

On a Geometric Generalization of the Upper Bound Theorem

Uli Wagner*

Einstein Institute of Mathematics
The Hebrew University of Jerusalem
Jerusalem, 91904, Israel
uli@math.huji.ac.il

Institut für Theoretische Informatik
ETH Zürich
8092 Zürich, Switzerland
uli@inf.ethz.ch

Abstract

We prove an upper bound, tight up to a factor of 2, for the number of vertices of level at most ℓ in an arrangement of n halfspaces in \mathbf{R}^d , for arbitrary n and d (in particular, the dimension d is not considered constant). This partially settles a conjecture of Eckhoff, Linhart, and Welzl.

Up to the factor of 2, the result generalizes McMullen's Upper Bound Theorem for convex polytopes (the case $\ell = 0$) and extends a theorem of Linhart for the case $d \leq 4$. Moreover, the bound sharpens asymptotic estimates obtained by Clarkson and Shor.

The proof is based on the h -matrix of the arrangement (a generalization, introduced by Mulmuley, of the h -vector of a convex polytope). We show that bounding appropriate sums of entries of this matrix reduces to a lemma about quadrupels of sets with certain intersection properties, and we prove this lemma, up to a factor of 2, using tools from multilinear algebra. This extends an approach of Alon and Kalai, who used linear algebra methods for an alternative proof of the classical Upper Bound Theorem. The bounds for the entries of the h -matrix also imply bounds for the number of i -dimensional faces, $i > 0$, at level at most ℓ .

Furthermore, we discuss a connection with crossing numbers of graphs that was one of the main motivations for investigating exact bounds that are valid for arbitrary dimensions.

1. Introduction

Levels in arrangements are fundamental objects in computational geometry. Let \mathcal{A} be a set of n closed affine halfspaces in \mathbf{R}^d , which we can think of as the constraints of a linear program. Then the level of a point $x \in \mathbf{R}^d$ is defined as the number of constraints that x violates, i.e., the

number of halfspaces that it is *not* contained in. The hyperplanes bounding the halfspaces define a decomposition of \mathbf{R}^d into convex polyhedral cells, or *faces*, of dimensions $i = 0, 1, \dots, d$. If we want to stress this decomposition, we also speak of the *arrangement* \mathcal{A} . Two points that lie in the same face of the arrangement have the same level, so it makes sense to speak of the level of a face. We will assume throughout that the arrangement is *simple*, i.e., that any d of the bounding hyperplanes intersect in exactly one point and that no $d + 1$ of them have a point in common.

The definitions carry over verbatim to arrangements of great hemispheres in the d -dimensional sphere \mathbf{S}^d . In such a spherical arrangement, the faces come in antipodal pairs, and if a face F has level ℓ and lies at the intersection of i bounding great $(d - 1)$ -spheres (in simple arrangements, these are precisely the $(d - i)$ -dimensional faces), then the antipodal face $-F$ has level $n - i - \ell$. By homogenization (i.e., by embedding \mathbf{R}^d as the affine hyperplane $\{x_{d+1} = 1\}$ into \mathbf{R}^{d+1}), an affine arrangement of n halfspaces in \mathbf{R}^d corresponds bijectively to an arrangement of n hemispheres in \mathbf{S}^d together with an additional distinguished “northern hemisphere” (which is not considered to be part of the arrangement). The spherical viewpoint is more symmetric and elegant, but for certain key notions that are available in the affine case, such as the h -matrix induced by a generic linear functional (see Section 2), suitable spherical analogues are still missing.

Levels in arrangements and their complexity have been studied extensively by discrete and computational geometers (often in the polar dual setting of k -sets of finite point sets, and also for arrangements of more general hypersurfaces). Of particular interest is the question of the maximum complexity of a single level: For an arrangement \mathcal{A} , let $v_\ell(\mathcal{A})$ denote the number of vertices (i.e., 0-dimensional faces) of the arrangement at level ℓ . With this notation, the precise question is this: What is the maximum of $v_\ell(\mathcal{A})$ over all arrangements \mathcal{A} of n affine halfspaces in \mathbf{R}^d (respectively, of hemispheres in \mathbf{S}^d). Notice that the sum $\sum_{\ell=0}^{n-d} v_\ell(\mathcal{A})$ equals $\binom{n}{d}$ for affine arrangements, respectively twice that number for spherical arrangements, and

*Research supported by the research training networks *Algebraic Combinatorics in Europe (ACE)* and *Combinatorial Structure of Intractable Problems (COMBSTRU)* and by the *Edmund Landau Minerva Center for Research in Mathematical Analysis and Related Areas*.

that in the latter case, $v_\ell = v_{n-d-\ell}$ for all ℓ . Consequently, $\binom{n}{d}$ is a trivial upper bound for each v_ℓ . If we allow the intersection of the halfspaces or hemispheres to be empty, then this upper bound is tight: It is not hard to show that if we apply *Gale duality* (see, e.g., [29], Chapter 5, for a definition) to a polar-to-neighborly arrangement (see below), then the resulting dual arrangement will have all vertices at the middle level. The maximum complexity of a single level becomes much more challenging to determine if we require that $\cap \mathcal{A} \neq \emptyset$; borrowing terminology from linear programming, we call this the *feasible case*. In this case the arrangement is the polar dual of a set P of n points in affine d -space (in the infeasible case, we can think of the dual as a “signed point set”), and vertices of level ℓ correspond to so-called ℓ -facets of the polar point set. In the feasible case, the trivial upper bound $\binom{n}{d}$ for the maximum number of vertices at a single level is no longer tight, and not even the asymptotics of this maximum is known, not even for fixed dimension. The currently best bounds are as follows: the maximum complexity of the ℓ -level in the plane is at most $O(n\sqrt[3]{\ell+1})$ [13] and at least $n \cdot e^{\Omega(\sqrt{\ell+1})}$ [40]. In higher dimensions, the best upper bounds are $O(n(\ell+1)^{3/2})$ in dimension 3 [36], $O(n^2(\ell+1)^{2-2/45})$ in dimension 4 [30], and $O(n^{\lfloor d/2 \rfloor}(\ell+1)^{\lfloor d/2 \rfloor - \varepsilon_d})$ in general dimension d , where $\varepsilon_d = 1/(4d-3)^d$ [46, 1].

($\leq \ell$)-Levels. The primary subject of this article is the complexity of the ($\leq \ell$)-level of an arrangement, i.e., the quantity $v_{\leq \ell} := \sum_{k=0}^{\ell} v_k$. This quantity is much better understood than the complexity of a single level, and for many applications, bounds for the ($\leq \ell$)-level are sufficient. The ($\leq \ell$)-level also arises naturally when studying a “relaxed version” of linear programming, where a limited number of constraints are allowed to be violated [11]. Clarkson and Shor [12] showed that for every arrangement of n hemispheres in d dimensions, and for $0 \leq \ell \leq n-d$,

$$v_{\leq \ell}(\mathcal{A}) \leq 2 \left(\frac{e}{\lfloor d/2 \rfloor} \right)^{\lfloor d/2 \rfloor} \binom{n}{\lfloor d/2 \rfloor} (\ell + \lfloor d/2 \rfloor)^{\lfloor d/2 \rfloor}.$$

For fixed dimension d , and $n \rightarrow \infty$, this bound is asymptotically tight, but it assumes its full strength only if n is large compared to d . As one motivation for the study of exact bounds that are valid for arbitrary n and d , we will discuss, at the end of this introduction, an application of ($\leq \ell$)-levels to crossing numbers of complete graphs for which we need a bound in the case $n = d+4$; for this choice of parameters, for instance, the Clarkson-Shor estimate is larger than the trivial bound.

The case $\ell = 0$ is the only one for which we know exact bounds that hold for all values of n and d . This is the content of McMullen’s [31] Upper Bound Theorem for convex polytopes:

Theorem 1.1. *For every arrangement \mathcal{A} of n hemispheres in d dimensions,*

$$v_0(\mathcal{A}) \leq v_0(\mathcal{C}_{n,d}^*),$$

where $\mathcal{C}_{n,d}^*$ is a polar-to-cyclic arrangement of hemispheres; i.e., the hemispheres are of the form $\{x \in \mathbf{S}^d : x_0 + t_1 x_1 + t_2 x_2 + \dots + t_d x_d \leq 0\}$ for n distinct real parameters $t_1, \dots, t_n \in \mathbf{R}$. Moreover, equality is attained iff \mathcal{A} is polar-to-neighborly, i.e., iff the intersection of any $\lfloor d/2 \rfloor$ bounding great $(d-1)$ -spheres contains a vertex at level 0.

We remark that the notion of a polar-to-cyclic arrangement without further specifications is only well-defined for spherical arrangements; in the affine case, the arrangement depends on the choice of the additional “northern hemisphere”. For the 0-level, however, the exact choice is immaterial, as long as the “equator”, i.e., the boundary of the northern hemisphere, does not intersect the 0-level.

The Upper Bound Theorem has been generalized in numerous ways. Most of these generalizations take the polar dual statement as their starting point: The number of facets of a simplicial d -polytope with n vertices is maximized by the cyclic polytope $\mathcal{C}_{n,d}$. The same upper bound for the number of facets holds for other $(d-1)$ -dimensional simplicial complexes on n vertices: For simplicial $(d-1)$ -spheres [38], for Eulerian complexes (provided n is sufficiently large) [23], and for several classes of simplicial manifolds and pseudomanifolds [33, 20, 34]. Another far-reaching extension is Kalai’s [22] so-called *Strong Upper Bound Theorem* concerning subcomplexes of the boundary complex of a simplicial polytope.

Most of these extensions have a topological or algebraic flavor. Here, we consider the following, more geometric, generalization concerning ($\leq \ell$)-levels:

Conjecture 1.2. *Let \mathcal{A} be an arrangement of n great hemispheres in \mathbf{S}^d . Then*

$$v_{\leq \ell}(\mathcal{A}) \leq v_{\leq \ell}(\mathcal{C}_{n,d}^*)$$

for $0 \leq \ell \leq (n-d-1)/2$. Equality holds iff \mathcal{A} is polar-to-neighborly.

This conjecture, which we refer to as *Spherical Generalized Upper Bound Conjecture (SGUBC)* was proposed by Eckhoff [14], Linhart [25], and Welzl [44], independently of one another.

We remark that there are many polar-to-neighborly simple arrangements that are not polar-to-cyclic (see [37] for a construction of superexponentially many). However, for the complexity of the ℓ -level, these differences do not matter: If \mathcal{N} is a polar-to-neighborly simple arrangement of hemispheres in \mathbf{S}^d , then $v_\ell(\mathcal{N}) = v_\ell(\mathcal{C}_{n,d}^*)$ for $0 \leq \ell \leq n-d$. This follows from the fact that polar-to-neighborliness is preserved under deletions, together with

a standard random sampling argument. The numbers of i -dimensional faces at level ℓ in a polar-to-cyclic spherical arrangement were computed explicitly by Andrzejak and Welzl [4]. In particular,

$$v_\ell(\mathcal{C}_{n,d}^*) = \begin{cases} 2^{\binom{\ell+\lceil d/2 \rceil-1}{\lceil d/2 \rceil-1}} \binom{n-\ell-\lceil d/2 \rceil}{\lceil d/2 \rceil-1} & \text{for odd } d, \\ \binom{\ell+d/2-1}{d/2-1} \binom{n-\ell-d/2}{d/2} + \binom{\ell+d/2}{d/2} \binom{n-\ell-d/2-1}{d/2-1} & \text{for even } d. \end{cases}$$

The SGUBC is known to be true in dimension $d = 2$ [2, 35]; and in dimension 3 provided the intersection of the hemispheres is nonempty [44]. The explicit upper bounds are $n(\ell + 1)$ and $\leq 2 \binom{\ell+2}{2} n - 2 \binom{\ell+3}{3}$ for dimensions 2 and 3, respectively.

As observed by Welzl, for arrangements of hemispheres, the restriction $\ell \leq (n - d - 1)/2$ is crucial if we are striving for exact bounds: For instance, in the case $d = 2$ and $k = 0$, let v_ℓ denote the number of vertices at level ℓ . In the spherical case, $v_\ell = v_{n-2-\ell}$. Thus, if n is even, then $2v_{\leq (n-2)/2} = 2 \binom{n}{2} + v_{(n-2)/2}$. Consequently, the bound $v_{\leq \ell} \leq n(\ell + 1)$ does not hold for $\ell = (n - 2)/2$, else we would get an linear upper bound of $v_{(n-2)/2} \leq n$ for the middle level, contradicting the known superlinear lower bounds mentioned above.

The SGUBC is also related to the so-called Generalized Lower Bound Theorem (which is part of a complete combinatorial characterization of the face numbers of simplicial convex polytopes, the g -Theorem, conjectured by McMullen and proved by Stanley [39] and by Billera and Lee [7]). For instance, Welzl [44] showed that the SGUBC for arrangements in \mathbf{S}^3 is equivalent, by Gale duality, to the GLBT for d -polytopes with $d + 4$ vertices.

In the present paper, we prove the following:

Theorem 1.3. *Let \mathcal{A} be an arrangement of n affine half-spaces in \mathbf{R}^d . Then*

$$v_{\leq \ell}(\mathcal{A}) \leq 2 \cdot v_{\leq \ell}(\mathcal{C}_{n,d}^*)$$

for $0 \leq \ell \leq n - d$ (where $\mathcal{C}_{n,d}^*$ denotes a polar-to-cyclic spherical arrangement as in the SGUBC).

We refer to Theorem 1.3 as 2AGUBT. A sharp version, without the factor of 2, was conjectured by Linhart, who proved it for $d \leq 4$; we refer to this sharp version as the *Affine Generalized Upper Bound Conjecture (AGUBC)*; it is sharp for $\ell < \lceil n/(d + 1) \rceil$, because it is equivalent to the SGUBC in that range. (This follows by taking the polar dual of a cyclic polytope with respect to a center point of the vertex set). Moreover, the AGUBC, if it is true, might even be sharp for all $\ell \leq (n - d)/2$; for instance, this is the case for $d = 2, 3$. At any rate, for every ℓ , the AGUBC

implies the SGUBC up to a factor of 2, since we can always capture at least half of the vertices of level $\leq \ell$ in a spherical arrangement by a suitable “northern hemisphere”.

We remark that the remaining multiplicative gap in the 2AGUBT very much appears to be an artifact of the proof. In an earlier version of this paper I claimed, in fact, to prove the AGUBC without the factor of 2, but during the preparation of this final version it became clear that the proof of the sharp version contained an error.

McMullen’s Upper Bound Theorem gives exact upper bounds for the maximum number of i -dimensional faces of the arrangement at level 0, for all $0 \leq i \leq d$. It is natural to extend the above generalizations to faces of intermediate dimensions. For $d \leq 4$, Wesp [45] extended Linhart’s proof of the AGUBT to faces of any dimension i , $0 \leq i \leq d$. Our method also gives bounds for the number of bounded i -faces at level $\leq \ell$; by induction on the dimension, we also obtain somewhat weaker bounds for the unbounded faces, see Corollary 2.5.

The structure of this paper. In Section 2, we define and study the h -matrix of a linear program. We state bounds on suitable sums of entries of this matrix and show how these imply bounds for the numbers $v_{\leq \ell}$ (or indeed faces of any dimension). Moreover, we prove a lemma that reduces these bounds to a problem about quadrupels of finite sets with certain intersection properties. In Section 3, we state our main result, Theorem 3.1, which provides exact upper bounds, up to a factor of 2, for this extremal problem. The proof of this theorem uses tools from multilinear algebra, which are reviewed in Section 4, and the proof is given in Section 5. The final Section 6 is devoted to open problems and concluding remarks.

A connection to crossing numbers. The *crossing number* of a graph G is the minimum number $\text{cr}(G)$ of crossings in any drawing of G in the plane \mathbf{R}^2 or on the 2-sphere \mathbf{S}^2 (where the vertices are represented as points and the edges as Jordan arcs connecting the points). Determining the crossing number of a given graph is an NP-hard problem [16], and even for some very basic classes of graphs it is not known what their crossing number is (see Chapter 9 of [10] for a recent survey of crossing numbers and related problems). Determining the crossing numbers of complete graphs and of complete bipartite graphs, respectively, are among the oldest problems in the area (in fact, the notion of crossing number was invented by Turán [41] when he posed the latter question, also known as “Turán’s Brick Factory Problem”). For complete graphs, it is conjectured [18] that

$$\text{cr}(K_n) \stackrel{(?)}{=} \frac{1}{4} \lfloor \frac{n}{2} \rfloor \lfloor \frac{n-1}{2} \rfloor \lfloor \frac{n-2}{2} \rfloor \lfloor \frac{n-3}{2} \rfloor. \quad (1)$$

Let us denote conjectured optimum by $c(n)$. Observe that $c(n) \sim \frac{3}{8} \binom{n}{4}$. There are constructions that achieve $c(n)$ crossings [19], but the best lower bound is $\text{cr}(K_n) \geq (3/10 + o(1)) \binom{n}{4}$. It is natural to ask if we can prove better bounds for more restricted kinds of drawings. If we require all edges to be straight-line segments, it is known [27] that at least $(3/8 + \varepsilon + o(1)) \binom{n}{4}$ crossings are necessary, for some $\varepsilon > 0$ (the currently best bound is $\varepsilon \approx 10^{-5}$, [6]). On the other hand, suppose we draw K_n by placing n points on the sphere \mathbf{S}^2 and connecting any two by the shorter geodesic arc between them. Let us define the *spherical geodesic crossing number* $\check{\text{cr}}(K_n)$ as the minimum number of crossings that can be achieved through such a drawing. The following is another simple consequence of Gale duality:

Proposition 1.4. $\check{\text{cr}}(K_n) \leq c(n)$. Moreover, equality holds iff the SGUBC holds for arrangements of n hemispheres in \mathbf{S}^{n-4} .

2 The h -Matrix of a Linear Program

One of the central notions in McMullen’s proof of the Upper Bound Theorem is the h -vector of a simple convex polytope. The j -th entry of the h -vector counts the number of vertices of out-degree j if we orient the graph of the polytope by a generic linear functional. Following Mulmuley [32], we define in this section a generalization for higher levels in arrangements.

Orientations induced by generic linear functionals. Let \mathcal{A} be an arrangement of n affine halfspaces in \mathbf{R}^d , and let φ be a linear functional $\mathbf{R}^d \rightarrow \mathbf{R}$. The pair (\mathcal{A}, φ) is called a *linear program*. We will assume that φ is *generic*, i.e., that $\varphi(v) \neq \varphi(w)$ for any two vertices $v \neq w$ (in fact, it would be sufficient to assume that φ is not constant on any edge). The functional φ defines an orientation on the graph of the arrangement (we also consider the semiinfinite rays as edges): we orient each edge in the direction in which φ increases.

Bases. By our assumption that the arrangement is simple, each vertex is the intersection of precisely d bounding hyperplanes of halfspaces in $\mathcal{A} = \{H_1, \dots, H_n\}$. We will frequently identify the vertices with the corresponding d -element subsets of the set $[n] = \{1, \dots, n\}$ of indices, which we call *bases*. Note that every vertex is incident to precisely d pairs of edges: If v is defined by d bounding hyperplanes, then omitting any one of these hyperplanes, we obtain an intersection of $d - 1$ hyperplanes, i.e., a line, that passes through v . Of the two edges on this line that are incident to v , one is directed towards v and the other one away

from v . In this sense, every vertex of the arrangement has in-degree and out-degree equal to d .

Out-labels and out-degrees. For our purposes, the following notions of in-degree and out-degree are more suitable: As we just saw, if a vertex corresponds to a set B of d halfspaces, then each $b \in B$ corresponds to a pair of edges incident to v . Exactly one of these two edges is contained in the halfspace b ; if this distinguished edge is directed away from v , then we say that b is an *out-label* of v (or equivalently, of B); otherwise it is an *in-label*. We denote the set of out-labels by $\text{out}(B)$ and define the *out-degree* of the vertex as the cardinality $|\text{out}(B)|$.

Conflicts and levels. Let v be a vertex of \mathcal{A} and let B be the corresponding basis. We define the set of *conflicts* of v as $\text{cfl}(B) := \{c \in [n] : v \notin H_c\}$. Thus, the level of v is precisely the number of its conflicts.

Definition 2.1 (*h -Matrix*). For $0 \leq j \leq d$ and $0 \leq \ell \leq n - d$, we define $h_{j,\ell}(\mathcal{A}, \varphi)$ as the number of vertices of out-degree j and level ℓ . We will often suppress \mathcal{A} and φ from the notation and simply write $h_{j,\ell}$.

The h -vector of a convex polytope is independent of the choice of the linear objective function φ . Mulmuley [32] showed that the same is true for the numbers $h_{j,\ell}(\mathcal{A}, \varphi)$, $0 \leq j \leq d$, provided the ℓ -level of the arrangement is bounded. In general, however, the entries of the h -matrix depend on φ .

Feasibility and boundedness. We say that the linear program (\mathcal{A}, φ) is *feasible* if the intersection $\cap \mathcal{A}$ of all halfspaces is nonempty. Of course, this is independent of the objective function φ , and we will also sometimes speak of feasibility or infeasibility of an arrangement.

We say the linear program (\mathcal{A}, φ) is *bounded* if there exists some vertex of out-degree 0 (but we do not require that that vertex be of level 0).

If the program is both feasible and bounded, then there exists a unique vertex v of both in-degree and level 0. We call this vertex (and sometimes the corresponding basis) the *optimum* of the LP, denoted by $\text{opt}(\mathcal{A}, \varphi)$.

LP-duality. We will use linear programming duality as a blackbox (see, for instance, [17]). If (\mathcal{A}, φ) is a linear program with n constraints in dimension d , then its *dual* $(\mathcal{A}, \varphi)^*$ is a linear program with n constraints in dimension $n - d$. The halfspaces of the dual program are labeled by the same set $[n]$ of indices. The dual of a generic LP is again generic, and passing to complimentary index sets gives a one-to-one correspondence between the bases of the primal and bases of the dual program, $B \leftrightarrow B^* := [n] \setminus B$. Under this correspondence, we have the following properties:

Lemma 2.2. Let $(\mathcal{A}^*, \varphi^*)$ be the dual of (\mathcal{A}, φ) . (Beware that both \mathcal{A}^* and φ^* depend on both \mathcal{A} and φ .)

1. For $B \in \binom{[n]}{d}$, $\varphi(B) = -\varphi^*(B^*)$ (the minus sign is only there because we want to have both the primal and the dual to be problems of maximizing a linear function).
2. Conflicts and out-labels are dual to each other,

$$\text{cfl}(B) = \text{out}(B^*) \quad \text{and} \quad \text{out}(B) = \text{cfl}(B^*).$$

Corollary 2.3. For $0 \leq j \leq d$ and $0 \leq \ell \leq n - d$,

$$h_{j,\ell}(\mathcal{A}, \varphi) = h_{\ell,j}((\mathcal{A}, \varphi)^*).$$

Because of this duality and because of the difficulty of the determining the complexity of a single level in a fixed dimension (respectively, because the latter is too large if we allow infeasible linear programs), we cannot expect to prove exact upper bounds for a single entry of the h -matrix or for the sum of entries in a single row or column.

Theorem 2.4. For every generic d -dimensional linear program \mathcal{A} with n constraints,

$$h_{\leq j, \leq \ell} \leq 2 \cdot \sum_{i=1}^j \binom{n-d-\ell+j}{i} \binom{d-j+\ell}{d-i}.$$

Lemma 2.6 below reduces Theorem 2.4 to a problem concerning quadrupels of sets with certain cardinality and intersection properties. Up to the factor of 2, the family $\mathcal{B}(j, d, \ell, n)$ provides an extremal example for the latter.

Proof of Theorem 1.3 from Theorem 2.4. Given the arrangement \mathcal{A} , we can choose any generic φ and get

$$v_{\leq \ell}(\mathcal{A}) = h_{\leq \lfloor d/2 \rfloor, \leq \ell}(\mathcal{A}, \varphi) + h_{\leq \lfloor \frac{d-1}{2} \rfloor, \leq \ell}(\mathcal{A}, -\varphi).$$

By applying the bounds for the h -matrix, we conclude

$$\begin{aligned} \frac{1}{2} v_{\leq \ell}(\mathcal{A}) &\leq \sum_{i=0}^{\lfloor d/2 \rfloor} \binom{n - \lceil \frac{d}{2} \rceil - \ell}{i} \binom{\lceil \frac{d}{2} \rceil + \ell}{d-i} \\ &\quad + \sum_{i=0}^{\lfloor \frac{d-1}{2} \rfloor} \binom{n - \lceil \frac{d+1}{2} \rceil - \ell}{i} \binom{\lceil \frac{d+1}{2} \rceil + \ell}{d-i} \end{aligned}$$

A routine computation shows that is the desired bound (we have to compare the results of our computations with the numbers $v_{\leq \ell}(\mathcal{C}^*(n, d))$, quoted in Section 1). \square

Higher-dimensional faces. Given a linear program (\mathcal{A}, φ) , let \vec{f}_{ℓ}^r denote the number of r -dimensional faces F at level ℓ of the arrangement \mathcal{A} that are bounded in the direction of φ (i.e., that have a finite φ -maximum, which will be attained a vertex), and let \bar{f}_{ℓ}^r denote the number of r -dimensional faces that are bounded both in the direction of φ and of $-\varphi$ (by genericity, these are precisely the r -faces that are bounded in the usual sense). A double-counting argument (charging every face to its “sink”, i.e., its φ -maximal vertex) [32] yields

$$\vec{f}_{\ell}^r = \sum_{j=0}^d \sum_{s=0}^r \binom{j}{r-s} \binom{d-j}{s} h_{j,\ell-s}(\mathcal{A}, \varphi). \quad (2)$$

This can be used to derive

$$\bar{f}_{\leq \ell}^r \leq 2 \cdot f_{\leq \ell}^r(\mathcal{C}_{n,d}^*).$$

Since the unbounded faces correspond to faces of a $(d-1)$ -dimensional arrangement (in the “hyperplane at infinity”), induction on the dimension yields:

Corollary 2.5. $f_{\leq \ell}^r(\mathcal{A}) \leq 2 \cdot \sum_{i=0}^d f_{\leq \ell}^{r-i}(\mathcal{C}_{n,d-i}^*)$.

The following lemma translates the problem of finding bounds for the numbers $h_{\leq j, \leq \ell}$ into a problem in extremal set theory:

Lemma 2.6. Consider a generic d -dimensional linear program (\mathcal{A}, φ) with n constraints. Let v_1, \dots, v_t be the vertices of out-degree at most j and level at most ℓ , listed in the order of increasing value of φ . For each vertex v_r , let B_r be the corresponding basis, and let $A_r = \text{out}(v_r)$ and $C_r = \text{cfl}(v_r)$ be the sets of labels of outgoing edges and of conflicts of the vertex, respectively. Then the following properties hold:

1. For every r , $A_r \subseteq B_r$, $C_r \subseteq \bar{B}_r := [n] \setminus B_r$, $B_r \cap \bar{B}_r = \emptyset$, $|A_r| \leq j$, $|B_r| = d$, $|C_r| \leq \ell$, and $|\bar{B}_r| = n - d$.

2. For $1 \leq r < s \leq t$,

$$A_r \cap (\bar{B}_s \setminus C_s) \neq \emptyset \quad \text{or} \quad (B_r \setminus A_r) \cap C_s \neq \emptyset.$$

In Theorem 3.1 below, we will prove that if $(A_r, B_r, C_r, \bar{B}_r)_{r=1}^t$ is any family of quadrupels of finite sets with the above properties, then the number t of quadrupels satisfies the bound $t \leq 2 \sum_{i=1}^j \binom{n-d-\ell+i}{i} \binom{d-j+\ell}{d-i}$. This implies Theorem 2.4.

3 Quadrupels of Sets

By Lemma 2.6, the following theorem implies the desired bound for the numbers $h_{\leq j, \leq \ell}$.

Theorem 3.1. Let j, k, ℓ, m be integers such that $0 \leq j \leq k$ and $0 \leq \ell \leq m$. Suppose that (A_r, B_r, C_r, D_r) , $1 \leq r \leq t$, are quadrupels of finite sets with the following properties:

1. For all r , $A_r \subseteq B_r$, $C_r \subseteq D_r$, $C_r \cap D_r = \emptyset$, $|A_r| \leq j$, $|B_r| = k$, $|C_r| \leq \ell$, and $|D_r| = m$.
2. For $1 \leq r < s \leq t$,

$$(A_r \cap (D_s \setminus C_s)) \cup ((B_r \setminus A_r) \cap C_s) \neq \emptyset. \quad (3)$$

Then the number t of quadrupels satisfies

$$t \leq 2 \cdot \sum_{i=0}^j \binom{m-\ell+j}{i} \binom{k-j+\ell}{k-i} \quad (4)$$

Remarks 3.2. 1. The theorem is completely symmetric in $A \leftrightarrow C$, $B \leftrightarrow D$, $j \leftrightarrow \ell$, $k \leftrightarrow m$, by reversing the ordering among the quadrupels, $r \leftrightarrow t - r + 1$. In the geometric setting, this symmetry corresponds to LP-duality.

2. In the case $\ell = 0$, the second part of the intersection condition (3) is trivially false for all r, s . Moreover, in the bound (4), all summands with $i < j$ are zero. Therefore, we obtain the following assertion about set pairs: Let $(A_r, D_r)_{r=1}^t$ be a family of pairs of sets with $|A_r| \leq j$, $|D_r| = m$, and $A_r \cap D_r = \emptyset$ for all r . If $A_r \cap D_s \neq \emptyset$ whenever $r < s$, then $t \leq 2 \cdot \binom{j+m}{j}$. This bound holds without the factor of 2 and is known as the *Set-Pair Lemma* or the *skew version of Bollobás's Theorem*. It was conjectured by Bollobás, who proved it, by purely combinatorial means, under the stronger, symmetric condition that $A_r \cap D_s \neq \emptyset$ for all $r \neq s$ [9]. The skew version was first proved by Frankl [15], using exterior algebra methods introduced by Lovász [26].
3. The same method was independently invented by Kalai [21], who used it, among other things, to prove that the upper bound in (4), without the factor of 2, holds for the number of k -element sets in a j -collapsible simplicial complex of dimension less than $j + m$ on a ground set of $m + k$ elements. However, there does not seem to be a direct way of deriving Theorem 3.1 from Kalai's result. Recall that a simplicial complex is j -collapsible if it can be reduced to the empty complex by a finite sequence of *elementary* j -collapses. An elementary j -collapse means removing from K an interval of faces $[F, M] := \{S \in K : F \subseteq S \subseteq M\}$, where M is an inclusion-maximal face of K and F is a *free* face, i.e., M is the unique inclusion-maximal face containing F . Consider the following sequence of quadrupels $(A, B \setminus A, C, D \setminus C)$ of subsets of

$\{1, \dots, 6\}$ (where we right 236, for instance, instead of $\{2, 3, 6\}$): $(2, 13, \emptyset, 456)$, $(2, 14, 3, 56)$, $(3, 14, \emptyset, 256)$, $(3, 24, 1, 56)$, $(3, 15, 4, 26)$, $(4, 15, \emptyset, 236)$, $(4, 25, 1, 36)$, $(4, 16, 5, 23)$, $(5, 16, \emptyset, 234)$, $(5, 26, 1, 34)$. These quadrupels satisfy the assumptions of Theorem 4 with $j = \ell = 1$ and $k = m = 3$. However, for every element a of the ground set $\{1, \dots, 6\}$, the union of the sets B containing a has cardinality at least 5. Thus, there is no 1-collapsible complex of dimension less than 4 that contains all the sets B .

4. Alon and Kalai [3] reproved the skew version of Bollobás Theorem and used it to prove the Upper Bound Theorem for convex polytopes. Our proof is inspired by their approach and by Kalai's paper [21]. For variants and further applications of the Set-Pair Lemma, see the surveys by Tuza [42, 43].

Example 3.3 (A Combinatorial Lower Bound). We define

$$\mathcal{B}(j, k, \ell, m) := \left\{ B \in \binom{[m+k]}{k} : |B \cap [m-\ell+j]| \leq j \right\}.$$

For $B \in \mathcal{B}(j, k, \ell, m)$, let $A = B \cap [m-\ell+j]$, $D = [m+k] \setminus B$, and $C = D \cap [m-\ell+j+1, \dots, m+k]$. We have $(A \cap (D' \setminus C')) \cup ((B \setminus A) \cap C') \neq \emptyset$ whenever $B \neq B'$, so the quadrupels (A, B, C, B) satisfy the conditions of Theorem 3.1 for any ordering; in other words, they provide a lower bound example, for the sharp version without the factor of 2, even in the symmetric (i.e., non-skew) case.

4 Algebraic Background

In this section, we briefly review the basic algebraic facts that we will require for the proof of Theorem 3.1. For an extensive introduction to (multi)linear-algebraic methods in combinatorics, we refer the reader to the notes by Babai and Frankl [5]. For a detailed treatment of multilinear algebra, see also [28].

Let F be an arbitrary field. All linear-algebraic notions (vector spaces, (multi)linear maps, algebras, etc.) will be with respect to this ground field. It will be convenient to choose F to be of characteristic 2, which will allow us to disregard signs in our computations.

Exterior Algebra. An *algebra* is a vector space L over that is additionally equipped with a map (or “*product*”) $L \times L \rightarrow L$, denoted by $(x, y) \mapsto x \cdot y$, that is bilinear and associative.

Let V be n -dimensional vector space over F . We denote its exterior algebra by $\bigwedge V = \bigoplus_{k=0}^n \bigwedge^k V$ (where $\bigwedge^0 V = F$ and $\bigwedge^1 V = V$) and the *exterior* or *wedge product* on $\bigwedge V$ by \wedge . We have $v \wedge v = 0$ for all $v \in V = \bigwedge^1 V$. In

characteristic 2, this implies commutativity: $x \wedge y = y \wedge x$ for all $x, y \in \wedge V$.

The elements of the form $v_1 \wedge \dots \wedge v_k$, $v_i \in V$, linearly span $\wedge^k V$. More precisely, we have the following: If $\{v_1, \dots, v_n\}$ is an ordered family of vectors in V and if $S = \{i_1 < \dots < i_k\} \subseteq [n]$, then we define $v_S := v_{i_1} \wedge \dots \wedge v_{i_k}$ (with the convention that $v_\emptyset = 1$). If $\{v_1, \dots, v_n\}$ is a basis of V , then $\{v_S : S \in \binom{[n]}{k}\}$ is a basis of $\wedge^k V$, which implies $\dim \wedge^k V = \binom{n}{k}$ and $\dim \wedge V = 2^n$. If $S, T \subseteq [n]$, then we have

$$v_S \wedge v_T = \begin{cases} v_{S \cup T} & \text{if } S \cap T = \emptyset, \\ 0 & \text{otherwise,} \end{cases}$$

If $f : V \rightarrow W$ is a linear map between vector spaces, then there is a linear map $\wedge f : \wedge V \rightarrow \wedge W$ that extends f (i.e., $\wedge f(v) = f(v)$ for all $v \in V$) and that satisfies $\wedge f(x \wedge y) = \wedge f(x) \wedge \wedge f(y)$. This map respects the direct sum decomposition, i.e., $\wedge f = \bigoplus_k \wedge^k f$ with $\wedge^k f : \wedge^k V \rightarrow \wedge^k W$. If f has the matrix representation $A = [a_{ij}]$ in bases $\{v_j\}_{j=1}^n$ and $\{w_i\}_{i=1}^m$ of V and W , respectively, then $\wedge^k f$ has the matrix representation $[\det A_{S|T}]$ in the bases $\{w_S\}$ and $\{v_T\}$, where S and T range over all k -element subsets of $[m]$ and of $[n]$, respectively, and $A_{S|T} = [a_{ij}]_{i \in S, j \in T}$.

Scalar and Left Interior Products. Suppose that V is a finite-dimensional vector space that is additionally equipped with a *scalar product*, i.e., a map $\langle \cdot, \cdot \rangle : V \times V \rightarrow F$, that is bilinear and *symmetric*, i.e., $\langle v, w \rangle = \langle w, v \rangle$ for all $v, w \in V$. We assume that the scalar product is *non-singular* in the sense that if a vector $v \in V$ satisfies $\langle v, w \rangle = 0$ for all $w \in V$, then $v = 0$. In fact, we will make the stronger assumption that there exists an *orthonormal basis* $\{e_1, \dots, e_n\}$ of V with respect to the scalar product (we have to assume this because the Gram-Schmidt process can break down over general fields; it works provided the characteristic is $\neq 2$ and all the field elements of the form $\langle v, v \rangle$, $v \in V$, are squares).

The scalar product on V can be extended to one on $\wedge V$, which we will denote by the same symbol, in such a way that $\{e_S\}$ is an orthonormal basis, where S ranges over all subsets of $[n]$. We define the *left interior product* $\wedge V \times \wedge V \rightarrow \wedge V$ as follows: For $x, y \in \wedge V$, $x \lrcorner y$ is uniquely determined by the requirement that

$$\langle x \lrcorner y, z \rangle = \langle y, z \wedge x \rangle$$

for all $z \in \wedge V$. (Thus, left interior multiplication by x is the adjoint of wedge multiplication by x on the right.) For subsets $R, S \subseteq [n]$, it follows that

$$e_R \lrcorner e_S = \begin{cases} \pm e_{S \setminus R} & \text{if } R \subseteq S, \\ 0 & \text{if } R \not\subseteq S; \end{cases}$$

in fact, we could use this as a definition of the left interior product, extending it to general elements by bilinearity.

We will need the following facts concerning the behaviour of wedge and left interior products. Again, working in characteristic 2 allows us to ignore signs. These properties are easily established for basis elements e_S and then follow by linearity for general elements.

Lemma 4.1. 1. For all $x, y, z \in \wedge V$, $x \lrcorner (y \lrcorner z) = (x \wedge y) \lrcorner z$.

2. Let $v \in V = \wedge^1 V$, and let $a, b \in \wedge V$. Then

$$v \lrcorner (a \wedge b) = (v \lrcorner a) \wedge b + a \wedge (v \lrcorner b).$$

3. From the previous, it follows by induction that for $v_1, \dots, v_n \in V$ and $S \subseteq [n]$,

$$e_S \lrcorner (a \wedge b) = \sum_{X, Y} (e_{X \lrcorner a}) \wedge (e_{Y \lrcorner b}),$$

where the summation runs over all ordered partitions (X, Y) of S (i.e., ordered pairs (X, Y) with $X \cap Y = \emptyset$ and $X \cup Y = S$).

4. For any partition $S = X \cup Y$, we have

$$e_S \wedge a \wedge b = \pm e_X \wedge a \wedge e_Y \wedge b.$$

Tensor Products If U and V are vector spaces, then $U \otimes V$ denotes their tensor product. If $\{u_i\}$ and $\{v_j\}$ are bases of U and of V , respectively, then $\{u_i \otimes v_j\}$ is basis of $U \otimes V$. Consequently, $\dim(U \otimes V) = \dim U \cdot \dim V$.

If U and V are both algebras, then the tensor product inherits a multiplication and also becomes an algebra: Generators are multiplied by the rule $(u \otimes v) \cdot (u' \otimes v') = (u \cdot u') \otimes (v \cdot v')$ and this is extended to general elements by bilinearity. Here, we abuse notation and denote the products in U , V , and $U \otimes V$ by the same symbol. If the factors U and V each carry several multiplicative structures (for instance, if they are exterior algebras, for which we have both the wedge product and the left interior product), we will specify which multiplication is meant in each factor.

If $f : U \rightarrow X$ and $g : V \rightarrow Y$ are linear maps, they induce a linear map $f \otimes g : U \otimes V \rightarrow X \otimes Y$ such that $(f \otimes g)(u \otimes v) = f(u) \otimes g(v)$ holds on generators. If $A = [a_{ij}]$ and $B = [b_{k\ell}]$ are matrix representations of f and g with respect to bases $\{u_i\}$, $\{x_j\}$, $\{v_k\}$, and $\{y_\ell\}$ of U , V , X , and Y , respectively then $f \otimes g$ is represented by the matrix $A \otimes B := [a_{ij} b_{k\ell}]$ with respect to the bases $\{u_i \otimes v_k\}$ and $\{x_j \otimes y_\ell\}$.

5 The Proof of the Main Theorem

Throughout this section, we consider a fixed system $(A_r, B_r, C_r, D_r)_{r=1}^t$ of set quadrupels that satisfy the assumptions of Theorem 3.1 with parameters j, k, ℓ, m . We want to show that $t \leq 2 \cdot |\mathcal{B}(j, k, \ell, m)|$ (as defined in Example 3.3). The general strategy the proof is encapsulated in the following lemma (which is implicit in [15]).

Lemma 5.1. *Let Λ be an algebra over a field F . Assume that we are given vectors $x_1, \dots, x_t, y_1, \dots, y_t \in \Lambda$ such that $x_r y_r \neq 0$ for $1 \leq r \leq t$ and $x_r y_s = 0$ for $1 \leq r < s \leq t$. Then the vectors x_1, \dots, x_t are linearly independent (analogously, so are the vectors y_r). In particular, if all vectors x_r lie in some linear subspace $\Sigma \subseteq \Lambda$, then $t \leq \dim \Sigma$.*

To prove this, assume that we have a nontrivial linear dependence $\sum_r \alpha_r x_r$'s, and multiply the equation by y_s , where $s = \max\{r : \alpha_r \neq 0\}$, to get a contradiction.

The ground field F . We fix a ground field F of characteristic 2. We choose F such that it contains infinitely many elements that are mutually algebraically independent over the field \mathbb{F}_2 with two elements, i.e., that do not satisfy any nonzero polynomial with coefficients in \mathbb{F}_2 . For instance, we can take F to be the field of rational functions in countably many variables x_1, x_2, x_3, \dots over \mathbb{F}_2 .

The algebra Λ and the subspace Σ . The choices for Λ and σ are rather natural in light of the bound that we wish to achieve.

Let U and W be vector spaces over F of dimensions $m - \ell + j$ and $k - j + \ell$, respectively. We take

$$\Lambda = \bigwedge U \otimes \bigwedge W$$

and

$$\Sigma = \bigoplus_{i=0}^j \left(\bigwedge^i U \otimes \bigwedge^{k-i} W \right).$$

Thus, $\dim \Sigma = |\mathcal{B}(j, k, \ell, m)|$.

The vectors x_r and y_r . Both the construction of suitable vectors $x_r \in \Sigma$, $y_r \in \Lambda$ associated with our set quadrupels and the proof of the desired properties are somewhat involved. For this reason, we present them in stages, as follows: First, we consider the special case that $j - |A_r| = \ell - |C_r|$ for all r . While this is a very restrictive assumption, the basic idea of the proof works in its simplest and most transparent form in this case. Next, we show how to modify the embedding so that it works for the case that either $j - |A_r| \geq \ell - |C_r|$ for all r or $j - |A_r| \leq \ell - |C_r|$ for all r . Since for each index r , one of these inequalities is

satisfied, this proof also works in the general case, but only yields a bound that is twice as large as desired.

5.1 The Case $j - |A_r| = \ell - |C_r|$

Proposition 5.2. *Let (A_r, B_r, C_r, D_r) , $1 \leq r \leq t$, be quadrupels of finite sets that satisfy the assumptions of Theorem 3.1. Assume furthermore that $j - |A_r| = \ell - |C_r|$ holds for all r . Then $t \leq |\mathcal{B}(j, k, \ell, m)|$.*

Note that the additional assumption implies (in fact, is equivalent to) $|A_r| + |D_r \setminus C_r| = m - \ell + j$ and $|B_r \setminus A_r| + |C_r| = k - j + \ell$ for all r . Furthermore, under the assumption, $|A_r| \leq j$ implies $|C_r| \leq \ell$, and vice versa, so we only have to explicitly assume one of these.

Generic vectors. Let q be a positive integer, and let V be a q -dimensional vector space over F with a fixed basis e_1, \dots, e_q . We call vectors $v_1, \dots, v_n \in V$ *generic* if their coordinates with respect to the basis $\{e_\mu\}$ are algebraically independent over \mathbb{F}_2 , i.e., $v_\nu = \sum_{\mu=1}^q \xi_{\mu\nu} e_\mu$ with algebraically independent $\xi_{\mu\nu}$.

Lemma 5.3. *If $v_1, \dots, v_n \in V$ are generic, then $v_S \neq 0$ for every $S \in \binom{[n]}{q}$.*

Proof. If X denotes the $q \times n$ -matrix $[\xi_{\mu\nu}]$ of coordinates, then $v_S = (\det X_{[q]|S}) e_{[q]}$, with the notation for minors as in Section 4. This determinant is a sum of distinct monomials of degree q in the $\xi_{\mu\nu}$'s and therefore does not vanish. \square

Proof of Proposition 5.2. We work with the algebra Λ and the subspace Σ as defined above. Without loss of generality, assume that all sets are subsets of $[n]$, for some suitable integer n . Let $u_1, \dots, u_n \in U$ and $w_1, \dots, w_n \in W$ be such that $u_S \neq 0$ and $w_T \neq 0$ for every $S \in \binom{[n]}{m-\ell+j}$ and every $T \in \binom{[n]}{k-j+\ell}$; for instance, by the preceding lemma, we can choose the u 's and w 's generically.

For each r , we define

$$x_r := u_{A_r} \otimes w_{B_r \setminus A_r}, \quad y_r := u_{D_r \setminus C_r} \otimes w_{C_r}.$$

We claim that the vectors x_r and y_r satisfy the assumptions of Lemma 5.1; the proposition then follows. We have

$$x_r \cdot y_s = (u_{A_r} \wedge u_{D_r \setminus C_r}) \otimes (w_{B_r \setminus A_r} \wedge w_{C_r}).$$

For $r < s$, the intersection condition $(A_r \cap (D_s \setminus C_s)) \cup ((B_r \setminus A_r) \cap C_s) \neq \emptyset$ implies that at least one of the two wedge products vanishes, and hence so does their tensor product. For $r = s$, we have $u_{A_r} \wedge u_{D_r \setminus C_r} = u_{A_r \cup (D_r \setminus C_r)}$ and $w_{B_r \setminus A_r} \wedge w_{C_r} = w_{(B_r \setminus A_r) \cup C_r}$. By genericity and by assumption on the cardinalities, neither of these wedge

products vanishes, and hence neither does their tensor product. This proves the claim and hence the proposition. Note that the additional assumption on the cardinalities was just needed to ensure that $x_r \cdot y_r \neq 0$. \square

5.2 The Case $|A_r| + |D_r \setminus C_r| \geq m - \ell + j$

Lemma 5.4. *Assume that (A_r, B_r, C_r, D_r) , $1 \leq r \leq t$, are quadrupels of finite sets that satisfy the assumptions of Theorem 3.1, and that all these sets are subsets of $[n]$, for some suitable integer n . Let e_1, \dots, e_n be the standard basis for F^n , and consider the algebra $\bigwedge F^n \otimes \bigwedge F^n$ with the multiplication induced by exterior multiplication in each factor. Then, for $1 \leq r < s \leq t$,*

$$\left(\sum_{R \subseteq A_r} e_R \otimes e_{B_r \setminus R} \right) \cdot \left(\sum_{S \subseteq D_s \setminus C_s} e_S \otimes e_{D_s \setminus S} \right) = 0. \quad (5)$$

Proof. We have

$$\sum_{R \subseteq A_r} e_R \otimes e_{B_r \setminus R} = \sum_{X \subseteq [n]} (e_X \lrcorner e_{A_r}) \otimes (e_X \wedge e_{B_r \setminus A_r})$$

and

$$\sum_{S \subseteq D_s \setminus C_s} e_S \otimes e_{D_s \setminus S} = \sum_{Y \subseteq [n]} (e_Y \lrcorner e_{D_s \setminus C_s}) \otimes (e_Y \wedge e_{C_s}).$$

Therefore, the product on the left-hand side of (5) equals

$$\sum_{X, Y} (e_X \lrcorner e_{A_r}) \wedge (e_Y \lrcorner e_{D_s \setminus C_s}) \otimes (e_X \wedge e_{B_r \setminus A_r}) \wedge (e_Y \wedge e_{C_s}).$$

Here, the sum ranges over all ordered pairs X, Y of subsets of $[n]$. However, if $X \cap Y \neq \emptyset$, then the wedge product on the right-hand side of the summand vanishes. Therefore, we may restrict our attention to pairs (X, Y) of disjoint sets. By grouping terms according to $Z = X \cup Y$ and by using linearity on the left-hand side of the tensor product, we get the sum

$$\sum_Z \underbrace{\left(\sum_{X, Y} (e_X \lrcorner e_{A_r}) \wedge (e_Y \lrcorner e_{D_s \setminus C_s}) \right)}_{(*)} \otimes e_Z \wedge e_{B_r \setminus A_r} \wedge e_{C_s}.$$

Here, the inner sum $(*)$ ranges over all ordered partitions (X, Y) of Z . Therefore, by Lemma 4.1, it equals $e_Z \lrcorner (e_{A_r} \wedge e_{D_s \setminus C_s})$. Consequently, the whole sum equals

$$\sum_{Z \subseteq [n]} e_Z \lrcorner (e_{A_r} \wedge e_{D_s \setminus C_s}) \otimes e_Z \wedge e_{B_r \setminus A_r} \wedge e_{C_s}.$$

In this sum, each summand equals

$$(e_Z \otimes e_Z) \circ \left((e_{A_r} \wedge e_{D_s \setminus C_s}) \otimes (e_{B_r \setminus A_r} \wedge e_{C_s}) \right),$$

where \circ is induced by left interior multiplication in the left factor of the tensor product and exterior multiplication in the right factor. Therefore, each summand vanishes, since $(e_{A_r} \wedge e_{D_s \setminus C_s}) \otimes (e_{B_r \setminus A_r} \wedge e_{C_s}) = 0$. Hence, the whole sum vanishes, which is what we wanted to prove. \square

Since Equation (5) is preserved under projections, we also have:

Corollary 5.5. *Let U, W be vector spaces over F , each of dimension at most n , and let $u_1, \dots, u_n \in U$ and $w_1, \dots, w_n \in W$. Then*

$$\left(\sum_{R \subseteq A_r} u_R \otimes w_{B_r \setminus R} \right) \cdot \left(\sum_{S \subseteq D_s \setminus C_s} u_S \otimes w_{D_s \setminus S} \right) = 0.$$

for $1 \leq r < s \leq t$.

Lemma 5.6. *Let U and W be vector spaces over F of dimensions p and q , respectively. Assume that $u_1, \dots, u_n \in U$ and $w_1, \dots, w_n \in W$ are such that the vectors $v_\nu = u_\nu + w_\nu \in V = U \oplus W$ are generic. Assume furthermore that $A \subseteq B$, $C \subseteq D$ are subsets of $[n]$, that B and D are disjoint and that $|B \cup D| = p + q$. If $|A| + |D \setminus C| \geq p$, then*

$$\left(\sum_{R \subseteq A} u_R \otimes w_{B \setminus R} \right) \cdot \left(\sum_{S \subseteq D \setminus C} u_S \otimes w_{D \setminus S} \right) \neq 0.$$

Proof. Without loss of generality, let $V = F^{p+q}$ with standard basis e_1, \dots, e_{p+q} and let U and W be the subspaces spanned by $\{e_1, \dots, e_p\}$ and $\{e_{p+1}, \dots, e_{p+q}\}$, respectively. Let $v_\nu = \sum_{\mu=1}^{p+q} \xi_{\mu\nu} e_\mu$, where $X = [\xi_{\mu\nu}]$ is a $(p+q) \times n$ -matrix whose entries are algebraically independent over \mathbb{F}_2 . We have

$$\begin{aligned} \left(\sum_{R \subseteq A} u_R \otimes w_{B \setminus R} \right) \cdot \left(\sum_{S \subseteq D \setminus C} u_S \otimes w_{D \setminus S} \right) \\ = \sum_{P \subseteq A \cup (D \setminus C)} u_P \otimes w_{(B \cup D) \setminus P}. \end{aligned}$$

In the sum on the right-hand side, all terms with $|P| \neq p$ vanish trivially. Thus, we are left with

$$\sum_{\substack{P \subseteq A \cup (D \setminus C) \\ |P|=p}} u_P \otimes w_{(B \cup D) \setminus P}.$$

Moreover, for $|P| = p$ and $|Q| = q$, we have $u_P = \det X_{[p]|P} \cdot e_{[p]}$ and $w_Q = \det X_{[p+1, \dots, p+q]|Q} \cdot e_{[p+1, \dots, p+q]}$. Expanding these determinants (recall that we work in characteristic 2), we see that the sum we wish to evaluate equals

$\lambda \cdot e_{[p]} \otimes e_{[p+1, \dots, p+q]}$, where

$$\begin{aligned} \lambda &= \sum_{P \in \binom{A \cup (D \setminus C)}{p}} \sum_{\substack{\sigma: [p] \rightarrow P, \\ \tau: [p+1, \dots, p+1] \rightarrow (B \cup D) \setminus P}} \prod_{\mu=1}^p \xi_{\mu, \pi(\mu)} \prod_{\mu=p+1}^{p+q} \xi_{\mu, \tau(\mu)} \\ &= \sum_{\substack{\pi: [p+q] \rightarrow B \cup D \\ \pi([p]) \subseteq A \cup (D \setminus C)}} \prod_{\mu=1}^{p+q} \xi_{\mu, \pi(\mu)}, \end{aligned}$$

and the sums run over all bijective maps σ, τ, π with the specified domains and ranges. Thus, λ is the sum of distinct monomials in the $\xi_{\mu\nu}$'s and therefore nonzero. \square

Proposition 5.7. *Let (A_r, B_r, C_r, D_r) , $1 \leq r \leq t$, be quadrupels of finite sets that satisfy the assumptions of Theorem 3.1. Assume furthermore that either $|A_r| + |D_r \setminus C_r| \geq m - \ell + j$ for all r or $|A_r| + |D_r \setminus C_r| \leq m - \ell + j$ for all r . Then*

$$t \leq |\mathcal{B}(j, k, \ell, m)| = \sum_{i=0}^j \binom{m - \ell + j}{i} \binom{k - j + \ell}{k - i}.$$

Proof. We consider the case that $|A_r| + |D_r \setminus C_r| \geq m - \ell + j$ holds for all r . (Otherwise $|A_r| + |D_r \setminus C_r| \leq m - \ell + j$ for all r , hence $|B_r \setminus A_r| + |C_r| \geq k - j + \ell$ for all r , and the result follows symmetrically.)

We let $V = F^{k+m}$ with standard basis e_1, \dots, e_{k+m} and define U and W be the span of $\{e_1, \dots, e_{m-\ell+j}\}$ and $\{e_{m-\ell+j+1}, \dots, e_{k+m}\}$, respectively. Let $v_1, \dots, v_n \in V$ be generic, and let u_i and w_i the projections of v_i onto U and W , respectively, $1 \leq i \leq n$. For $1 \leq r \leq t$, we define $x_r := \sum_{R \subseteq A_r} u_R \otimes w_{B_r \setminus R}$ and $y_r := \sum_{S \subseteq D_r \setminus C_r} u_S \otimes w_{D_r \setminus S}$. By Corollary 5.5 and Lemma 5.6, these vectors satisfy the assumptions of Lemma 5.1 with Λ and Σ as before, and the bound follows. \square

Proof of Theorem 3.1. The Main Theorem 3.1 follows immediately from Proposition 5.7, since from any sequence of quadrupels, we can always extract a sequence of at least half the length that satisfies the additional conditions on the cardinalities. \square

6 Concluding Remarks

The most obvious open problem is get rid of the factor of 2 in Theorem 3.1 (and consequently in Theorems 1.3 and 2.4, and in Corollary 2.5).

A second problem is to prove geometric lower bounds for the numbers $v_{\leq \ell}$ in the range $\lceil n/(d+1) \rceil \leq \ell \leq (n-d-1)/2$. In order to show that the AGUBC (without the factor of 2) is sharp, i.e., that $v_{\leq \ell}(\mathcal{A}) = v_{\leq \ell}(\mathcal{C}_{n,d}^*)$ can be achieved for some affine arrangement \mathcal{A} , we have to answer the following question: Is there a neighborly d -polytope on n vertices whose vertex set is ℓ -centered around

some point $o \in \mathbf{R}^d$? Here, a set S of n points in \mathbf{R}^d is $(\ell+1)$ -centered around some point $o \in \mathbf{R}^d$ if every affine hyperplane through o contains at least ℓ points of S on either side (in each of the two open halfspaces).

Another intriguing question is how to define a good analogue of the h -matrix for the spherical case. One conceivable approach is the following: If the $(\leq \ell)$ -level of an affine arrangement \mathcal{A} is bounded, then based on Equation 2 one can show that the numbers $h_{j,k}(\mathcal{A}, \varphi)$, $0 \leq j \leq d$ and $k \leq \ell$, can be expressed as linear combinations of the numbers f_m^r , of r -dimensional faces at level m , $0 \leq r \leq d$, $m \leq \ell$ in a way that is independent of φ (see [32]). In particular, the $(\leq \ell)$ -portion of the h -matrix is independent of φ . It is tempting to use the resulting formulas to define the h -matrix for the spherical case. However, there are simple examples for which this definition yields negative entries, so most likely the entries of this h -matrix no longer have a combinatorial interpretation. For this reason, it seems questionable if this is a promising definition.

The definitions of levels and of the h -matrix find their most natural setting in the context of oriented matroids (or equivalently, by the Folkman-Lawrence representation theorem, of pseudo-hemisphere arrangements). We refer the reader to the book [8] as a general reference for oriented matroids. The original Upper Bound Theorem also holds for oriented matroids. In fact, the 0-level of an oriented matroid forms a simplicial sphere, so the Upper Bound Theorem holds by Stanley's results. Moreover, for the particular case of oriented matroids, a more elementary proof is available (essentially McMullen's original proof can be suitably generalized, see Kleinschmidt and Onn [24]). Our methods immediately carry over to so-called *Euclidean* oriented matroids and yield an analogue of the 2AGUBT in that case. It would be interesting to find a way to extend the 2AGUBT to the non-Euclidean case, or even better, to prove the AGUBC or the SGUBC for general oriented matroids.

Acknowledgements. I am indebted to Gil Kalai and Emo Welzl for helpful comments and discussions.

References

- [1] N. Alon, I. Bárány, Z. Füredi, and D. J. Kleitman. Point selections and weak ϵ -nets for convex hulls. *Combin. Probab. Comput.*, 1(3):189–200, 1992.
- [2] N. Alon and E. Gyöfi. The number of small semispaces of a finite set of points in the plane. *J. Comb. Theory, Ser. A*, 41:154–157, 1986.
- [3] N. Alon and G. Kalai. A simple proof of the Upper Bound Theorem. *European Journal of Combinatorics*, 6:211–214, 1985.
- [4] A. Andrzejak and E. Welzl. In between k -sets, j -facets, and i -faces: (i, j) -partitions. *Discrete Comput. Geom.*, 29(1):105–131, 2003.

- [5] L. Babai and P. Frankl. Linear algebra methods in combinatorics. Lecture Notes. University of Chicago, 1992.
- [6] J. Balogh and G. Salazar. Improved bounds for the number of ($\leq k$)-sets, convex quadrilaterals, and the rectilinear crossing number of K_n . In J. Pach, editor, *Graph Drawing. 12th International Symposium, GD 2004, New York, NY, September 29-October 2, 2004. Revised selected papers*, volume 3383 of *Lecture Notes in Computer Science*, pages 25–35. Springer-Verlag, Berlin, 2005.
- [7] L. J. Billera and C. W. Lee. A proof of the sufficiency of McMullen’s conditions for f -vectors of simplicial convex polytopes. *Journal of Combinatorial Theory*, 31(3):237–255, 1981.
- [8] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler. *Oriented matroids*, volume 46 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, second edition, 1999.
- [9] B. Bollobás. On generalized graphs. *Acta Math. Acad. Sci. Hungar.*, 16:447–452, 1965.
- [10] P. Brass, W. Moser, and J. Pach. *Research problems in discrete geometry*. Springer, New York, 2005.
- [11] T. M. Chan. Low-dimensional linear programming with violations. *SIAM J. Comput.*, 34(4):879–893 (electronic), 2005.
- [12] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry. II. *Discrete Comput. Geom.*, 4(5):387–421, 1989.
- [13] T. K. Dey. Improved bounds on planar k -sets and related problems. *Discrete Comput. Geom.*, 19:373–382, 1998.
- [14] J. Eckhoff. Helly, Radon, and Carathéodory type theorems. In P. M. Gruber and J. M. Wills, editors, *Handbook of Convex Geometry*, pages 389–448. North-Holland, Amsterdam, 1993.
- [15] P. Frankl. An extremal problem for two families of sets. *European J. Combin.*, 3(2):125–127, 1982.
- [16] M. R. Garey and D. S. Johnson. Crossing number is NP-complete. *SIAM J. Algebraic Discrete Methods*, 4(3):312–316, 1983.
- [17] B. Gärtner and J. Matoušek. *Understanding and Using Linear Programming*. Universitext. Springer-Verlag, 2006.
- [18] R. K. Guy. A combinatorial problem in geometry. *Nabla (Bull. Malayan Math. Soc.)*, 7:68–72, 1960.
- [19] R. K. Guy, T. Jenkyns, and J. Schaer. The toroidal crossing number of the complete graph. *J. Comb. Theory*, 4:376–390, 1968.
- [20] P. Hersh and I. Novik. A short simplicial h -vector and the upper bound theorem. *Discrete Comput. Geom.*, 28(3):283–289, 2002.
- [21] G. Kalai. Intersection patterns of convex sets. *Israel J. Math.*, 48(2-3):161–174, 1984.
- [22] G. Kalai. The diameter of graphs of convex polytopes and f -vector theory. In *Applied geometry and discrete mathematics*, volume 4 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 387–411. Amer. Math. Soc., Providence, RI, 1991.
- [23] V. Klee. On the number of vertices of a convex polytope. *Canad. J. Math.*, 16:701–720, 1964.
- [24] P. Kleinschmidt and S. Onn. Signable posets and partitionable simplicial complexes. *Discrete Comput. Geom.*, 15(4), 1996.
- [25] J. Linhart. The upper bound conjecture for arrangements of half-spaces. *Beitr. Algebra Geom.*, 35(1):29–35, 1994.
- [26] L. Lovász. Flats in matroids and geometric graphs. In *Combinatorial surveys (Proc. Sixth British Combinatorial Conf., Royal Holloway Coll., Egham, 1977)*, pages 45–86. Academic Press, London, 1977.
- [27] L. Lovász, K. Vesztegombi, U. Wagner, and E. Welzl. Convex quadrilaterals and k -sets. In *Towards a theory of geometric graphs*, volume 342 of *Contemp. Math.*, pages 139–148. Amer. Math. Soc., Providence, RI, 2004.
- [28] M. Marcus. *Finite-Dimensional Multilinear algebra. Part I*. Marcel Dekker Inc., New York, 1973. Pure and Applied Mathematics, Vol. 23.
- [29] J. Matoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.
- [30] J. Matoušek, M. Sharir, S. Smorodinsky, and U. Wagner. k -Sets in 4 dimensions. *Discrete Comput. Geom.*, 2006.
- [31] P. McMullen. On the upper-bound conjecture for convex polytopes. *J. Comb. Theory, Ser. B*, 10:187–200, 1971.
- [32] K. Mulmuley. Dehn-sommerville relations, upper bound theorem, and levels in arrangements. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 240–246, 1993.
- [33] I. Novik. Upper bound theorems for homology manifolds. *Israel J. Math.*, 108:45–82, 1998.
- [34] I. Novik. Remarks on the upper bound theorem. *J. Combin. Theory Ser. A*, 104(1):201–206, 2003.
- [35] G. W. Peck. On “ k -sets” in the plane. *Discrete Math.*, 56(1):73–74, 1985.
- [36] M. Sharir, S. Smorodinsky, and G. Tardos. An improved bound for k -sets in three dimensions. *Discrete Comput. Geom.*, 26(2):195–204, 2001. ACM Symposium on Computational Geometry (Hong Kong, 2000).
- [37] I. Shemer. Neighborly polytopes. *Israel J. Math.*, 43(4):291–314, 1982.
- [38] R. P. Stanley. The upper bound conjecture and Cohen-Macaulay rings. *Studies in Appl. Math.*, 54(2):135–142, 1975.
- [39] R. P. Stanley. The number of faces of a simplicial convex polytope. *Advances in Mathematics*, 35(3):236–238, 1980.
- [40] G. Tóth. Point sets with many k -sets. In *Proceedings of the Sixteenth Annual Symposium on Computational Geometry*, 2000.
- [41] P. Turán. A note of welcome. *J. Graph Theory*, 1:7–9, 1977.
- [42] Z. Tuza. Applications of the set-pair method in extremal hypergraph theory. In *Extremal problems for finite sets (Visegrád, 1991)*, volume 3 of *Bolyai Soc. Math. Stud.*, pages 479–514. János Bolyai Math. Soc., Budapest, 1994.
- [43] Z. Tuza. Applications of the set-pair method in extremal problems. II. In *Combinatorics, Paul Erdős is eighty, Vol. 2 (Keszthely, 1993)*, volume 2 of *Bolyai Soc. Math. Stud.*, pages 459–490. János Bolyai Math. Soc., Budapest, 1996.
- [44] E. Welzl. Entering and leaving j -facets. *Discrete Comput. Geom.*, 25(3):351–364, 2001.
- [45] G. Wesp. The upper bound conjecture for arrangements of halfspaces. *Beitr. Algebra Geom.*, 42(2):431–437, 2001.
- [46] R. T. Živaljević and S. T. Vrećica. The colored Tverberg’s problem and complexes of injective functions. *J. Combin. Theory Ser. A*, 61(2):309–318, 1992.