Some Combinatorial and Algorithmic Applications of the Borsuk-Ulam Theorem

Sambudda Roy
Computer Science
Rutgers University

William Steiger
Computer Science
Rutgers University

September 19, 2005

Abstract

The Borsuk-Ulam theorem has many applications in algebraic topology, algebraic geomtry, and combinatorics. Here we study some combinatorial consequences, typically asserting the existence of a certain combinatorial object. We also focus on the computational complexity of algorithms that search for the object.

1 Introduction and Summary

The Borsuk-Ulam theorem states that if f is a continuous function from the unit sphere in R^n into R^n, there is a point x where $f(x) = f(-x)$; i.e., some pair of antipodal points has the same image. The recent book of Matoušek [21] is devoted to explaining this theorem, its background, and some of its many consequences in algebraic topology, algebraic geometry, and combinatorics. Borsuk-Ulam is considered a great theorem because it has several different equivalent versions, many different proofs, many extensions and generalizations, and many interesting applications.

A familiar consequence is the ham-sandwich theorem (given d finite continuous measures on R^d, there exist a hyperplane that simultaneously bisects them), along with some of its extensions and generalizations to partitioning continuous measures [2], [6], [7], [8], [10]. In many cases we can derive combinatorial statements that give discrete versions of these results. This, in turn, raises algorithmic issues about the computational complexity of finding the asserted combinatorial object. For example Lo et. al. [20] gave a direct combinatorial proof of the discrete ham-sandwich theorem and described algorithms to compute ham-sandwich cuts for point sets. Various generalizations and extensions were considered in [1], [2], [3], [4], [5], [9], [10], [11], [12], [17], [18], [19], [23], and [25].

A recent interesting example extends a result of Bárány and Matoušek [7], who combined Borsuk-Ulam with equivariant topology to show that three finite, continuous measures on R^2 can be equipartitioned by a *2-fan*, the region spanned by two half-lines incident at a point. Bereg [10] strengthened this statement, proved a discrete version for measures concentrated on a given set of points, and described a beautiful algorithm to find such a partitioning. In Section 2 we show his algorithm to be nearly optimal via a lower bound for this task.

In Section 3 we study equitable partitions of a set of points in R^2 by a pair of lines, and in Section 4 we consider some other partioioning questions, including the necklace problem.

2 Equipartitioning Two Measures by a 3-Fan

A *two-fan* in the plane is a point P (called the center) and two rays, ρ_1 and ρ_2 incident with P. This structure partitions R^2 into two connected regions. Bárány and Matoušek [7] had proved that given three finite, measures on the plane, there is always an equitable partition by a two fan; i.e., there exists a two fan whose two regions have exactly half of each measure. Bereg [10] later considered a discrete version and proved that there are many two fans with equitable partitions of a given input point set. Specifically, given $2r$ red points, $2b$ blue points, and $2g$ green points in general position in R^2, and a line ℓ, there exists a point $P \in \ell$ and a two fan centered at P for which there are r red, b blue, and g green points in both of the regions induced by the fan. He described an $O(n(\log n)^2)$ algorithm to find such a two fan, $n = r + b + g$ being one half the number of points in the problem. Here we show the algorithm to be nearly optimal, by proving

Lemma 1 *Let S be a given set of points in R^2, an of them red, bn of them blue, and cn of them green. For a given point P, $\Omega(n \log n)$ steps are required by any algebraic decision tree that can decide if there is an equitable two-fan for S with center at P.*

Proof: Let T be an algebraic decision tree that can decide for a set S with $\Theta(n)$ data points, and a point P, whether there is an equitable 2-fan for S centered at P. We will show that the subset X of inputs where T returns a **NO** answer has at least $n!$ path connected components.

We take $P = (0, 0)$ and data points $Q_i = (r_i \cos(\theta_i), r_i \sin(\theta_i))$, $r_i \neq 0$. For each such point we only need its argument $\theta \in (0, 2\pi)$, since this alone determines whether the point is in some 2-fan centered at P. Inputs will consist of $N = 16n - 8$ points, described by the components of $\underline{z} = (z_0, \ldots, z_{N-1}) \in R^N$. Points z_j are blue if $j = 0$ or 7 mod 8, red if $j = 1, 4, 5$ or 6 mod 8, and green if $j = 2$ or 3 mod 8. Thus each input describes a set S with $8n - 4$ red points, $4n - 2$ blue points, and $4n - 2$ green points.

The canonical input is the point $\underline{z}^* = (\theta_0, \ldots, \theta_{N-1}) \in R^N$, where

$$\theta_i = \frac{\pi}{2} + \left(\frac{i+1}{N+1}\right) 2\pi, \quad i = 0, \ldots, N. \tag{1}$$

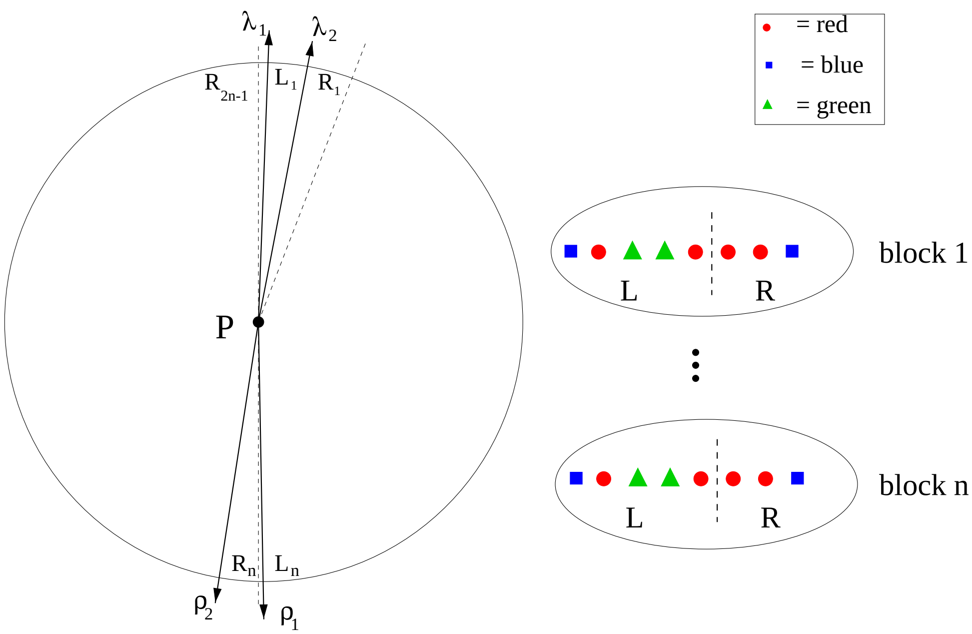

Figure 1. The canonical input $\underline{z}^*$.

These $16n-8$ points are divided into $2n-1$ blocks of 8 points each. The blocks have a LEFT part and a RIGHT part. Block i has left part $L_i = [\theta_{8i-8}, \theta_{8i-7}, \theta_{8i-6}, \theta_{8i-5}, \theta_{8i-4}]$ and right part $R_i = [\theta_{8i-3}, \theta_{8i-2}, \theta_{8i-1}]$ for which, replacing the entries by their colors, is $L_i = [b\,r\,g\,g\,r]$ and $R_i = [r\,r\,b]$.

Let λ_1 be a ray separating R_{2n-1} and L_1, and ρ_1, a ray separating L_n and R_n. This two-fan has $n-1$ blocks and the left half of block n. Also it contains $4n-2$ red points (half), $2n-1$ blue points (half) and $2n$ green points (one more than half). The next (clockwise) two-fan that equitably partitions the red and the blue points is λ_2 (a ray separating L_1 and R_1) and ρ_2 (a ray separating R_n and L_{n+1}) but now there are $n-2$ green points (one less than half).

In fact ALL two-fans that split both red and blue points evenly are either deficient by one green or in excess by one green. Specifically, for each $j = 1, \ldots, n$, we have the two-fans with rays λ_{2j-1} and ρ_{2j-1} containing blocks $[L_j\, R_j\, \cdots\, R_{j+n-2}\, L_{j+n-2}]$ ($2n$ green points); also for each $j = 1, \ldots, n$, we have the two-fans with rays λ_{2j} and ρ_{2j} containing blocks $[R_j\, L_{j+1}\, \cdots\, L_{j+n-1}\, R_{j+n-1}]$ ($2n-2$ green points). Thus, with input $\underline{z}^*$, T must return a NO answer to the equitable two-fan query.

We consider only a restricted set I of inputs to T. A point $\underline{z} = (z_0, \ldots, z_{N-1}) \in I$ if $z_j = \theta_j$ for $j \neq 2 \bmod 16$, and otherwise $z_j \in (0, 2\pi)$; only the first green point in the odd numbered blocks is free. For a permutation $\underline{\pi} = (\pi_1, \ldots, \pi_n)$ of $(1, \ldots, n)$, define $\underline{z}_\pi \in I$ by

$$z_{16j-14} = \theta_{16\pi_j - 14};$$

$\underline{z}_\pi$ describes the same N points as $\underline{z}^*$ except that the first green points of the odd-blocks of $\underline{z}^*$ appear in permuted order in $\underline{z}_\pi$. Clearly $\underline{z}_\pi \in X$ for every permutation. We also claim that if $\underline{\pi}$ and $\underline{\rho}$ are distinct permutations, then $\underline{z}_\pi$ and $\underline{z}_\rho$ are in different path connected components of I. We move along a continuous path p in I from $\underline{z}_\pi$ to $\underline{z}_\rho$

(holding z_j fixed, $j \neq 2 \bmod 16$). As we do, some green point leaves its half-block, say L_k, and moves to an adjacent half-block, R_{k-1} or R_k. Let j be the block where this first occurs on p and let $\underline{z}(t) \in p$ denote the corresponding point in R^N. The input described by this point is no longer in X because either $[R_{k-n-2}\, L_{k-n-1} \cdots L_{k-1}\, R_{k-1}]$ or $[R_k\, L_{k+1} \cdots L_{k+n-1}\, R_{k+n-1}]$ are now equitable partitions of all three colors. Therefore $X \cap I$ has $n!$ path connected components, and the lemma now follows by Ben-Or's theorem [16]. ■

The claimedlower bound for finding an equitable 2-fan is somewhat misleading: it may be possible in $o(n \log n)$ time to determine a point Q, and two lines incident at Q for which one of the four 2-fans is equitable. Lemma 1 just says that given only Q, $n \log n$ steps are needed to know if there is an equitable 2-fan with apex Q, and to find one if YES. Thus any search algorithm (like Bereg's) that tests candidate apex points P must have complexity at least $n \log n$.

3 Equitable Partitioning by Orthogonal Lines

Given n points in general position in R^2, Willard [24] asked for a pair of non-parallel lines ℓ_1 and ℓ_2 that equitably partition the points; i.e., in each of the four open quadrants they define, there are at most $n/4$ points. An efficient algorithm for this was implied by results in Cole, Sharir, and Yap [14], and an optimal $O(n)$ algorithm follows immediately using Megiddo's separated, discrete ham-sandwich cut [22]. In fact we can even insist that the lines are orthogonal: this is implied by a result of Bárány and Matoušek [8]) that uses Borsuk-Ulam along with equivariant topology. Here we give an easy, direct combinatorial proof that there is an orthogonal four-partitioning, and we describe the computational complexity of finding one. Specifically we prove

Theorem 1 *Given a set S of n points in general position in R^2, there exist orthogonal lines ℓ_1 and ℓ_2 that equipartition S, and they may by found in $\Theta(n \log n)$ RAM steps.*

Proof: A halving line for S has at most $|S|/2$ points in its open halfspaces. For the existence, w.l.o.g. we may assume $n = 4j + 1$ is odd and start with ℓ_1 as a vertical halving line incident with the point $P \in S$ with median $x-$coordinate, and ℓ_2 as a horizontal halving line incident with the point $Q \in S$ of median $y-$coordinate. Also suppose the open upper left quadrant has the most points, say a, which we assume is $> j$ or this partition is already equitable. We will rotate ℓ_1 and ℓ_2 counter clockwise through $\pi/2$ radiands, always keeping them orthogonal, and keeping ℓ_1 a halving line: i.e., ℓ_1 rotates about P until it first meets a point - say $P_1 \in S$. Next we rotate about P_1 until ℓ_1 meets meets another point, P_2, etc. Except for the moments when ℓ_1 is incident with two points, P_i and P_{i+1}, it is always a halving line for S. During this process, as ℓ_2 passes points of S, we will move it (always maintaining its orthogonality to ℓ_1) so at most half the points of S are in either of *its* open halfspaces. At the end of the rotation the upper-left quadrant has become the original lower left, and now has $< j$ points (because ℓ_2 is halving). Since its cardinality changes by ± 1 at each "event" in the rotation, there is a position where it has exactly j points.

The complexity statement follows from the following two results.

Lemma 2 *Given a set S of $\Theta(n)$ points in general position in R^2 and a point $Q \in R^2$, $\Omega(n \log n)$ steps are required by any algebraic decision tree that can decide if there is an equitable partitioning of S by orthogonal lines incident with Q.*

Proof: The argument is a construction sharing several features with that of Lemma 1, so we will be terse. We take Q to be the origin. Let $N = 32k + 9$ and take N points on the unit circle with arguments given by

$$\theta_j = \frac{2\pi j}{N},\ j = 1, \ldots, N;$$

since N is odd, no two are antipodal. The points of S will be those θ_j where $j = 1, 2, 3, 4, 5 \bmod 8$, so $n = |S| = 20k + 5$. (It may help to think of $4k + 1$ groups of equally-spaced points, 8 points per group, plus one extra point. Each open quadrant has $k + 1/4$ groups. Points of S occupy the first 5 places in a group; the last 3 are empty.)

If two orthogonal lines through $Q = (0, 0)$ equipartition S there can be at most $5k + 1$ points in any open quadrant. Let ℓ_1 and ℓ_2 be orthogonal lines through Q. It is easy to see that

1. as they are rotated about Q (maintaining orthogonality), if *neither* is incident with a point of S, then exactly two quadrants each contain $5k + 2$ points of S, and the other two contain a total of $10k + 1$ points;

2. in addition, as either ℓ_1 (or ℓ_2) rotates across $\theta_i, \ldots, \theta_{i+7}$, there is a position where the four open quadrants contain $5k + 2$, $5k + 2$, $5k + 1$, and $5k$ pts.

Thus, at the canonical input

$$\underline{z}^* = (\theta_1, \ldots, \theta_5, \theta_9, \ldots, \theta_{13}, \ldots, \theta_{32k+1} \ldots, \theta_{32k+5}) \in R^{20k+5},$$

the decision tree must answer NO. On the other hand take k even and consider the restricted set of inputs I where $z_j \in (0, 2\pi)$, and $z_i = \theta_i$ for $i \neq 5 \bmod(16)$. For each $\underline{\pi} = (\pi_1, \ldots, \pi_{2k})$, a permutation of $(1, \ldots, 2k)$, define $\underline{z}_\pi \in I$ by

$$z_{16j-5} = \theta_{16\pi_j - 5};$$

$\underline{z}_\pi$ and $\underline{z}_\rho$ are in different connected components because on a continuous path $p(t)$ in I from $\underline{z}_\pi$ to $\underline{z}_\rho$, one of the middle points in an even numbered group is first to enter an adjacent group, and at the input described by that $p(t)$, there is an allignment of ℓ_1 and ℓ_2 that is incident with a point of S, and where each open quadrant has at most $5k + 1$ points of S, a YES input. ■

As with Lemma 1, it may be possible in $o(n \log n)$ time to determine orthogonal lines ℓ_1, ℓ_2 that equipartition the n given data points. The Lemma just says that given only Q, $n \log n$ steps are needed to know if there exist orthogonal lines incident at Q which do the job, so any search-based algorithm that tests candidate points must have complexity at least $n \log n$.

Lemma 3 *Given a set S with n points in general position in R^2, in $O(n \log n)$ RAM steps we can find orthogonal lines ℓ_1 and ℓ_2 that equitably partition the points.*

Proof: Dualizing the existence proof in Theorem 1, there is a point $P_1 = (-a, y_1)$, $a > 0$, that is dual to ℓ_1, and a point $P_2 = (1/a, y_2)$, dual to ℓ_2 ($\ell_2 \perp \ell_1$), so that (i) at most half of the lines in $\mathcal{L}$ (dual to the n points in S) are above (below) P_1, (ii) at most half are above (below) P_2, and (iii) $\lfloor n/4 \rfloor$ are above *both* P_1 and P_2. We will search for P_1, starting with a point $Q = (-c, d)$ on the median level, and $c > 0$ chosen so that Q is to the left of the vertex of $A(\mathcal{L})$ with min x-coordinate. The cost for Q is $O(n \log n)$ and we test it in $O(n)$ time by computing $Q' = (-1/c, d')$ on the median level, and counting $N(Q)$, the number of lines above *both* Q and Q', stopping with $P_1 = Q$ if $N(Q) = \lfloor n/4 \rfloor$.

We could now use slope selection to perform a binary search on the vertices of $A(\mathcal{L})$ and after $\log \binom{n}{2}$ search steps, each evaluating $N(\cdot)$, we determine P_1 at a cost of $O(n(\log n)^2)$. On the other hand in linear time we could determine vertical lines $x = t_1$ and $x = t_2$ with the property that P_1 is in the strip they determine, but there are at most εn^2 vertices of $A(\mathcal{L})$ that are also within the strip, $\varepsilon > 0$ small. To do this we evaluate N at a point on $x = t_1$ and the median level and also at a point on $x = t_2$ and the median level, and accept (t_1, t_2) if N is bigger than $n/4$ at one point and smaller at the other. Now it is easy to prune a constant fraction of the lines of $\mathcal{L}$ which cannot determine P_1 and then recursing within the strip on $\mathcal{L}'$, the unpruned lines, we obtain P_1 in time $O(n \log n)$: the strip in the next recursive step is determined in time $O(n)$, because we evaluate $N(\cdot)$ with respect to the original lines, and there are $O(\log n)$ steps before only a constant number of lines remain for an $O(n)$ brute force conclusion. ■

In the special case where the n points of S are in convex position, their radial order is the same from every point interior to conv(S). This fact simplifies the problem enough to allow us to find an orthogonal equipartitioning in linear time.

Lemma 4 *An equitable orthogonal partitioning for n points in convex position can be found in time $O(n)$.*

Proof: We may consider the points to be on the unit circle. From an arbitrary point Q with $|Q| = 1$, draw a halving line ℓ through Q. Next find another halving line ℓ' for S orthogonal to ℓ, and count the number of points in the four reulting quadrants. The counts have to be a, $n/2 - a$, a, $n/2 - a$. Now, choose a quadrant which has more than $n/4$ points and select the middle point in it. Now we want to draw another pair of orthogonal lines through this point - we halve the set through this point (the green line in Fig. 1), and then we draw another halving line perpendicular to this direction.

We again measure how good this new partitioning is : if its not an equipartitioning, we can mark out sets A and B (and their antipodal sets A_1, B_1 of the same corresponding sizes) such that sets A and B have at least $n/8$ points each. Depending on the present partitioning and the previous one, we can decide which way to conduct our binary search (in Fig. 2 the direction of search is indicated by the arrow).

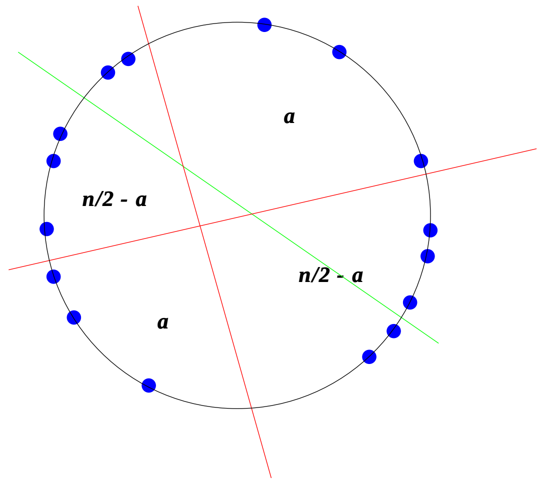

Figure 1: Orthogonal partitioning

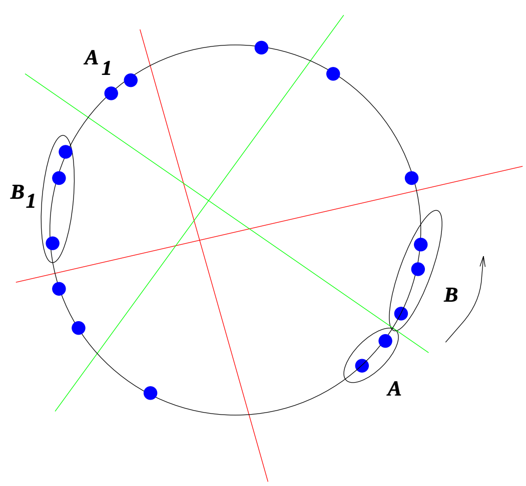

Figure 2: Orthogonal partitioning

But now we observe, that the set A_1 in the figure (of size at least $n/4$) will not ever enter into calculations any more (that is, for the solution orthogonal partitioning, it belongs wholly to one quadrant).

So, from now on, we don't need to count the points of A_1 separately. If the arrow were in the other direction, we could have discounted B_1 from future computations. In either case, we can prune out $n/8$ of the points.

One more iteration would go thus : delete the points of A_1 from consideration and look for a "halving" line which partitions the points into $n/2$ and $n/2-|A_1|$. Also look for a similar line in the perpendicular direction, and continue.

So, we see, that at each step we can prune away a constant fraction of the points ($1/8^{th}$). Also each step takes linear time (two halvings of point sets and a selection). Altogether, this takes linear time hence. ■

References

[1] J. Akiyama and N. Alon. Disjoint Simplices and Geometric Hypergraphs. in :*Combinatorial Mathematics; Proc. of the Third International Conference , New York, NY 1985 (G. S. Blum, R. L. Graham and J. Malkevitch, eds.)*, Annals of the New York Academy of Sciences, Vol. 555, 1-3 (1989).

[2] N. Alon. Some Recent Combinatorial Applications of Borsuk-Type Theorems. in *Algebraic, Extremal and Metric Combinatorics, M. M. Deza, P. Frankl and I. G. Rosenberg eds.*, Cambridge Univ. Press, Cambridge, England, 1-12 (1988).

[3] N. Alon. Splitting Necklaces. *Advances in Mathematics 63*, 247-253 (1987).

[4] N. Alon and D. West. The Borsuk-Ulam Theorem and Bisection of Necklaces. *Proc. Amer. Math. Soc. 98*, 623-628 (1986).

[5] D. Avis. On the Partitionability of Point Sets in Space. *Proc. First ACM Symposium on Computational Geometry*, 116-120, 1985.

[6] I. Bárány. Geometric and Combinatorial Applications of Borsuk's Theorem. *New Trends in Discrete and Computational Geometry, J. Pach, ed. (vol 10 of Algorithms and Combinatorics)*, 235-249, Springer-Verlag, Berlin (1993).

[7] I. Bárány and J. Matoušek. Simultaneous Partitions of Measures by k-Fans. *Discrete and Computational Geometry 25*, 317-334 (2001). v

[8] I. Bárány and J. Matoušek. Equipartition of Two Measures by a 4-Fan. *Discrete and Computational Geometry 27*, 293-301 (2002).

[9] S. Bespamyatnikh, D. Kirkpatrick, and J. Snoeyink. Generalizing Ham-Sandwich Cuts to Equitable Subdivisions. *Discrete and Computational Geometry 24*, 605-622 (2000).

[10] S. Bereg. Equipartitions of Measures by 2-Fans. *Discrete and Computational Geometry* , (2005).

[11] P. Bose, E. Demaine, F. Hurtado, J. Iaconc, S. Langerman, and P. Morin. Geodesic Ham-Sandwich Cuts. *Proc.* 20^{th} *ACM Symp. on Computational Geometry*, 1-9 (2004).

[12] P. Bose and S. Langerman. Weighted Ham-Sandwich Cuts. *Lecture Notes in Computer Science (JCDCG 2004: Japan Conference on Discrete and Computational Geometry, Tokyo, Japan, M. Kano and J. Akiyama (Eds.))*, (to appear), Springer-Verlag, (2005).

[13] R. Cole, J. Salowe, W. Steiger, and E. Szemerédi. An Optimal Time Algorithm for Slope Selection. *SIAM J. Comp. 18*, (1989) 792-810.

[14] R. Cole, M. Sharir, and C. Yap. On k-Hulls and Related Topics. *SIAM J. Computing 16*, 61–77, 1987.

[15] D. Dobkin and H. Edelsbrunner. Ham-Sandwich Theorems Applied to Intersection Problems. *Proc.* 10^{th} *International Workshop on Graph Theoretic Concepts in Comp. Sci.*, 88-99 (1984).

[16] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag, Berlin, 1987.

[17] H. Ito, H. Uehara, and M. Yokoyama. Two Dimensional Ham-Sandwich Theorem for Partitioning into Three Convex Pieces. *Lecture Notes in Computer Science 1763 (JCDCG 1998: Japan Conference on Discrete and Computational Geometry, Tokyo, Japan, M. Kano and J. Akiyama (Eds.))*, Springer-Verlag, (2000).

[18] A. Kaneko and M. Kano. Balanced Partitions of Two Sets of Points in the Plane. *Comp. Geo.. Theory and Application 13*, 253-261 (1999).

[19] S. Langerman and W. Steiger. Optimization in Arrangements. *Lecture Notes in Computer Science 2607 (STACS 2003: 20th Annual Symposium on Theoretical Aspects of Computer Science, Berlin, Germany, H.Alt, M.Habib (Eds.))*, 50-61, Springer-Verlag, (2003).

[20] C.-Y. Lo, J. Matoušek and W. Steiger. Algorithms for Ham-sandwich Cuts. *Discrete & Computational Geometry 11*, 433–452, 1994.

[21] J. Matoušek. *Using the Borsuk-Ulam Theorem*. Springer-Verlag, Heidelberg, (2003).

[22] N. Megiddo. Partitioning with Two Lines in the Plane. *J. Algorithms 6*, 430–433, 1985.

[23] T. Sakai. Balanced Convex Partitions of Measures in R^2. *Graphs and Combinatorics 18*, 169-192 (2002).

[24] D. Willard. Polygon Retrieval. *SIAM J. Computing 11*, 149–165, 1982.

[25] F. Yao, D. Dobkin, H. Edelsbrunner, and M. Paterson. Partitioning Space for Range Queries. *SIAM J. Computing 18*, 371-384 (1989).