On the Computational Complexity of some problems from Combinatorial Geometry

Christian Knauer[1] ⋆, Hans Raj Tiwary[2], and Daniel Werner[3] ⋆⋆

[1] Institut für Informatik, Universität Bayreuth, Germany
christian.knauer@uni-bayreuth.de
[2] Institut füt Mathematik, Technische Universität Berlin, Germany
tiwary@math.tu-berlin.de
[3] Institut für Informatik, Freie Universität Berlin, Germany,
dwerner@inf.fu-berlin.de

Abstract. We study several canonical decision problems that arise from the most famous theorems from combinatorial geometry. We show that these are W[1]-hard (and NP-hard) if the dimension is part of the input by fpt-reductions (which are actually ptime-reductions) from the d-Sum problem. Among others, we show that computing the minimum size of a *Caratheodory set* and a *Helly set* and certain decision versions of the *Ham-Sandwich cut problem* are W[1]-hard. Our reductions also imply that the problems we consider cannot be solved in time $n^{o(d)}$ (where n is the size of the input), unless the Exponential-Time Hypothesis (ETH) is false.
Our technique of embedding d-Sum into a geometric setting is conceptually much simpler than direct fpt-reductions from purely combinatorial W[1]-hard problems (like the clique problem) and has great potential to show (parameterized) hardness and (conditional) lower bounds for many other problems.

Keywords: combinatorial geometry, ham-sandwich cuts, parameterized complexity, geometric dimension, exponential-time hypothesis.

1 Introduction

Many theorems from combinatorial geometry are of the following type: Given n objects that have a certain property, then there are already $d+1$ of them that have this property. Two examples of this are Caratheodory's Theorem [6] and Helly's Theorem [19].

Caratheodory's Theorem states, in one of its several formulations, that whenever a point p is contained in the convex hull of a point set in $\mathbb{R}^d$, then it is already contained in the convex hull of a subset of size at most $d+1$. A minimal set containing p in the convex hull is called a *Caratheodory set* for p. Here, we will consider the corresponding decision problem:

⋆ This research was supported by the German Science Foundation (DFG) under grant Kn 591/3-1.
⋆⋆ This research was funded by Deutsche Forschungsgemeinschaft within the Research Training Group (Graduiertenkolleg) "Methods for Discrete Structures"

Definition 1. *(d-CARATHEODORY-SET) Given a point set in $\mathbb{R}^d$, are there d points whose convex hull contains the origin?*

Stated in a dual setting, this gives another well known theorem: If n convex sets in $\mathbb{R}^d$ have an empty intersection, then by Helly's Theorem there are already $d+1$ whose intersection is empty. The decision problem can be stated as follows:

Definition 2. *(d-HELLY-SET) Given n convex sets $P_1, \ldots, P_n$ in $\mathbb{R}^d$, do any d of them have an empty intersection?*

The canonical decision versions of Caratheodory's and Helly's Theorem have not explicitly been considered in the literature so far. This is quite surprising, as they are interesting to people from computational as well as discrete geometry. However, they have been studied in the context of Linear Programming as the problem of finding the smallest minimal set of inequalities that is infeasible.

Definition 3. *(d-MIN-IIS) Given n inequalities in $\mathbb{R}^d$, do any d of them have an empty intersection?*

The d-MIN-IIS has been studied before, mainly because of its connection to the NP-complete MAXIMUM-FEASIBLE-SUBSYSTEM problem, where one is given an infeasible linear program and one has to find a feasible subsets of constraints of maximum size. Amaldi et al. [2] show that d-MIN-IIS is NP-hard by a (transitive) reduction from DOMINATING-SET. However, the dimension depends on the number of elements, so it does not reveal anything with respect to this parameter d.

The Ham-Sandwich Theorem as a corollary of the Borsuk-Ulam Theorem (see, e.g., Matoušek [24]) states that for any d finite point sets in $\mathbb{R}^d$ there is a hyperplane that bisects all of the sets at once, i.e., has at most half of the points on each side. Computing a ham-sandwich cut efficiently is an important problem and has been studied extensively (see Edelsbrunner and Waupotitsch [10], Matoušek et al. [25], Yu [31]). For general dimensional, the fastest known algorithm [25] runs in time roughly $O\left(n^{d-1}\right)$.

The Ham-Sandwich problem is not a decision problem, as, given an instance, we know that there always exists a solution, but still it is not known how to find it efficiently. Such problems are captured by the complexity class PPAD, see Papadimitrou [28]. It is an important open question whether computing a ham-sandwich cut is PPAD complete. In this paper we show that a natural "incremental" approach for computing the ham-sandwich cut will not work unless $W[1] = P$.

One way to find a ham-sandwich cut incrementally could be to take any point, decide whether there is some ham-sandwich cut through it, and perform a dimension reduction until the hyperplane is determined. This gives rise to the following decision problem:

Definition 4. *(d-HAM-SANDWICH) Given d sets $P_1, \ldots, P_d$ in $\mathbb{R}^d$ and a point $a \in \mathbb{R}^d$, is there a Ham-Sandwich cut that passes through a?*

We show that d-Ham-Sandwich is $W[1]$-hard and therefore most likely no polynomial algorithm (FPT or otherwise) exists for this problem.

Our reductions use a new technique of embedding of d-Sum into the d-dimensional space. Thereto, a d-Sum instance is encoded into sets of points (or hyperplanes, respectively), and the property of d elements summing up to 0 is expressed by an equivalent geometric property of the point set, e.g., allowing a ham-sandwich cut through the origin.

Overview. The main results of this paper are following:

In Sec. 3, 4 we prove the following:

Theorem 1. *The problems d-Caratheodory-Set and d-Helly-Set are W[1]-hard with respect to the parameter d and NP-hard.*

All these proofs follow from slight modifications of the hardness proof for the first problem (FIRST PROBLEM?) ¡— REMOVE?.

Using the same proof techniques we obtain following corollaries.

Corollary 1. *The problem d-Min-IIS is W[1]-hard with respect to the dimension.*

Observe that this problem becomes polynomial-time solvable if we ask for $d+1$ halfspaces by first solving the corresponding linear program and afterwards applying Helly's Theorem.

Corollary 2. *Deciding whether a point q is in general position*[4] *with respect to P is W[1]-hard with respect to d and NP-hard.*

For the Ham-Sandwich problem, a little more work has to be done. The resulting sets have to be *balanced*, such that ham-sandwich cuts through the origin correspond to d elements summing up to 0. After this construction, we are able to derive the following result:

Theorem 2. *The d-Ham-Sandwich cut problem is W[1]-hard with respect to the dimension and NP-hard.*

Combining our reductions with a result of Pătraşcu and Williams [29], Theorems 1 and 2 immediately gives:

Corollary 3. *The problems d-Caratheodory-Set, d-Helly-Set and d-Ham-Sandwich cannot be solved in time $n^{o(d)}$ (where n is the size of the input), unless the Exponential-Time Hypothesis (ETH) is false*[5]*.*

[4] No hyperplane that contains d points from P also contains q.

[5] The Exponential Time Hypothesis [20] conjectures that n-variable 3-CNFSAT cannot be solved in $2^{o(n)}$-time.

Related work. The study of computational variants of theorems from discrete geometry is not new. Several problems that arise from theorems in discrete geometry have received a lot of attention, most notably computation of (approximate) center- and Tverberg points in the plane as well as in higher dimension. In the plane, surprisingly one can compute a centerpoint in linear time [21], whereas for higher dimension, no exact polynomial time algorithms are known. See Agarwal et al. [1] and Miller and Sheehy [26] for recent progress. There too, like in the present paper, the corresponding decision problem is considered, i.e., to decide whether a given point is a center point. This problem has been shown to be co-NP complete by Teng [30] if d is part of the input.

We have already discussed the previous studies about the computation of ham-sandwich cuts. Perhaps surprisingly, the computation of smallest sets arising from Caratheodory's and Helly's theorem has not been explicitly studied even though it has been studied under the guise of IIS in the context of Linear Programming.

Even though the dimension of geometric problems is a natural parameter for studying their parameterized complexity, only relatively few results of this type are known: Langerman and Morin [22] gave fpt-algorithms for the problem of covering points with hyperplanes, while the problem of computing the volume of the union of axis parallel boxes has been shown to be W[1]-hard by Chan [7]. Cabello et al. [5, 4] have developed a technique that has been applied succesfully to show W[1]-hardness for a number of problems from various application areas like shape matching [3], clustering [4, 16], and discrepancy-computation [17]. We refer to Giannopoulos et al. [18] for a survey on other parameterized complexity results for geometric problems.

For a general introduction to combinatorial geometry, we recommend Matoušek [23] and Ziegler [32].

Parameterized complexity. Parameterized complexity theory provides a framework for the study of algorithmic problems by measuring their complexity in terms of one or more parameters, explicitly or implicitly given by their underlying structure, in addition to the problem input size. For an introduction to the field of parameterized complexity theory, we refer to the textbooks of Flum and Grohe [14], Niedermeier [27] and Downey and Fellows [9].

The dimension d of geometric problems in $\mathbb{R}^d$ is a natural parameter for studying their parameterized complexity. In terms of parameterized complexity theory the question is whether these problems are *fixed-parameter tractable* with respect to d. Proving a problem to be W[1]-hard with respect to d, gives a strong evidence that an fpt-algorithm (i.e., an algorithm that runs in time $O\left(f(d) \cdot n^c\right)$ for some fixed c and an arbitrary function f) does not exist. W[1]-hardness is often established by fpt-reductions from the clique problem in general graphs, which is known to be W[1]-complete [9]. Below we use a different approach by giving *conceptually much simpler* fpt-reductions from the d-Sum problem [15, 12]:

Definition 5. *(d-Sum) Given n integers, are there d (not necessarily distinct) numbers that sum up to 0?*

This problem is NP-hard [12] and can be solved in (roughly) $O(n^{d/2})$ time. It can be shown to be W[1]-hard with respect to d from a simple reduction from the subset-sum problem which was shown to be W[1]-hard by Downey and Koblitz [13]. Recently it has been shown [29] (without using parameterized complexity explicitly) that, unless the ETH fails, d-Sum cannot be solved in time $n^{o(d)}$.

Reductions from 3-Sum seem somewhat more "natural" for computational geometers: Gajentaan and Overmars [15] introduced the 3-Sum problem for the purpose of arguing that certain problems in planar geometry "should" take $\Omega(n^2)$ time; showing 3-Sum-hardness for such problems is considered a routine task today. Surprisingly this approach apparently has not been used to show W[1]-hardness of geometric problems in $\mathbb{R}^d$ until now.

Basic notation. For a hyperplane h and a point set P, let h_P^+ denote the set of points of P that lie strictly on the positive side of h, and analogously h_P^-. For a point p, by $(p)_i$ we denote the i–th coordinate of p. Finally, for a number x as usual let

$$\operatorname{sign}(x) := \begin{cases} 1 & x \geq 0 \\ -1 & x < 0. \end{cases}$$

2 Affine Containment

We start with a problem for which we think the hardness proof is the most straightforward. This proof will subsequently be modified to show the main theorems.

Definition 6. *(d-Affine-Containment) Given a set of points P in $\mathbb{R}^d$, is the origin contained in the affine hull of any d points?*

Recall that $x \in \operatorname{affHull}(\{p_1, \ldots, p_j\})$ iff there exist $\alpha_i, 1 \leq i \leq j$ such that $\sum \alpha_i = 1$ and $\sum \alpha_i p_i = x$.

For a given set $S = \{s_1, \ldots, s_n\}$, we will create a point-set in $\mathbb{R}^{d+1}$ in which $d+1$ points span an affine plane through the origin if and only d of these numbers sum up to 0.

Let e_i denote the i–th unit vector. Set

$$p_i^j := \frac{1}{s_i} \cdot e_j + e_{d+1} = \left(0, \ldots, \frac{1}{s_i}, \ldots, 0, \ldots, 1\right)^T$$

and $q := -\sum_{i=1}^{d} e_i$.

The set P consists of all points $p_i^j, 1 \leq j \leq d, 1 \leq i \leq n$ and the point q. The size of the point set is thus $n \cdot d + 1$.

Lemma 1. *There are d elements that sum up to 0 iff there are $d+1$ points in P whose affine hull contains the origin*[6].

Proof. $\Rightarrow$: Let $\sum_{j=1}^{d} s_{i_j} = 0$. We choose points $x_j = p_{i_j}^j$, $1 \leq j \leq d$ and $x_{d+1} = q$. Let $\alpha_j = s_{i_j}$ and $\alpha_{d+1} = 1$. Then

$$\sum_{j=1}^{d+1} \alpha_j x_j = \sum_{j=1}^{d} s_{i_j} p_{i_j}^j + q = \sum_{j=1}^{d} e_j + \left(\sum_{j=1}^{d} s_{i_j} \right) e_{d+1} - \sum_{j=1}^{d} e_j = \mathbf{0}$$

and

$$\sum_{j=1}^{d+1} \alpha_j = \sum_{j=1}^{d} s_{i_j} + \alpha_{d+1} = 1.$$

That means that $\mathbf{0}$ is in $\operatorname{affHull}\left(\left\{p_{i_1}^1, \ldots p_{i_d}^d, q\right\}\right)$.

$\Leftarrow$: Let $\mathbf{0} \in \operatorname{affHull}(\{x_1, \ldots, x_d\})$, i.e., let $\sum_{j=1}^{d+1} \alpha_j x_j = \mathbf{0}$ and $\sum \alpha_j = 1$. As all points but q lie on the hyperplane $x_{d+1} = 1$, one of the points, without loss of generality x_{d+1}, must be q. Because of $(q)_{d+1} = 0$, and $(x)_{d+1} = 1$ for all $x \neq q$, by computing the $(d+1)$-st coordinate we get

$$0 = \sum_{j=1}^{d} (\alpha_j x_j)_{d+1} = \sum_{j=1}^{d} \alpha_j (x_j)_{d+1} = \sum_{j=1}^{d} \alpha_j \tag{1}$$

and thus $\alpha_{d+1} = 1 - \sum_{j=1}^{d} \alpha_j = 1$.

Further, as $\sum_{j=1}^{d+1} \alpha_j x_j = \mathbf{0}$, the other points must satisfy

$$\sum_{j=1}^{d} \alpha_j x_j = -\alpha_{d+1} q = \sum_{j=1}^{d} e_j.$$

Any x_j is nonzero for only one other coordinate except the $(d+1)$-st, and as $(q)_j = -1$ for all $j < d+1$, for each j there must be at least one point that is nonzero at coordinate j (in particular, also $\alpha_j \neq 0$). Thus, there are exactly d such points. Without loss of generality assume that x_j is the point that is nonzero in coordinate j, so $(x_j)_j = \frac{1}{s_{i_j}}$ for some i_j. This means that $\alpha_j \frac{1}{s_{i_j}} - 1 = 0$, and thus $\alpha_j = s_{i_j} \in S$, which implies (Eqn. 1) that we have d elements in S summing up to 0. $\square$

Theorem 3. *d-Affine-Containment is W[1]-hard with respect to the dimension and NP-hard.*

[6] Recall that the dimension is also $d+1$.

3 Caratheodory sets

In order to use the previous construction to prove the first part of Theorem 1, we have to modify it such that all coefficients can be chosen positive. Observe that $\mathbf{0} \in \mathrm{conv}(P)$ iff $\mathbf{0} = \sum_{p \in P} \alpha_p p$ for any $\alpha_p \geq 0, \sum \alpha_p > 0$ (proof: divide by $\sum \alpha_p$). To this end we now define

$$p_i^j = \frac{1}{|s_i|} \cdot e_j + \mathrm{sign}(s_i) \cdot e_{d+1}$$

and q as above. The set P again consists of all the points $p_i^j, 1 \leq j \leq d, 1 \leq i \leq n$ and q.

Lemma 2. *There are d elements in S that sum up to 0 iff the origin lies in the convex hull of $d+1$ points of P.*

Proof. $\Rightarrow$: Let $\sum_{j=1}^{d} s_{i_j} = 0$. Setting $\alpha_j = |s_{i_j}| > 0, x_j = p_{i_j}^j$ for $1 \leq j \leq d$ and $\alpha_{d+1} = 1$, $x_{d+1} = q$ again yields

$$\sum_{j=1}^{d+1} \alpha_j x_j = \sum_{j=1}^{d} |s_{i_j}| p_{i_j}^j + q = \sum_{j=1}^{d} e_j + \left(\sum_{j=1}^{d} \mathrm{sign}(s_{i_j}) |s_{i_j}| \right) e_{d+1} - \sum_{j=1}^{d} e_j = \mathbf{0}.$$

$\Leftarrow$: Let $\sum_{j=1}^{d+1} \alpha_j x_j = \mathbf{0}, \alpha_j \geq 0$. As all points lie in the positive halfspace $\sum^d e_j^* x > 0$, q must be one of the points of the convex combination. We can assume $x_{d+1} = q$ and $\alpha_{d+1} = 1$. Further, by the same argument as in Lemma 1, there must be at least d other points for the total sum to become $\mathbf{0}$. Again, without loss of generality let $(x_j)_j \neq 0$. As $(q)_j = -1$ for all $1 \leq j \leq d$, this means that $\alpha_j \frac{1}{|s_{i_j}|} = 1$ for some i_j, and thus $\alpha_j = |s_{i_j}|$. Further, because of the $(d+1)$-st coordinate, we get

$$0 = \sum_{j=1}^{d} \alpha_j \mathrm{sign}(s_{i_j}) = \sum_{j=1}^{d} \mathrm{sign}(s_{i_j}) \cdot \left| s_{i_j} \right| = \sum_{j=1}^{d} s_{i_j}$$

and thus we have d elements summing up to 0. $\square$

Thereby we have shown the first part of Theorem 1.

Remark. Observe that if we project all points onto the unit sphere, all the above properties still hold: Clearly, $\mathbf{0} \in \mathrm{conv}(P)$ iff $\mathbf{0} \in \mathrm{conv}\left(\pi_{S^{d-1}}(P)\right)$. Thus, we can even assume all points to lie in convex position and thereby get a slightly stronger result:

Theorem 4. *The following problem is W[1]-hard and NP-hard: Given a V-polytope in $\mathbb{R}^d$, is the origin contained in the convex hull of any d vertices?*

4 Helly sets

Starting from the result in the previous section, we will now show how to prove the hardness for the d-Helly-Set problem. Using a duality transform, for a given set P in $\mathbb{R}^d$, we will construct a set of convex sets (that are actually halfspaces) such that d have an empty intersection if and only if there are d points in P that contain the origin in their convex hull. A similar construction (which is used to prove Caratheodory's Theorem from Helly's Theorem) can be found in Eggleston [11, Chapter 2.3].

Consider a set P of points $p_1, \cdots, p_n \in \mathbb{R}^d$ whose convex hull contains the origin. For each point $p \in P$ set consider the halfspace

$$p^* = \left\{ x \mid p^T x \geq 1 \right\}.$$

Define P^* to be the set of all these halfspaces corresponding to the points in P. We show that any Caratheodory set of P (for the origin) corresponds to a Helly set (a set of halfspaces with empty intersection) of P^* of the same size. Since checking if the minimum Caratheodory set has cardinality at most d is W[1]-hard, it then follows that checking if the minimum Helly set is of cardinality at most d is also $W[1]$-hard.

Let $Q \subseteq P$ and let V be a $d \times |Q|$ matrix whose columns represent the vectors in Q. Further, let cone(V) denote the conic hull of the vectors, i.e., the set $\left\{ \sum_{q \in Q} \alpha_q q \mid \alpha_q \geq 0 \right\}$.

Using the fact that cone(V) is pointed if and only if $V^T x \leq \mathbf{0}$ is a full-dimensional cone, we can now show the main lemma of this section, which is a variant of Gordan's Theorem, see e.g. Dantzig and Thapa [8, Theorem. 2.13]:

Lemma 3. *Let $Q \subseteq P$ and let V be a $d \times |Q|$ matrix whose columns represent the vectors in Q. Then $\mathbf{0} \in conv(V)$ if and only if the system of inequalities $V^T x \geq \mathbf{1}$ is infeasible.*

Proof. $\Rightarrow$: Suppose that $V^T x \geq \mathbf{1}$ is feasible. Then there exists a vector $\alpha \in \mathbb{R}^d$ such that $V^T \alpha \leq -\mathbf{1}$. That is, $V^T \alpha < \mathbf{0}$ and thus $V^T x \leq \mathbf{0}$ is a full-dimensional cone. Therefore, cone(V) is pointed. But this means that $\mathbf{0} \notin$ conv(V).

$\Leftarrow$: Now suppose $\mathbf{0} \notin$ conv(V), then cone(V) is pointed and therefore $V^T x \leq \mathbf{0}$ is a full-dimensional cone. Thus, there exists $\alpha \in \mathbb{R}^d$ such that $V^T \alpha < \mathbf{0}$, and so for a large enough $\lambda > 0$, $V^T(-\lambda\alpha) > \mathbf{1}$ and hence $V^T x \geq \mathbf{1}$ is feasible. □

Thus, any set $Q \subseteq P$ of points whose convex hull contains the origin corresponds to a set $Q^* \subseteq P^*$ of convex set (inequalities) of the same size that has an empty intersection, and vice versa. This finishes the proof of the second part of Thm. 1.

As the convex sets in this case are even halfspaces, we can derive the stronger result of Corollary 1.

5 Ham-Sandwich

Using the construction from Sec. 2, we will now prove that the decision version for the HAM-SANDWICH is W[1]-hard.

A hyperplane h is said to *bisect* a set Q if $|h_Q^+| \leq \left\lfloor \frac{|Q|}{2} \right\rfloor$ and $|h_Q^-| \leq \left\lfloor \frac{|Q|}{2} \right\rfloor$. A *Ham-Sandwich cut* of d point sets $P_1, \ldots, P_d$ in $\mathbb{R}^d$ is a hyperplane h that bisects each of the sets. In particular, if the number of points in each set is odd, the hyperplane has to pass through at least one of the points from each set.

Def. 4 asks whether there is a cut that goes through a given point q. Via translation we can obviously assume q to be the origin. This will be called a *linear Ham-Sandwich cut.*

In order to show Theorem 2 we will create $d+1$ sets $P_1, \ldots, P_{d+1}$. The set P_{d+1} will consist of the single point $q = \sum_{j=1}^{d} e_j$ (which is $-q$ in the above notion). The sets P_j will be the union of the two set R_j and B_j. R_j contains all points of the form p_i^j, defined exactly as in Sec. 2, i.e.,

$$R_j := \left\{ p_i^j \mid 1 \leq i \leq n \right\}.$$

for $p_i^j = \frac{1}{s_i} e_j + e_{d+1}$. If we choose a linear hyperplane through one of these points, the number of points on each side will (most likely) not be the same. So in addition to these, for each of these sets we need $n-1$ *balancing* points B_j to ensure that any linear hyperplane passing through any of these points has equally many points of P_j on both sides (c.f. Figure 1). Thus, the set $P = \bigcup P_j$ is of size $d\,(2n-1)+1$.

Construction of the Balancing-set The idea is to add a point set similar to the mirror image of the original set R_j. This way any hyperplane that has *many* of the original points on, say, the positive side, will contain *few* of the mirrored points on the positive side, and vice versa.

By making the total number of points in each set P_j odd, we will ensure that any Ham-Sandwich cut must pass through one of the points from P_j. Further, by the construction of the balancing set, it will not be possible to choose a linear cut through q that also goes through any of these balancing points, thereby getting the correspondence between subsets of S and linear cuts through q.

For this, we will choose the mirror-image of a set of $n-1$ points that lie *between two successive points* in R_j (recall that all points from R_j lie on a line; this is why we use the construction from Sec. 2). Thereto, let S be in ascending order with respect to $s_i \prec s_j$ iff $1/s_i < 1/s_j$ (or, equivalently: $1/s_i < 1/s_j$ for $i < j$).

Then, let $\varepsilon_j = \frac{1}{2^j}$ and

$$b_i^j := -\left(\frac{1}{s_i - \varepsilon_j} \right) \cdot e_j - e_{d+1}.$$

This the mirror image of a point slightly to the right of p_i^j, for $1 \leq i < n$; see Figure 1. Let B_j consist of all balancing points of the form b_i^j and set

$$P_j := R_j \cup B_j.$$

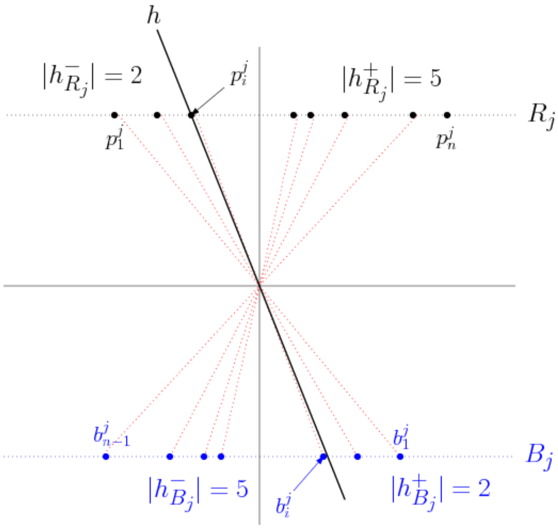

Fig. 1. The set P_j: points and balancing points

The main lemma. Now we come to prove the main lemma, namely that the point set allows a linear ham-sandwich cut if and only if there are d elements that sum up to 0, based on the following two simple lemmas. The first one states that any (not necessarily linear) ham-sandwich cut must intersect exactly one point from each set P_j, whereas the second one guarantees that any linear hyperplane that contains a point from R_j will bisect P_j.

Lemma 4. *Any linear ham-sandwich cut intersects exactly one point from each P_j, $1 \leq j \leq d+1$.*

Proof. For $P_{d+1} = \{q\}$ this is trivial. We show that for any linear Ham-Sandwich cut $h = (h_1, \ldots, h_{d+1})$ we have $h_i \neq 0$ for all i: First, if h_{d+1} were 0, because the cut must pass through at least one point from each set, we would have $h_j = 0$ for all j. Thus, $h_{d+1} \neq 0$. Further, as $h_j(p^j)_j = -h_{d+1}(p^j)_j \neq 0$ for some $p^j \in P_j$, we also must have $h_j \neq 0$ for all j.

Thus, no cut can pass through more than one point of any set P_j: If

$$h_j(p)_j + h_{d+1}(p)_{d+1} = h \cdot p = 0 = h \cdot p' = h_j(p')_j + h_{d+1}(p')_{d+1}$$

for two points $p, p' \in P_j$, then $p = p'$ or $h_j = 0$, a contradiction. □

Lemma 5. *Any linear hyperplane intersecting a single point from R_j bisects the set P_j.*

Proof. Let $h \cdot p_i^j = 0$ and without loss of generality $h \cdot p_k^j < 0$ for all $1 \leq k < i$. Then also $h \cdot -b_k^j < 0$ and thus $h \cdot b_k^j > 0$ for all $1 \leq k < i$. Further, $h \cdot p_k^j > 0$ for all $k > i$ and $h \cdot b_k^j < 0$ for $k \geq i$. So

$$|h^-_{P_j}| = |h^-_{R_j}| + |h^-_{B_j}| = i - 1 + n - i = \left\lfloor \frac{|P_j|}{2} \right\rfloor = |h^+_{P_j}|.$$

□

Lemma 6. *There are d elements in S that sum up to 0 if and only if there is a linear Ham-Sandwich cut.*

Proof. $\Rightarrow$: Let $\sum_{j=1}^{d} s_{i_j} = 0$. We have to find a linear hyperplane $h \cdot x = 0$ such that for each set P_j it holds that $|h^+_{P_j}|, |h^-_{P_j}| \leq \left\lfloor \frac{|P_j|}{2} \right\rfloor$. Choose $h_j = s_{i_j}$ for $1 \leq j \leq d$ and $h_{d+1} = -1$. Because $\sum^d s_{i_j} = 0$, we have $h \cdot q = \sum^d s_{i_j} = 0$ (so the one element set P_{d+1} is bisected). Further,

$$h \cdot p^j_{i_j} = h_j \cdot 1/s_{i_j} + h_{d+1} \cdot 1 = 1 - 1 = 0.$$

Because of Lemma 5, this means that all sets are bisected, and thus we have a linear Ham-Sandwich cut.

$\Leftarrow$: Let h be a linear Ham-Sandwich cut. All h_i must be nonzero (Lemma 4), so we can assume $h_{d+1} = -1$. For each j, we have $h \cdot p^j = 0$ for exactly one point $p^j \in P_j$. This means that

$$0 = h \cdot p^j = h_j(p^j)_j + h_{d+1}(p^j)_{d+1} = h_j(p^j)_j - 1(p^j)_{d+1} = h_j(p^j)_j - 1,$$

and so either $h_j = s_{i_j}$ or $h_j = s_{i_j} - \varepsilon_j$ for some i_j. Because for any $\emptyset \neq J \subset \{1, \ldots, d\}$ we have $0 < \sum_{j \in J} \varepsilon_j < 1$ and S is a set of integers, if one (or more) of the h_j were of the latter form, the total sum can never be an integer, and in particular not 0. But this is required for q to lie on h.

Thus, $h_j = s_{i_j} \in S$ for some i_j, and as q also lies on the hyperplane, we get

$$0 = hq = \sum_{j=1}^{d} h_j = \sum_{j=1}^{d} s_{i_j},$$

i.e., there are d elements in S that sum up to 0. □

From this Theorem 2 follows.

Remarks. In the previous construction, the origin (i.e., the point for which we want to solve the decision version) is not part of any of the sets, but this is easily fixed: Set $P_{d+1} = \{\mathbf{0}, q/2, q\}$. Then any Ham-Sandwich cut through $\mathbf{0}$ also has to go through the other two points (otherwise there would be too many points on the one side) and thus must also contain q. On the other hand, whenever there are no such d elements that sum up to 0, all Ham-Sandwich cuts are (truly) affine hyperplanes through $q/2$. This gives a slightly stronger result:

Theorem 5. *The following problem is W[1]-hard with respect to the dimension and NP-hard: Given d point sets in $\mathbb{R}^d$ and a point $q \in \bigcup P_i$, is there a Ham-Sandwich cut through q?*

The hardness proof can easily be adapted to show the following:

Corollary 4. *The following problems are W[1]-hard with respect to the dimension and NP-hard:*

1. *(d-Strong-Ham-Sandwich) Given $d+1$ point sets in $\mathbb{R}^d$, is there a hyperplane that bisects all sets?*
2. *(d-Colorful-Hyperplane) Given d point sets in $\mathbb{R}^d$, is there a linear hyperplane that contains at least one point from each set?*

Proof. 1. Follows by adding $\mathbf{0}$ as a set.
2. Follows from the fact that any Ham-Sandwich cut through the origin will contain a point from each of the sets.

□

References

1. P. K. Agarwal, M. Sharir, and E. Welzl. Algorithms for center and Tverberg points. In *Journal of the ACM*, volume 5, pages 1–20, New York, NY, USA, 2008. ACM.
2. E. Amaldi, M. E. Pfetsch, and L. E. T. Jr. On the maximum feasible subsystem problem, iiss and iis-hypergraphs. *Math. Program.*, 95(3):533–554, 2003.
3. S. Cabello, P. Giannopoulos, and C. Knauer. On the parameterized complexity of d-dimensional point set pattern matching. In *Proc. 2nd Int. Workshop on Parameterized and Exact Computation (IWPEC)*, volume 4169 of *LNCS*, pages 175–183, 2006.

4. S. Cabello, P. Giannopoulos, C. Knauer, D. Marx, and G. Rote. Geometric clustering: fixed-parameter tractability and lower bounds with respect to the dimension. *ACM Transactions on Algorithms*, 2009. to appear.
5. S. Cabello, P. Giannopoulos, C. Knauer, and G. Rote. Geometric clustering: fixed-parameter tractability and lower bounds with respect to the dimension. In *Proc. 19th Ann. ACM-SIAM Sympos. Discrete Algorithms*, pages 836–843, 2008.
6. C. Caratheodory. ber den Variabilitätsbereich der Fourierschen Konstanten von positiven harmonischen Funktionen. *Rend. Circ. Mat. Palermo*, 32:193–217, 1911.
7. T. M. Chan. A (slightly) faster algorithm for Klee's measure problem. In *Proc. 24th Annual Symposium on Computational Geometry*, pages 94–100, 2008.
8. G. B. Dantzig and M. N. Thapa. *Linear Programming 2: Theory and Extensions.* Springer, 2003.
9. R. G. Downey and M. R. Fellows. *Parameterized Complexity.* Monographs in Computer Science. Springer, November 1999.
10. H. Edelsbrunner and R. Waupotitsch. Computing a Ham-sandwich Cut in Two Dimensions. *J. Symb. Comput.*, 2(2):171–178, 1986.
11. H. G. Eggleston. *Convexity.* Cambridge University Press, 1963.
12. J. Erickson. New Lower Bounds for Convex Hull Problems in Odd Dimensions. *SIAM J. Comput.*, 28(4):1198–1214, 1999.
13. M. R. Fellows and N. Koblitz. Fixed-Parameter Complexity and Cryptography. *AAECC-10: Proceedings of the 10th International Symposium on Applied Algebra, Algebraic Algorithms and Error-Correcting Codes*, pages 121–131, 1993.
14. J. Flum and M. Grohe. *Parameterized Complexity Theory.* Texts in Theoretical Computer Science. Springer, 2006.
15. A. Gajentaan and M. H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom. Theory Appl.*, 5(3):165–185, 1995.
16. P. Giannopoulos, C. Knauer, and G. Rote. The parameterized complexity of some geometric problems in unbounded dimension. In *Parameterized and exact computation. 4th international workshop, IWPEC 2009*, volume 5917 of *Lecture Notes in Computer Science*, pages 198–209, 2009.
17. P. Giannopoulos, C. Knauer, M. Wahlström, and D. Werner. Hardness of discrepancy and related problems parameterized by the dimension. In *26th European Workshop on Computational Geometry (EuroCG)*, pages 173–176, 2010.
18. P. Giannopoulos, C. Knauer, and S. Whitesides. Parameterized Complexity of Geometric Problems. *Computer Journal*, 51(3):372–384, 2008.
19. E. Helly. Über Mengen konvexer Korper mit Gemeinschaftlichen Punkten. *Jahresbericht Deutsch. Math. Verein.*, 32:175–176, 1923.
20. R. Impagliazzo and R. Paturi. On the Complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
21. S. Jadhav and A. Mukhopadhyay. Computing a centerpoint of a finite planar set of points in linear time. In *SCG '93: Proceedings of the ninth annual symposium on Computational geometry*, pages 83–90. ACM, 1993.
22. S. Langerman and P. Morin. Covering Things with Things. *Discrete and Computational Geometry*, 33(4):717–729, 2005.
23. J. Matoušek. *Lectures on Discrete Geometry.* Springer-Verlag New York, Inc., 2002.
24. J. Matoušek. *Using the Borsuk-Ulam Theorem - Lectures on Topological Methods in Combinatorics and Geometry.* Springer, 2003.
25. J. Matoušek, C.-Y. Lo, and W. Steiger. Ham-sandwich cuts in R^d. *STOC '92: Proceedings of the twenty-fourth annual ACM Symposium on Theory of Computing*, pages 539–545, 1992.

26. G. L. Miller and D. R. Sheehy. Approximate center points with proofs. In *SCG '09: Proceedings of the 25th annual symposium on Computational geometry*, pages 153–158, New York, NY, USA, 2009. ACM.
27. R. Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford Lecture Series in Mathematics and its Applications. Oxford University Press, 2006.
28. C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48(3):498–532, 1994.
29. M. Pătraşcu and R. Williams. On the Possibility of Faster SAT Algorithms. *Proc. 21st ACM/SIAM Symposium on Discrete Algorithms (SODA)*, 2010.
30. D.-H. Teng. *Points, spheres, and separators: A unified geometric approach to graph partitioning*. PhD thesis, School of Computer Science, Carnegie Mellon University, 1992.
31. F. Yu. On the Complexity of the Pancake Problem. *Electron. Notes Theor. Comput. Sci.*, 167:95–115, 2007.
32. G. M. Ziegler. *Lectures on Polytopes*. Springer, 1998.