider the case of well-separated masses and point sets, as some of the ideas are needed for our containment proof. The part where the most new ideas are used is Section 4, where we show the containment of DISCRETEPIZZACUTTING in PPA.

2 Hardness results

In this section, we give a proof of existence of Consensus halvings and necklace splittings using the planar pizza cutting theorem. This proof gives a natural reduction to the corresponding algorithmic problems, and thus a variety of hardness results follow.

► **Lemma 8.** *The planar pizza cutting theorem implies the Hobby-Rice theorem.*

Proof. Consider the moment curve γ in \mathbb{R}^2 , that is, the curve parametrized by (t, t^2) . Note that any line ℓ intersects γ in at most two points, call them g_1 and g_2 , with g_1 being to the left of g_2 (in the case of a single or no intersection, we may consider $g_1 = -\infty$ or $g_2 = \infty$

13:6 The Complexity of Sharing a Pizza

or both). Let x_1 and x_2 be the projections of g_1 and g_2 to the x -axis under the projection $\pi(x, y) := x$. Consider now a half-plane h bounded by ℓ and consider its intersection with γ . If h lies below ℓ , then $h \cap \gamma$ projects to the interval $[x_1, x_2]$. If h lies above ℓ , then $h \cap \gamma$ projects to $(-\infty, x_1] \cup [x_2, \infty)$. Similarly, if ℓ is vertical, $h \cap \gamma$ projects either to $(-\infty, x_1]$ or $[x_2, \infty)$. Thus, in all cases $h \cap \gamma$ projects to an interval or the complement of an interval, which, in a slight abuse of notation, we denote by $\pi(h)$.

Given a valuation function v , we now want to define a mass distribution μ in the plane. For this, it is enough to just define μ on all half-planes. This we can do using the above observations: for any half-plane h , we define $\mu(h) := v(\pi(h))$.

This way, we have defined n mass distributions. Now, we do the same thing, but shift the interval $[0, 1]$, which is the support of the valuation functions, to the interval $[2, 3]$. More formally, we consider the projection $\varphi(x, y) := x - 2$ and, given a valuation function v , define a mass distribution η as $\eta(h) := v(\varphi(h))$.

We have now defined $2n$ mass distributions. By the pizza cutting theorem, there exists an arrangement $A = (\ell_1, \dots, \ell_n)$ of n lines which simultaneously bisect these mass distributions. Consider the intervals I_1 and I_2 on γ defined by $t \in [0, 1]$ and $t \in [2, 3]$. As there are at most $2n$ intersections of A and γ , by the pigeonhole principle there are at most n intersections in one of them, say I_1 . Let i_1, \dots, i_n be the projections of these intersections under the projection π . We claim that i_1, \dots, i_n simultaneously bisect v_1, \dots, v_n .

To show this, consider some valuation function v_i . By construction, we now have that $v_i(\mathcal{I}^+) = \mu_i(R^+(A)) = \mu_i(R^-(A)) = v_i(\mathcal{I}^+)$, which proves the claim. In the case where I_2 has at most n intersections, we can do the same argument, replacing π with φ and μ with η . ◀

Note that all the steps in the proof can be computed in polynomial time. Thus, as CONSENSUSHALVING is FIXP-hard [18], we immediately get the following:

► **Corollary 9.** *PIZZACUTTING is FIXP-hard.*

In the discrete setting, the above proof can be phrased even simpler: for each point x in $[0, 1]$, we just define two points (x, x^2) and $(x + 2, (x + 2)^2)$. As NECKLACESPLITTING is PPA-hard [24], we get that

► **Corollary 10.** *DISCRETEPIZZACUTTING is PPA-hard.*

Clearly, the construction in the proof above also works for more than n valuation functions. In [18], it was shown that deciding whether $n + 1$ valuation functions can be bisected with n cuts is $\exists\mathbb{R}$ -hard. It thus follows that it is $\exists\mathbb{R}$ -hard to decide whether $2n + 2$ masses can be bisected by n lines. However, in the case where n is even, we can also only use π and map all valuation functions to a single interval on γ , which analogously proves the Hobby-Rice theorem for even n . From an asymptotic point of view, this restriction to even values does not matter, so using this reduction, we get the following slightly stronger statement.

► **Corollary 11.** *PIZZACUTTINGDECISION is $\exists\mathbb{R}$ -hard.*

Finally, it was shown in [9, 35] that finding the minimal number of cuts required to split a necklace is NP-hard. We again get the analogous result for discrete pizza cuttings.

► **Corollary 12.** *Finding the minimal number of lines that simultaneously bisect a family of $2n$ point sets is NP-hard.*

3 Well-separated point sets

In this section we consider well-separated mass distributions and point sets. We generalize the α -Ham-Sandwich theorem to pizza cuttings.

► **Theorem 13.** *Let μ_1, \dots, μ_{nd} be nd well-separated mass distributions in \mathbb{R}^d . Given a vector $\alpha = (\alpha_1, \dots, \alpha_{nd})$, with each $\alpha_i \in [0, 1]$, there exists an arrangement A of n oriented hyperplanes such that for each $i \in \{1, \dots, nd\}$ we have $\mu_i(R^+(A)) = \alpha_i \mu_i(\mathbb{R}^d)$.*

Proof. By the α -Ham-Sandwich theorem [4], for any d well-separated mass distributions μ_1, \dots, μ_d and any vector $(\alpha_1, \dots, \alpha_d)$, there exists a unique single hyperplane cutting each mass distribution in the required ratio. By the definition of well-separatedness, this hyperplane does not intersect the support of any other mass distribution. Partition the point set into n parts of d point sets each. For each part, pick some oriented hyperplane which intersects the support of all masses in this part. This defines an arrangement B of oriented hyperplanes. For each mass μ_i , consider the intersection of its support with the positive side of the hyperplane intersecting it, as well as with $R^+(B)$. If these two coincide, set $\alpha'_i := \alpha_i$, otherwise set $\alpha'_i := 1 - \alpha_i$. Now, taking the α -Ham-Sandwich cut for the vectors α' gives the required arrangement. ◀

We call such an arrangement an α -Pizza cut. From the discrete α -Ham-Sandwich theorem [43], we analogously get the discrete version of the above.

► **Corollary 14.** *Let P_1, \dots, P_{nd} be nd well-separated point sets in \mathbb{R}^d . Given a vector $\alpha = (k_1, \dots, k_{nd})$, where each k_i is an integer with $0 \leq k_i \leq |P_i|$, there exists an arrangement A of n oriented hyperplanes such that for each $i \in \{1, \dots, nd\}$ we have $|P_i \cap R^+(A)| = k_i$.*

In [14], it was shown that the problem of computing an α -Ham-Sandwich cut for point sets is in UEOPL. As our α -Pizza cuts are just a union of α -Ham-Sandwich cuts, their result generalizes to our setting.

► **Corollary 15.** *The problem of computing an α -Pizza cut for point sets is in UEOPL.*

► **Remark 16.** The computation of the vector α' does not directly translate into the setting of UEOPL. However, in [14], they start with an arbitrary hyperplane, and then rotate it in a well-defined fashion to the solution. This immediately translates to our setting: we just start with an arbitrary arrangement, and the choice of the direction of rotation works analogously to the choice in [14].

Further, Bereg [7] has shown that an α -Ham-Sandwich cut for d point sets of m points total in \mathbb{R}^d can be computed in time $m2^{O(d)}$. In particular, if d is fixed, this algorithm runs in linear time. Again, we get the same result for α -Pizza cuts.

► **Corollary 17.** *Let P_1, \dots, P_{nd} be well-separated points sets in \mathbb{R}^d with $\sum_{i=1}^{nd} |P_i| = m$. Then an α -Pizza cut for P_1, \dots, P_{nd} can be computed in time $m2^{O(d)}$.*

Proof. Partition the point sets into parts $P_{(i-1)d+1}, \dots, P_{id}$, for $i \in \{1, \dots, n\}$. Compute the vector α' in time $O(nd)$. As $m \geq nd$, the runtime of the second part dominates the total runtime. For each $i \in \{1, \dots, n\}$, use Bereg's algorithm to compute the α -Ham-Sandwich cut for $P_{(i-1)d+1}, \dots, P_{id}$. It follows that the solution is an α -Pizza cut. The runtime of the algorithm is

$$\sum_{i=1}^n (|P_{(i-1)d+1}| + \dots + |P_{id}|) 2^{O(d)} = \sum_{i=1}^{nd} |P_i| 2^{O(d)} = m 2^{O(d)}. \quad \blacktriangleleft$$

4 Containment results

In this section, we prove that DISCRETEPIZZACUTTING is in PPA. We do this by giving a new proof of the discrete planar pizza cutting theorem, which allows for an algorithm in PPA. Before we go into the details of the proof, we briefly sketch the main ideas. The main structure of our proof is similar to and inspired by the original proof of Hubard and Karasev [26], but we replace their topological arguments with easier ones that only use the combinatorics of point sets, but hence do not work for the more general case of mass distributions.

The main idea is that we continuously transform well-separated point sets into the point sets which we want to bisect. In the beginning of this process, we have several bisections, namely one for each partition of the labels of the point sets into pairs, and this number is odd. During the process, we pull these bisecting arrangements along. Every time the orientation of some triple of points changes, it can happen that one of the arrangements is not bisecting anymore. The main step in the proof is, that in these cases, we can always slightly change this arrangement so that it is again bisecting, or that there is a second arrangement that also is not bisecting anymore. In other words, some bisections might vanish during the process, but if they do, then they always vanish in pairs. This step is also where our proof differs from the one by Hubard and Karasev.

Once we have this, the remainder of the proof is rather simple: as we started with an odd number of bisections, and they always vanish in pairs, the number of bisections is always odd, and thus in particular at least 1. Further, we can build a graph where each vertex corresponds to a point set in our process with an arrangement, where the vertices are connected whenever one arrangement is the pulled along version of the other one, or when both point sets are the same and the arrangements are the two arrangements that vanish at this point of the process. (In the end, some of these connections will be paths instead of single edges.) Adding an additional vertex which we connect to all starting arrangements, we get a graph in which the only odd-degree vertices are this additional vertex and the final solutions.

4.1 A proof of the discrete planar pizza cutting theorem

We now proceed to give a detailed proof of the discrete planar pizza cutting theorem (Corollary 5). Let P_1, \dots, P_{2n} be the point sets we wish to bisect. Let further $m := \sum_{i=1}^{2n} |P_i|$.

► **Lemma 18.** *We may assume that each P_i contains an odd number of points.*

Proof. For each point set P_i with an even number of points, add some arbitrary new point q_i to P_i , such that the point sets are still in general position. Take some bisecting arrangement A of the resulting point sets. As all of these point sets consist of an odd number of points in general position, each of them must contain a point p_i which lies on a line of A . As each line in A can pass through at most two points by the general position assumption, there is exactly one such point in each point set. Now, remove q_i again. If $p_i = q_i$, the arrangement still bisects P_i . Otherwise, one side, without loss of generality R^+ , contains one point too few. Rotate the line through p_i slightly so that p_i lies in R^+ . Now the arrangement again bisects P_i . ◀

So, from now on we may assume that each P_i contains an odd number of points. Let Q_1, \dots, Q_{2n} be point sets of the same size, that is, $|Q_i| = |P_i|$, that are well-separated. Match each point $q \in Q_i$ with a point $p \in P_i$ and consider for each such pair the map $\varphi(t) := tp + (1-t)q$. These maps define a map $\Phi(t)$, which assigns to each t the point set defined by the $\varphi(t)$'s.

During this process, some orientations of triples of points change. We may assume, that no two triples change their orientation at the same time. Further, as each point crosses each line spanned by two other points at most once, the total number of orientation changes is in $O(m^3)$. Thus, the interval $[0, 1]$ is partitioned into $O(m^3)$ subintervals, in each of which the orientations of all triples stay the same, that is, the *order type* is invariant. At the boundaries of these subintervals, some three points are collinear.

Taking a representative of each subinterval and the point sets at their boundaries, we thus get a sequence of point sets

$$Q_i = P_i^{(0)}, P_i^{(0.5)}, P_i^{(1)}, \dots, P_i^{(k)}, P_i^{(k.5)}, P_i^{(k+1)}, \dots, P_i^{(C)} = P_i,$$

where $C \in O(m^3)$ is the number of orientation changes, $P^{(k)} := \bigcup_{i=1}^{2n} P_i^{(k)}$ is a point set in general position, and $P^{(k.5)} := \bigcup_{i=1}^{2n} P_i^{(k.5)}$ is a point set with exactly one collinear triple. We will mainly work with the sets $P^{(k)}$, the sets $P^{(k.5)}$ are just used to make some arguments easier to understand.

As argued before, for each $P^{(k)}$, any bisecting arrangement A contains exactly one point $p_i^{(k)}(A)$ of each $P_i^{(k)}$ on one of its lines. In particular, A is defined by a set of pairs of points $(p_i^{(k)}(A), p_j^{(k)}(A))$, where each pair defines one line of the arrangement.

► **Lemma 19.** *For $P^{(0)}$, the number of bisecting arrangements is odd.*

Proof. It is known that for 2 separated point sets, each of odd size, there is a unique Ham-Sandwich cut (see e.g. [43]). We have seen in the proof of Theorem 13, that for well-separated point sets, each bisecting arrangement corresponds to n Ham-Sandwich cuts, each for a pair of point sets. Thus, the number of bisecting arrangements is the same as the number of partitions of $2n$ elements into pairs. This number is $(2n - 1)!! = (2n - 1)(2n - 3) \cdots 3 \cdot 1$, which is a product of odd numbers, and thus odd. ◀

Let us now follow some arrangement A through the process. More precisely, Let $A^{(0)}$ be a bisecting arrangement for $P^{(0)}$, defined by pairs of points $(p_i^{(0)}(A), p_j^{(0)}(A))$. We now consider the sequence of arrangements $A^{(k)}$, where each $A^{(k)}$ is defined by the corresponding pairs of points $(p_i^{(k)}(A), p_j^{(k)}(A))$. Clearly, if $A^{(k)}$ is bisecting and the orientation change from $P^{(k)}$ to $P^{(k+1)}$ is not a point moving over a line in the arrangement $A^{(k)}$, then $A^{(k+1)}$ is still bisecting. For the other changes, we need the following lemma. Here, we say that an arrangement $A^{(k.5)}$ is *almost bisecting* for $P^{(k.5)}$, if it bisects each $P_i^{(k.5)}$ except for one, for which two points are on a line of the arrangement, and for the remaining points one side contains exactly one point more.

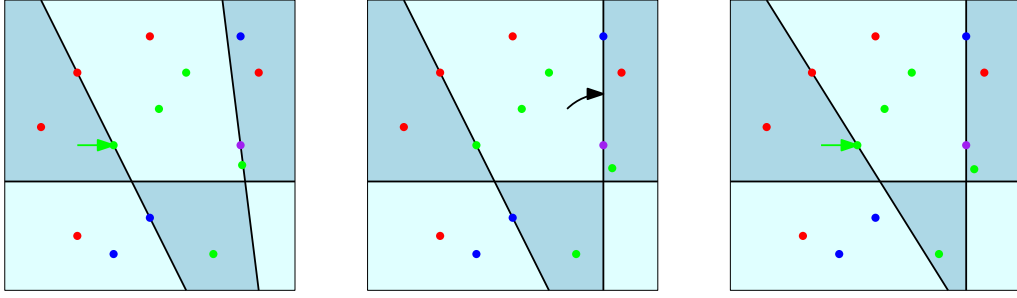
► **Lemma 20.** *Let $A^{(k)}$ and $A^{(k+1)}$ be such that $A^{(k)}$ is bisecting and $A^{(k+1)}$ is not. Then there either exists a sequence of arrangements*

$$A^{(k)}, A^{(k.5)} = A_0^{(k.5)}, A_1^{(k.5)}, \dots, A_L^{(k.5)} = B^{(k.5)}, B^{(k+1)},$$

where each $A_i^{(k.5)}$ is almost bisecting, $B^{(k+1)}$ is bisecting and $B^{(k)}$ is not bisecting, or there exists a sequence of arrangements

$$A^{(k)}, A^{(k.5)} = A_0^{(k.5)}, A_1^{(k.5)}, \dots, A_L^{(k.5)} = B^{(k.5)}, B^{(k)},$$

where each $A_i^{(k.5)}$ is almost bisecting, $B^{(k)}$ is bisecting and $B^{(k+1)}$ is not bisecting. Further, in the second case, the sequence for $B^{(k)}$ is the reverse of the sequence for $A^{(k)}$.



■ **Figure 3** Going from $A^{(k.5)}$ (left) via $B^{(k.5)}$ (middle) to $B^{(k+1)}$ (right). In this example, $B^{(k+1)}$ does not bisect the blue point set anymore.

In this lemma, the first case corresponds to a bisecting arrangement that can be changed to a new bisecting arrangement, whereas the second case corresponds to two bisecting arrangements disappearing at the same time.

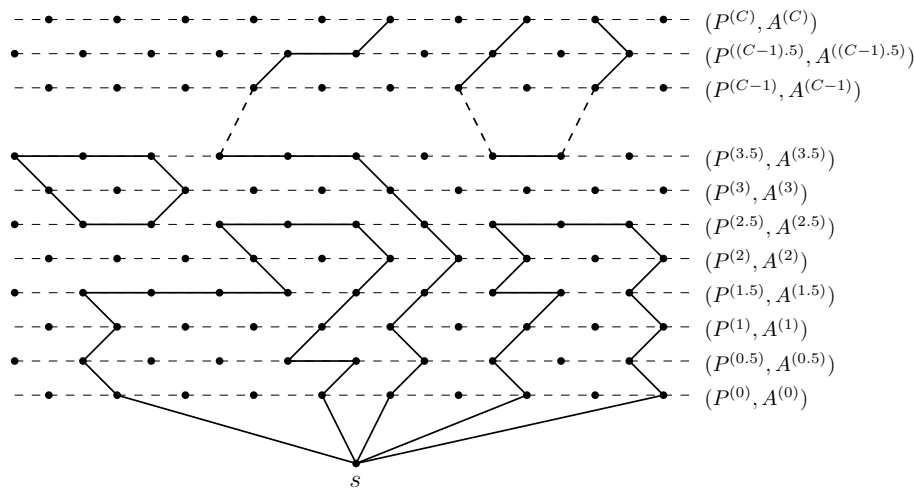
Proof. As mentioned above, the situation of the lemma can only occur if the orientation change from $P^{(k)}$ to $P^{(k+1)}$ corresponds to a point $q^{(k)}$ moving over a line ℓ of the arrangement $P^{(k)}$. Without loss of generality, let ℓ be defined by the points $p_1^{(k)} \in P_1^{(k)}$ and $p_2^{(k)} \in P_2^{(k)}$. There are two cases we consider: either $q^{(k)}$ is in the same point set as one of the points on ℓ , without loss of generality $q^{(k)} \in P_1^{(k)}$, or it is in a different point set, without loss of generality $q^{(k)} \in P_3^{(k)}$. We start with the second case. For an illustration of that case, see Figure 3.

Case 1: $q^{(k)} \in P_3^{(k)}$. Orient the arrangement in such a way that $q^{(k)}$ is in $R^+(A^{(k)})$. In $A^{(k.5)}$ the line ℓ contains three points, namely $p_1^{(k)}$, $p_2^{(k)}$ and $q^{(k)}$. Note that $A^{(k.5)}$ is almost bisecting, and $R^+(A^{(k.5)})$ is the smaller side for $P_3^{(k.5)}$. There exists another line ℓ' in $A^{(k.5)}$ which contains a second point $p_3^{(k.5)}$ of $P_3^{(k.5)}$. Rotate ℓ' around the other point on it such that $p_3^{(k.5)}$ is in $R^+(A^{(k.5)})$. Note that this direction of rotation is unique. Continue the rotation until ℓ' hits another point $q_j^{(k.5)} \in P_j^{(k.5)}$ which is not on some line of $A^{(k.5)}$. The resulting arrangement is now $A_1^{(k.5)}$.

Note that $A_1^{(k.5)}$ is again almost bisecting, and that the point set which is not exactly bisected is $P_j^{(k.5)}$. While $j \notin \{1, 2, 3\}$, we can now find another point in $P_j^{(k.5)}$ lying on a line of the arrangement $A_l^{(k.5)}$, rotate this line into the correct direction until a new point is hit, to get a new arrangement $A_{l+1}^{(k.5)}$. As all of these arrangements are different, and there are only finitely many possible arrangements, at some point we will have $j \in \{1, 2, 3\}$. This gives the arrangement $B^{(k.5)}$. There are now several options to consider: there are 3 different ways how $p_1^{(k)}$, $p_2^{(k)}$ and $q^{(k)}$ can lie on ℓ , j can be 1, 2 or 3, and the smaller side of $P_j^{(k.5)}$ can be $R^+(B^{(k.5)})$ or $R^-(B^{(k.5)})$. In all of the cases, the only change from the arrangements $B^{(k)}$ to $B^{(k+1)}$ are the defining points of ℓ . It follows that exactly one of $B^{(k)}$ and $B^{(k+1)}$ can be bisecting, depending on whether the point on ℓ is in $R^+(B^{(k+1)})$ or not.

Case 2: $q^{(k)} \in P_1^{(k)}$. In this case, we just have $A^{(k.5)} = B^{(k.5)}$.

Finally, the last claim follows from the fact that the directions of rotations from $A_l^{(k.5)}$ to $A_{l+1}^{(k.5)}$ are unique. ◀



■ **Figure 4** A schematic drawing of the graph whose leaves correspond to bisecting arrangements.

► **Remark 21.** During the process, it can also happen that two bisecting arrangement appear at the same time. In this case, we immediately get the analogous lemma, just reversing k and $(k + 1)$.

With all these lemmas at hand, we can now finally give a

Proof of the discrete planar pizza cutting theorem. Let $P = (P_1, \dots, P_{2n})$ be in general position. By Lemma 18, we may assume that each P_i contains an odd number of points. Consider the sequence $P^{(0)}, P^{(1)}, \dots, P^{(C)} = P$ of point sets. By Lemma 19, $P^{(0)}$ has an odd number of bisecting arrangements. When going from $P^{(k)}$ to $P^{(k+1)}$, either all bisecting arrangements stay bisecting or, by Lemma 20, one of them can be changed into a new bisecting arrangement or exactly two bisecting arrangements appear or disappear. It follows that for each k , $P^{(k)}$ has an odd number of bisecting arrangements. In particular, also P has an odd number of bisecting arrangements, and thus at least 1. ◀

4.2 Containment in PPA

In order to show that DISCRETEPIZZACUTTING is in PPA, we need to define a graph where the neighborhoods of each vertex can be efficiently computed and all odd-degree vertices, except the starting vertex, correspond to bisecting arrangements. In the following, we describe such a graph. For an illustration of the graph, see Figure 4.

The vertex set. Our vertex set consists of a starting vertex s , as well as vertices of the form $(P^{(k)}, A^{(k)})$ and $(P^{(k.5)}, A^{(k.5)})$. For vertices $(P^{(k)}, A^{(k)})$, $A^{(k)}$ is an arrangement which is not necessarily bisecting, but whose lines contain exactly one point of each $P_i^{(k)}$. Similarly, for vertices $(P^{(k.5)}, A^{(k.5)})$, the arrangement $A^{(k.5)}$ is not necessarily almost bisecting, but its lines contain at least one point of each $P_i^{(k.5)}$ and exactly 2 for one of them. In particular, one of the lines of $A^{(k.5)}$ is the line ℓ through the unique three collinear points.

The edge set. The starting vertex s is connected to all vertices $(P^{(0)}, A^{(0)})$, for which $A^{(0)}$ is bisecting. A vertex of the form $(P^{(k)}, A^{(k)})$, $k < L$ is connected to the vertices $(P^{((k-1).5)}, A^{((k-1).5)})$ and $(P^{(k.5)}, A^{(k.5)})$ if and only if $A^{(k)}$ is bisecting. Otherwise, it is not connected to any other vertex. Similarly, the vertices $(P^{(C)}, A^{(C)})$ are connected

13:12 The Complexity of Sharing a Pizza

to $(P^{((C-1)\cdot 5)}, A^{((C-1)\cdot 5)})$ if and only if $A^{(C)}$ is bisecting, and not connected to anything otherwise. Also a vertex of the form $(P^{(k\cdot 5)}, A^{(k\cdot 5)})$ where $A^{(k\cdot 5)}$ is not almost bisecting is not connected to any other vertex. For vertices $(P^{(k\cdot 5)}, A^{(k\cdot 5)})$ with $A^{(k\cdot 5)}$ almost bisecting we distinguish two cases. If ℓ does not contain a point of the point set which is contained twice in the lines of $A^{(k\cdot 5)}$, we connect $(P^{(k\cdot 5)}, A^{(k\cdot 5)})$ to the two vertices which correspond to the rotations defined in the proof of Lemma 20. Otherwise, we connect $(P^{(k\cdot 5)}, A^{(k\cdot 5)})$ to the single vertex that we get by such a rotation, as well as to either $(P^{(k)}, A^{(k)})$ or $(P^{(k+1)}, A^{(k+1)})$, depending on which one has the bisecting arrangement (note that we get from Lemma 20 that exactly one of the two arrangements is bisecting).

Note that all vertices have degree 0 or 2, except s and the vertices $(P^{(C)}, A^{(C)})$ with $A^{(C)}$ bisecting, that is, the vertices corresponding to solutions. All the vertices corresponding to solutions have degree 1. The starting vertex has degree $(2n - 1)!!$, which is exponential, so we cannot compute its neighborhood in polynomial time. This is however not an issue, as already noted in the paper where PPA is defined [36], provided that we can decide for any two vertices in polynomial time whether they are connected, and there is an efficiently computable *pairing function*, that is, a function $\phi(v, v')$ which takes as input a vertex v and an outgoing edge $\{v, v'\}$ and computes another vertex v'' such that $\{v, v''\}$ is also an edge. If v has even degree, then $\phi(v, v') \neq v'$ for all v' . If v has odd degree, then there is exactly one v' for which $\phi(v, v') = v'$. The exact statement from [36] is as follows:

► **Lemma 22** ([36]). *Any problem defined in terms of an edge recognition algorithm and a pairing function is in PPA.*

It remains to show that we have both these ingredients.

► **Lemma 23.** *For any two vertices, we can decide in polynomial time whether they are connected.*

Proof. It follows from the construction that for any vertex except the starting vertex s , the neighborhood can be computed in polynomial time. In particular, it can also be checked whether two such vertices are connected. As for the starting vertex s , some other vertex can only be connected to it if its point set is $P^{(0)}$ and its arrangement is bisecting. The first is encoded in the label of the vertex and the second can easily be checked in polynomial time. ◀

► **Lemma 24.** *There is a pairing function ϕ which can be computed in polynomial time.*

Proof. As all vertices except s have degree 0, 1 or 2, and the neighborhoods can be computed in polynomial time, the pairing function follows trivially for these vertices. As for s , we note that its neighbors correspond to partitions of the $2n$ point sets into pairs. In other words, its neighbors can be interpreted as perfect matchings in the complete graph K_{2n} with vertex set $\{w_1, \dots, w_{2n}\}$. It is thus enough to describe a pairing function for these perfect matchings. We do this in an algorithmic fashion.

Let M be some perfect matching. Consider the vertices w_1 and w_2 . Assume first that they are not connected to each other, that is, we have two distinct edges $\{w_1, w_a\}$ and $\{w_2, w_b\}$ in M . In that case, define the pairing $M' := M \setminus \{\{w_1, w_a\}, \{w_2, w_b\}\} \cup \{\{w_1, w_b\}, \{w_2, w_a\}\}$, that is, flip the edges incident to w_1 and w_2 . Note that the pairing of M' is again M , that is, it indeed is a pairing.

If w_1 and w_2 are connected to each other, repeat the same process with w_3 and w_4 , and so on. This process defines a pairing for each matching except $\{\{w_1, w_2\}, \dots, \{w_{2n-1}, w_{2n}\}\}$. Further, the algorithm clearly runs in polynomial time. We thus get a valid pairing function. ◀

We thus get from Lemma 22 that DISCRETEPIZZACUTTING is in PPA. Together with Corollary 10, we can thus conclude our main result.

► **Corollary 25.** *DISCRETEPIZZACUTTING is PPA-complete.*

5 Conclusion

We have shown several complexity results related to the pizza cutting problem. Our main result is that DISCRETEPIZZACUTTING is PPA-complete. While we have only considered the planar case, the problem as well as our arguments extend to higher dimensions. More precisely, the proof of the Hobby-Rice theorem using the pizza cutting theorem can be adapted to work for any dimension of the pizza cutting theorem. Further, the proof of the discrete pizza cutting theorem can be adapted to any dimension where the number of initial solutions is odd. This is the case if the dimension is a power of 2 [26]. Thus, in all these dimensions the analogous versions of DISCRETEPIZZACUTTING are also PPA-complete, assuming the dimension is fixed. On the other hand, it is an open problem whether the pizza cutting theorem also holds in other dimensions.

We have also shown that PIZZACUTTING is FIXP-hard. It is an interesting problem whether the problem is in FIXP or even harder. One way to show that the problem is in FIXP is to find a proof for the planar pizza cutting theorem using Brouwer’s fixpoint theorem. Such a proof would have the potential to generalize to any dimension, resolving the above question.

References

- 1 Oswin Aichholzer, Nieves Atienza, Ruy Fabila-Monroy, Pablo Perez-Lantero, Jose M Diaz-Báñez, David Flores-Peñaloza, Birgit Vogtenhuber, and Jorge Urrutia. Balanced islands in two colored point sets in the plane. *arXiv preprint*, 2015. [arXiv:1510.01819](#).
- 2 Noga Alon. Splitting necklaces. *Advances in Mathematics*, 63(3):247–253, 1987.
- 3 Noga Alon and Douglas B West. The Borsuk-Ulam theorem and bisection of necklaces. *Proceedings of the American Mathematical Society*, 98(4):623–628, 1986.
- 4 Imre Bárány, Alfredo Hubard, and Jesús Jerónimo. Slicing convex sets and measures by a hyperplane. *Discrete & Computational Geometry*, 39(1-3):67–75, 2008.
- 5 Luis Barba, Alexander Pilz, and Patrick Schnider. Sharing a pizza: bisecting masses with two cuts. *arXiv preprint*, 2019. [arXiv:1904.02502](#).
- 6 Sergey Bereg. Equipartitions of measures by 2-fans. In *International Symposium on Algorithms and Computation*, pages 149–158. Springer, 2004.
- 7 Sergey Bereg. Computing generalized ham-sandwich cuts. *Information Processing Letters*, 112(13):532–534, 2012.
- 8 Sergey Bereg, Ferran Hurtado, Mikio Kano, Matias Korman, Dolores Lara, Carlos Seara, Rodrigo I. Silveira, Jorge Urrutia, and Kevin Verbeek. Balanced partitions of 3-colored geometric sets in the plane. *Discrete Applied Mathematics*, 181:21–32, 2015.
- 9 P Bonsma, Th Epping, and Winfried Hochstättler. Complexity results on restricted instances of a paint shop problem for words. *Discrete Applied Mathematics*, 154(9):1335–1343, 2006.
- 10 Xi Chen, Decheng Dai, Ye Du, and Shang-Hua Teng. Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 273–282. IEEE, 2009.
- 11 Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009.
- 12 Xi Chen, David Durfee, and Anthi Orfanou. On the complexity of nash equilibria in anonymous games. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 381–390, 2015.

13:14 The Complexity of Sharing a Pizza

- 13 Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. The complexity of non-monotone markets. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 181–190, 2013.
- 14 Man-Kwun Chiu, Aruni Choudhary, and Wolfgang Mulzer. Computational Complexity of the α -Ham-Sandwich Problem. In *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- 15 Bruno Codenotti, Amin Saberi, Kasturi Varadarajan, and Yinyu Ye. The complexity of equilibria: Hardness results for economies via a correspondence with games. *Theoretical Computer Science*, 408(2-3):188–198, 2008.
- 16 Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- 17 Argyrios Deligkas, John Fearnley, and Themistoklis Melissourgos. Pizza sharing is PPA-hard. *arXiv preprint*, 2020. [arXiv:2012.14236](https://arxiv.org/abs/2012.14236).
- 18 Argyrios Deligkas, John Fearnley, Themistoklis Melissourgos, and Paul G Spirakis. Computing exact solutions of consensus halving and the Borsuk-Ulam theorem. *Journal of Computer and System Sciences*, 117:75–98, 2021.
- 19 Herbert Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag New York, Inc., New York, NY, USA, 1987.
- 20 Edith Elkind, Leslie Ann Goldberg, and Paul Goldberg. Nash equilibria in graphical games on trees revisited. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 100–109, 2006.
- 21 Kousha Etessami and Mihalis Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597, 2010.
- 22 John Fearnley, Spencer Gordon, Ruta Mehta, and Rahul Savani. Unique end of potential line. *Journal of Computer and System Sciences*, 114:1–35, 2020.
- 23 Aris Filos-Ratsikas and Paul W Goldberg. Consensus halving is PPA-complete. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 51–64, 2018.
- 24 Aris Filos-Ratsikas and Paul W Goldberg. The complexity of splitting necklaces and bisecting ham sandwiches. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 638–649, 2019.
- 25 Charles R. Hobby and John R. Rice. A moment problem in L_1 approximation. *Proceedings of the American Mathematical Society*, 16(4):665–670, 1965. URL: <http://www.jstor.org/stable/2033900>.
- 26 Alfredo Hubbard and Roman Karasev. Bisecting measures with hyperplane arrangements. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 169, pages 639–647. Cambridge University Press, 2020.
- 27 Mikio Kano and Jorge Urrutia. Discrete geometry on colored point sets in the plane—a survey. *Graphs and Combinatorics*, 37(1):1–53, 2021.
- 28 Roman N Karasev, Edgardo Roldán-Pensado, and Pablo Soberón. Measure partitions using hyperplanes with fixed directions. *Israel Journal of Mathematics*, 212(2):705–728, 2016.
- 29 Christian Knauer, Hans Raj Tiwary, and Daniel Werner. On the computational complexity of Ham-Sandwich cuts, Helly sets, and related problems. In *Symposium on Theoretical Aspects of Computer Science (STACS2011)*, volume 9, pages 649–660, 2011.
- 30 Stefan Langerman. personal communication, 2017.
- 31 Chi-Yuan Lo, Jiří Matoušek, and William Steiger. Algorithms for ham-sandwich cuts. *Discrete & Computational Geometry*, 11(1):433–452, 1994.
- 32 Chi-Yuan Lo and William Steiger. An optimal time algorithm for ham-sandwich cuts in the plane. In *Proc. 2nd Canadian Conference on Computational Geometry (CCCG 1990)*, pages 5–9, 1990.
- 33 Jiri Matoušek. *Using the Borsuk-Ulam Theorem: Lectures on Topological Methods in Combinatorics and Geometry*. Springer Publishing Company, Inc., 2007.

- 34 Ruta Mehta. Constant rank bimatrix games are PPAD-hard. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 545–554, 2014.
- 35 Frédéric Meunier. Discrete splittings of the necklace. *Mathematics of Operations Research*, 33(3):678–688, 2008.
- 36 Christos H Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and system Sciences*, 48(3):498–532, 1994.
- 37 Alexander Pilz and Patrick Schnider. Bisecting three classes of lines. *Computational Geometry*, 98:101775, 2021.
- 38 Edgardo Roldan-Pensado and Pablo Soberon. A survey of mass partitions. *arXiv preprint*, 2020. [arXiv:2010.00478](https://arxiv.org/abs/2010.00478).
- 39 Aviad Rubinfeld. Inapproximability of Nash equilibrium. *SIAM Journal on Computing*, 47(3):917–959, 2018.
- 40 Patrick Schnider. Equipartitions with Wedges and Cones. *arXiv preprint*, 2019. [arXiv:1910.13352](https://arxiv.org/abs/1910.13352).
- 41 Patrick Schnider. Ham-sandwich cuts and center transversals in subspaces. *Discrete & Computational Geometry*, 64(4):1192–1209, 2020.
- 42 Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. Finding clearing payments in financial networks with credit default swaps is PPAD-complete. *LIPICs: Leibniz International Proceedings in Informatics*, 67, 2017.
- 43 William Steiger and Jihui Zhao. Generalized ham-sandwich cuts. *Discrete & Computational Geometry*, 44(3):535–545, 2010.
- 44 A. H. Stone and J. W. Tukey. Generalized “sandwich” theorems. *Duke Math. J.*, 9(2):356–359, June 1942.
- 45 Vijay V Vazirani and Mihalis Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. *Journal of the ACM (JACM)*, 58(3):1–25, 2011.