Grid orientations, (d,d+2)-polytopes, and arrangements of pseudolines

Author(s):
Felsner, Stefan; Gärtner, Bernd; Tschirschnitz, Falk

Publication date:
2005-09

Permanent link:
https://doi.org/10.3929/ethz-b-000031881

Originally published in:
Discrete & Computational Geometry 34(3), https://doi.org/10.1007/s00454-005-1187-x

Discrete Comput Geom 34:411–437 (2005)
DOI: 10.1007/s00454-005-1187-x

Discrete & Computational
Geometry

Grid Orientations, $(d, d+2)$-Polytopes, and Arrangements of Pseudolines*

Stefan Felsner,[1] Bernd Gärtner,[2] and Falk Tschirschnitz[2]

[1]Institut für Mathematik, Technische Universität Berlin,
D-10623 Berlin, Germany
felsner@math.tu-berlin.de

[2]Institute of Theoretical Computer Science, ETH Zürich,
CH-8092 Zürich, Switzerland
gaertner@inf.ethz.ch
falk@tschirschnitz.de

Abstract. We investigate the combinatorial structure of linear programs on simple d-polytopes with $d+2$ facets. These can be encoded by *admissible* grid orientations. Admissible grid orientations are also obtained through orientation properties of a planar point configuration or by the dual line arrangement. The point configuration and the polytope corresponding to the same grid are related through an extended Gale transform.

The class of admissible grid orientations is shown to contain nonrealizable examples, i.e., there are admissible grid orientations which cannot be obtained from a polytope or a point configuration. It is shown, however, that every admissible grid orientation is induced by an arrangement of pseudolines. This later result is used to prove several nontrivial facts about admissible grid orientations.

1. Introduction

In this paper we investigate the graphs of simple d-polytopes with $d+2$ facets (also called $(d, d+2)$-polytopes), oriented by means of linear or *abstract* objective functions. Abstract objective functions are designed to capture the combinatorial essence of linear objective functions and have been studied by several researchers [1], [27], [28], [17], [20].

The vertex-edge graph of a $(d, d+2)$-polytope is an (n, m)-*grid* for some n, m with $n+m = d+2$. This is a graph whose vertex set is the Cartesian product of an n-set and

* The second author acknowledges support from the Swiss Science Foundation (SNF), Project No. 200021-100316/1.

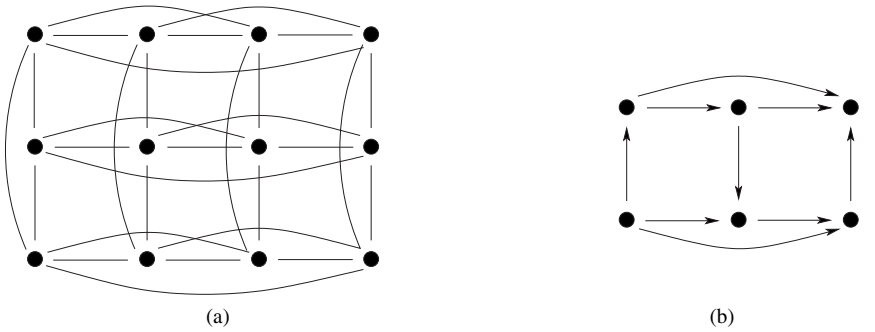

(a) (b)

Fig. 1. (a) A (3, 4)-grid and (b) the *double-twist*.

an m-set, with edges connecting all pairs of vertices that differ in exactly one coordinate, see Fig. 1(a). Therefore, we are concerned with orientations of grids.

Via an extended Gale transform, a $(d, d+2)$-polytope P, along with a linear function f, maps to a planar point configuration in such a way that the orientations of certain point triples encode the order in which the vertices of P appear under f. This is the setup of *One line and N points* [11]. We introduce a dual scenario of *No point and N lines* in which the grid orientation can be read off a planar arrangement of red and blue lines.

If f is an abstract objective function (meaning that it orders the vertices—and thus orients the graph—in such a way that each nonempty face of P induces a subgraph with a unique sink), we only get a point configuration (or a line arrangement) if f is linear. In the general case, the vertex order can still be encoded by (abstract) oriented triples, but it is not clear that they have a geometric interpretation. In fact, a more recent result of Holt and Klee shows that abstract objective functions are missing a crucial combinatorial feature of linear functions: whenever the graph of a (not necessarily simple) d-polytope is oriented by means of a generic[1] linear function, there must be d directed paths from the unique source to the unique sink with pairwise disjoint interiors (we simply call them vertex-disjoint in the following). Of course, the same statement applies to all faces, the number of paths being the dimension of the face [18]. Imposing this *Holt–Klee condition* on our grid orientations, in addition to the axioms of abstract objective functions, we arrive at the class of *admissible grid orientations*. The *double-twist* of Fig. 1(b) is the smallest example of a nonadmissible orientation induced by an abstract objective function.

Even though admissible grid orientations seem close to linear functions from a combinatorial point of view, they still fail to capture the geometry: with a construction involving a deformed Pappus-configuration we show that admissible grid orientations exist that are not induced by linear functions. Nevertheless, the main result of this paper shows that not much is missing: we prove that an arrangement of red and blue *pseudolines* exists that encodes the orientation. Because it also holds that any such arrangement induces an admissible grid orientation, we get a complete characterization of the latter. For simple $(d, d+2)$-polytopes, we therefore exactly understand the difference be-

[1] We call a linear function generic if it is not constant on any edge.

tween vertex orderings induced by actual linear functions and the orderings we get from *Holt–Klee functions* that just behave like linear functions under the known combinatorial criteria. Because there are many more pseudoline arrangements than line arrangements, our characterization implies that the fraction of linear functions among the Holt–Klee functions tends to zero as d tends to infinity. This was previously shown by Develin for $(d, 2d)$-polytopes, more precisely, for d-cubes [6].

Mihalisin and Klee have shown that for 3-polytopes, Holt–Klee functions induce the same class of vertex orderings as linear functions [23]. Morris gave an example of a 4-cube orientation which is induced by a Holt–Klee function but not by a linear function [24]. Our Pappus-configuration yields an orientation of a simple $(d, d+2)$-polytope with this property, for $d = 7$. This is smallest possible, a consequence of the fact that any arrangement of fewer than $d + 2 = 9$ pseudolines is *stretchable* [4].

The paper is organized as follows. In Section 2 we formally introduce admissible grid orientations as combinatorial models for linear functions on $(d, d+2)$-polytopes, and we show how they arise from *One line and N points*, equivalently, from *No point and N lines*. Section 3 contains our main result, the characterization of admissible grid orientations in terms of arrangements of red and blue pseudolines. Actually, we show more: every grid orientation which is induced by an abstract objective function can be represented by a red–blue arrangement of curves with the property that curves of different color cross exactly once. In the admissible case, it also holds that there are no multiple crossings between curves of the same color.

In Section 4 a number of interesting properties of admissible grid orientations is derived from the characterization. For example, it is shown that any two admissible and more generally any two *unique sink orientations* of the (n, m)-grid (Definition 2.1) can be transformed into each other by a sequence of edge-flips, with all intermediate orientations being admissible or unique sink, respectively. Section 5 constructs an admissible orientation of the $(3, 6)$-grid that is not induced by a linear function. Moreover, we show that any admissible orientation of the $(2, m)$-grid is linearly induced, for all m.

The concluding Section 6 briefly discusses grids of higher dimension and addresses an interpretation of our result in the context of (*partial*) *chirotopes*.

2. Admissible Grid Orientations

2.1. *From Polytopes to Grid Graphs*

All d-polytopes with $d + 1$ facets are combinatorially equivalent to the standard d-simplex Δ_d, defined as $\Delta_d = \text{conv}\{e_1, \ldots, e_{d+1}\}$, where e_i is the ith unit vector in $\mathbb{R}^{d+1}$. The 1-skeleton, or vertex-edge graph, of the d-simplex is the complete graph K_{d+1}, and acyclic orientations are transitive tournaments, i.e., permutations. It follows easily that every acyclic orientation of the graph of the simplex is *induced* by some linear function f. This means that there is an arc (v, w) between adjacent vertices if and only if $f(v) > f(w)$.

Having just one extra facet compared with a simplex of the same dimension, the $(d, d+2)$-polytopes still have a simple, though nontrivial, structure. The vertex-edge graph of a $(d, d+2)$-polytope is a grid: a graph consisting of n rows and m columns,

$n + m = d + 2$, with the property that the vertices in any fixed row or column induce a complete graph, see Fig. 1(a). This is implied by the following:

Lemma 2.1. *Any simple d-polytope with $d + 2$ facets is combinatorially equivalent to the product of simplices $\Delta_{n-1} \times \Delta_{m-1}$, where $n + m = d + 2$ and $n, m > 1$.*

Proof. We use a well-known theorem by Grünbaum [15], [16, Result 5.1.1]: *every d-polytope with $e \geq d + 1$ facets is the intersection of an $(e - 1)$-simplex with a d-dimensional flat*. Thus, for P being a simple $(d, d + 2)$-polytope, there is a $(d + 1)$-simplex $\Delta \subseteq \mathbb{R}^{d+1}$ and a hyperplane $h \subseteq \mathbb{R}^{d+1}$ such that $P = \Delta \cap h$. Because P is simple, we can slightly perturb h without changing the combinatorial structure of $\Delta \cap h$, so we may assume that no vertex of Δ lies in h. Then the vertex set of Δ gets partitioned by h into two sets S and T with $|S| = n$, $|T| = m$, $n + m = d + 2$. We have $n, m > 1$, because P is not a simplex. Furthermore, all intersections of faces of Δ with h are faces of P, which implies that the 1-skeleton of P contains an (n, m)-grid, with rows indexed by S and columns indexed by T. Because this grid is already a d-regular graph, it must coincide with the 1-skeleton. Thus, P shares its 1-skeleton with $\Delta_{n-1} \times \Delta_{m-1}$, and because the combinatorial structure of a simple polytope is determined by its 1-skeleton [5], [19], the lemma follows. □

Now we fix some $(d, d + 2)$-polytope P and its grid $G(P)$. A subgraph of the grid induced by all vertices belonging to some subset of rows *and* some subset of columns is called a *subgrid*. The proof of the previous lemma also yields the fact that subgrids of $G(P)$ correspond bijectively to the faces of P. In particular, the vertices of the grid correspond to the vertices of P.

A generic linear function f induces an orientation of $G(P)$: we have $v \to w$ if $f(v) > f(w)$. A grid orientation which is induced in this way has the following three properties:

The orientation is acyclic,	(1)
every nonempty subgrid has a unique sink, and	(2)
no subgrid is isomorphic to the double-twist depicted in Fig. 1(b).	(3)

Property (1) is obvious, and (2) follows from the fact that the minimum of a generic linear function over a nonempty face is attained at a unique vertex of the face. Property (3), finally, is a consequence of the Holt–Klee condition: a (2, 3)-subgrid of $G(P)$ corresponds to a three-dimensional face of P, so there must be three vertex-disjoint paths from source to sink [18]. The double-twist has only two such paths.

Definition 2.1. A grid orientation with property (2) is called the *unique sink orientation*. If in addition, (1) and (3) hold, the orientation is called *admissible*.

We remark that properties (1) and (2) are just the axioms of orientations induced by *abstract objective functions* [1], [27], [28], [17], [20], specialized to our grid setting. In this specific setting, it also holds that (2) implies acyclicity [10].

A natural question (we gave the answer away already) is whether every admissible grid orientation is actually induced by a linear function on a $(d, d+2)$-polytope. Since this is an important property, we give it a name.

Definition 2.2. A grid orientation is called *realizable* if it is induced by a linear function on a $(d, d+2)$-polytope.

In realizable grid orientations, the Holt–Klee condition imposes a restriction on *any* subgrid, while our notion of admissible grid orientations takes into account only the restrictions coming from certain three-dimensional faces. In Section 4 we show that the other restrictions are implied.

2.2. *One Line and N Points*, *No Point and N (Pseudo)lines*

Grid graphs can also be used to encode the setting of *One line and N points* [11]. Formally, let $S = \{s_1, \ldots, s_N\}$ be a set of N points in general position in the plane (i.e., no three on a common line), and let ℓ be a vertical line which is disjoint from S. For such a configuration we use the notation (S, ℓ). For ease of reference we think of the points left of ℓ as blue points while the points right of ℓ are red. A two-colored pair of points $e \in \binom{S}{2}$, i.e., a pair separated by line ℓ, is called an *ℓ-edge*.

For an ℓ-edge e, we let $\text{below}(e)$ be the set of points in S which are below the line spanned by the two points of e. Given an ℓ-edge e and a point $s \in \text{below}(e)$, there is a unique ℓ-edge $\{s, s_e\}$ with $s_e \in e$. This edge is denoted by $\text{pivot}(e, s)$, see Fig. 2.

A configuration (S, ℓ) naturally induces a digraph G whose nodes are the ℓ-edges. From the node $e \in E$ we have an outgoing edge to each $e' = \text{pivot}(e, s)$ for $s \in \text{below}(e)$. Since each node is determined by a point on the left and one on the right side of ℓ, the underlying undirected graph is a grid graph, and it is not hard to see that G is an admissible grid orientation.

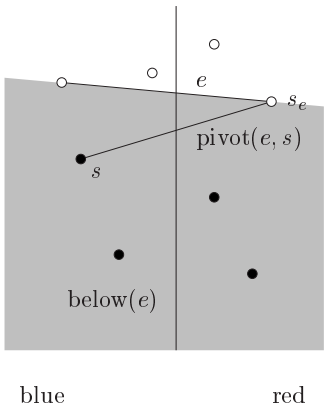

Fig. 2. One line and N points.

In fact, the setup of *One line and N points* has been conceived in order to study *realizable* grid orientations according to Definition 2.2: an extended Gale transform maps instances (S, ℓ) to instances (P, f) (and vice versa), where P is a $(d, d+2)$-polytope and f is a generic linear function. The crucial property of this transform is that the grid orientation induced by (S, ℓ) as described above coincides with the one induced by its image (P, f) as described in the previous subsection [11]. This yields the following:

Corollary 2.1. *A grid orientation is realizable if and only if it is induced by an instance* (S, ℓ) *of* one line and N points, *with at least two points of S on either side of* ℓ.

To motivate the introduction of our *No point and N lines* setup, fix an instance (S, ℓ) of *One line and N points* and consider the problem of finding the ℓ-edge e^* with the lowest intersection point y_{e^*} with line ℓ. Equivalently, we want the highest line with all points from S above.

Introducing coordinates and assuming that ℓ is the y-axis, i.e., the vertical through the origin, the problem of finding e^* can be translated to a linear program in two variables κ, δ:

$$\begin{array}{lll} \text{maximize} & \delta & \\ \text{subject to} & p_y \geq \kappa p_x + \delta & \text{for all} \quad (p_x, p_y) \in S. \end{array}$$

The same linear program models the following dual (polar) problem: With a point (p_x, p_y) in S associate the line $y = (-p_x)x + p_y$. Let the line associated with a point inherit its color. Blue lines are the lines with positive slope, red lines have negative slope. The quest is for the highest point below all lines. Equivalently, we seek the intersection point of a red and a blue line with the least y-coordinate. The general setting of *No point and N lines* consists of a set of N (nonvertical) lines partitioned into blue lines of positive slope and red lines of negative slope, and the problem is to find the red–blue intersection point ξ which is lowest.

The *No point and N lines* model has the advantage that the generalization from lines to pseudolines is quite apparent. Let $\mathcal{A}$ be a simple arrangement of N pseudolines in the Euclidean plane. To be simple means that there is no vertex of the arrangement where three or more pseudolines cross (this assumption is dual to the general position assumption in the (S, ℓ) model). To define the appropriate partition of the pseudolines into red and blue pseudolines we make another assumption: there should be a vertical line ℓ_x crossed by all pseudolines, such that left of their crossing with ℓ_x the pseudolines are x-monotone.[2] Let ℓ_0 be a vertical line left of ℓ_x which has no vertex of the arrangement to its left. Label the pseudolines from top to bottom in the order of their intersection with ℓ_0. This labeling is an analog to the ordering of a set of lines by increasing slope. Given n with $1 \leq n < N$, we color the first n pseudolines $L_1, \ldots, L_n$ red (drawn as solid curves throughout the paper), and the remaining $N - n$ pseudolines $L^{n+1}, \ldots, L^N$ blue (dashed).

[2] It is common to assume that the pseudolines of an arrangement are x-monotone throughout the range; in view of our proof of Theorem 3.1, we only ask for this weaker condition.

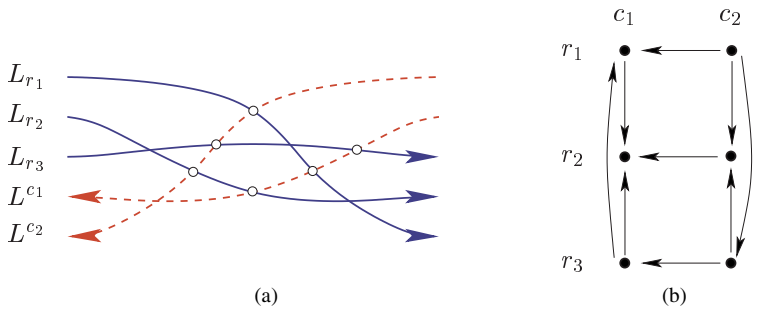

Fig. 3. A (sub)arrangement of three red and two blue pseudolines along with the corresponding (sub)grid orientation.

Consider the bottom cell of the arrangement $\mathcal{A}$ (the cell containing the portion of ℓ_0 below all pseudolines). The lines appearing on the boundary of this cell (the *lower envelope*) include L^N and L_1 and are ordered by decreasing labels from left to right. Therefore, there is a unique red–blue intersection ξ at the border of this cell—this is the "lowest" red–blue vertex of the arrangement, see Fig. 3(a).

A red line induces a total ordering (from left to right) of its intersections with blue lines. Symmetrically, a blue line induces a total ordering (from right to left) of its intersections with red lines. If there are n red lines and m blue lines, this yields an orientation of the (n, m)-grid. Clearly, subarrangements lead to subgrid orientations. The construction is illustrated in Fig. 3. When p, q are two red–blue intersection points of the arrangement, we write $p \to q$ whenever there is a directed path from the grid vertex corresponding to p to the one corresponding to q.

Lemma 2.2. *The orientation of the (n, m)-grid induced by an arrangement $\mathcal{A}$ of $N = n + m$ pseudolines is admissible.*

Proof. Acyclicity follows from the fact that there is a *sweep* of $\mathcal{A}$ from top to bottom, starting with a pseudoline above all vertices of the arrangement, its intersection with the pseudolines given by the order $L_n, \ldots, L_1, L^{n+m}, \ldots, L^{n+1}$. While the sweepline moves towards the bottom of the arrangement, meeting one vertex at a time, its intersection with any red line advances to the right, while blue lines are swept from right to left. This means that the ordering of the vertices as they are encountered during the sweep topologically sorts the grid orientation [8].

The unique sink property of subgrids is a consequence of the following observation: Let ξ be the unique red–blue intersection of $\mathcal{A}$ incident to the lower envelope. If p is any other red–blue intersection, then we have $p \to \xi$, or there is a red–blue intersection point q with $p \to q \to \xi$, see Fig. 4. In particular p cannot be a sink, hence, ξ is the unique sink of the grid. This argument carries over to subarrangements and the corresponding subgrids.

It remains to show that there is no induced double-twist. To yield a double-twist, there have to be two red lines, say, which both cross three blue lines L^i, L^j, L^k in this order. In addition, the order in which the red lines are crossed by L^i, L^j, L^k changes *twice* as we go through them in this order. This is impossible, see Fig. 5. □

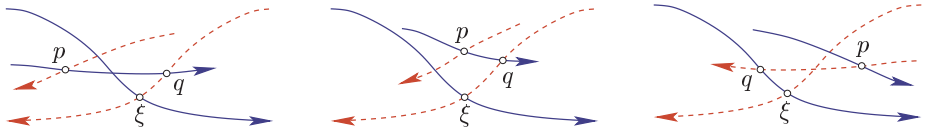

Fig. 4. Proof of Lemma 2.2: a path of length 3 from p to ξ.

In Section 3 we prove that the converse of Lemma 2.2 also holds: every admissible grid orientation is induced by a red–blue arrangement of pseudolines.

We conclude this section with some remarks on pseudoconfigurations of points and their relation to arrangements of pseudolines.

A *uniform pseudoconfiguration of points* of rank 3 is a pair $(\mathcal{A}, S)$ where S is a planar point set of size N, and $\mathcal{A}$ is an arrangement of $\binom{N}{2}$ pseudolines $\lambda_{ss'}$ through all pairs of points $s, s' \in S$. Uniform means that no three points in S lie on the same pseudoline. Let ℓ be a vertical line which is disjoint from S. We refer to such a planar pseudoconfiguration as the configuration $(\mathcal{A}, S, \ell)$, see Fig. 6 for an example.

Obviously, this is a generalization of *One line and N points*. The notions of ℓ-edge, below(e) as well as pivot(e, s) extend in a canonical fashion, so every instance $(\mathcal{A}, S, \ell)$ induces a grid orientation as before. By known duality results for arrangements of pseudolines, the concepts of red–blue arrangements of pseudolines and of pseudoconfigurations of points are completely equivalent [13], [2]. In particular, we obtain admissible grid orientations in both scenarios.

3. Pseudorealizability of Admissible Grid Orientations

Our findings so far can be summarized as follows. We have encountered six different settings which give rise to grid orientations. Three of them are geometric and mutually equivalent—they characterize the realizable case:

$$\left\{\begin{array}{l}(d, d+2)\text{-polytopes (with a linear function)}\\ \text{configurations of points (with a vertical line)}\\ \text{arrangements of lines (two-colored by slope)}\end{array}\right\}$$
$$\Updownarrow$$
realizable grid orientations

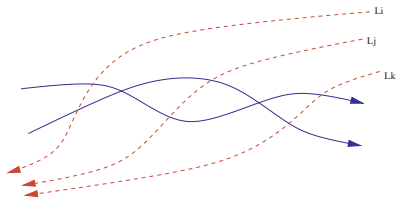

Fig. 5. Proof of Lemma 2.2: there is no induced double-twist.

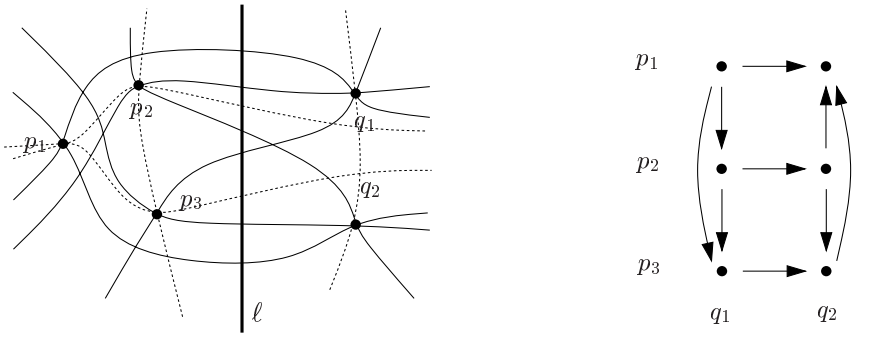

Fig. 6. One line and a pseudoconfiguration of N points, along with its induced admissible grid orientation.

The remaining three are also equivalent to each other and induce admissible orientations:

$$\left\{\begin{array}{c}(d, d+2)\text{-polytopes (with a Holt–Klee function)}\\ \text{pseudoconfigurations of points (with a vertical line)}\\ \text{arrangements of pseudolines (two-colored by line at infinity)}\end{array}\right\}$$
$$\Downarrow$$
$$\text{admissible grid orientations}$$

The goal of this section is to prove that the latter implication can be replaced by an equivalence. This means, arrangements of pseudolines are a universal model for admissible grid orientations.

Theorem 3.1. *Any admissible grid orientation is induced by a red–blue arrangement of pseudolines.*

To prove this theorem, we first describe a simple construction that encodes any unique sink grid orientation (Definition 2.1) by a red–blue arrangement of curves. We then show that any two curves of different color intersect exactly once, and that we can read the grid orientation off the resulting intersection pattern in exactly the same way as we have done earlier for the case of a pseudoline arrangement (see Fig. 3). This result is of interest in itself and can be used to gain nontrivial insights about general unique sink grid orientations (see Section 4).

While two curves of the same color might intersect several times in the general case, admissible grid orientations are shown to lead to arrangements in which any pair of curves induces *at most one* crossing—we get the desired arrangement of pseudolines.

Here is the construction. Given an (n, m)-grid orientation G, we identify the rows of G with the numbers $R = \{0, \ldots, n-1\}$ and the columns with $C = \{0, \ldots, m-1\}$. The set $R \times C$ then corresponds to the set of vertices of G.

Definition 3.1. Let $(r, c) \in R \times C$ be a vertex and define

$$\begin{aligned}\mathrm{Out}^{\downarrow}_G(r, c) &:= \{r' \in R \mid (r, c) \to_G (r', c)\},\\ \mathrm{Out}^{\rightarrow}_G(r, c) &:= \{c' \in C \mid (r, c) \to_G (r, c')\}\end{aligned}$$

as the sets of rows and columns to which (r, c) has outgoing edges. The mapping

$$\begin{aligned} \mathrm{ri}_G\colon\ R \times C &\to R \times C, \\ (r, c) &\mapsto (|\operatorname{Out}_G^{\downarrow}(r, c)|, |\operatorname{Out}_G^{\rightarrow}(r, c)|) \end{aligned}$$

is the *refined index* of G.

If G is a unique sink orientation, any row and any column is linearly ordered by G. It follows that for any pair $r \in R$, $\gamma \in C$, there is a unique column c with $\mathrm{ri}_G(r, c) = (*, \gamma)$. Similarly, for any pair $c \in C$, $\rho \in R$, we find $r \in R$ with $\mathrm{ri}_G(r, c) = (\rho, *)$.

To obtain the arrangement of curves representing G, every row is mapped to a red y-monotone polygonal path, traversing the construction from top to bottom. Columns become blue x-monotone paths, passing from right to left. The path for row r is the y-monotone path L_r interpolating the point set

$$\{\mathrm{ri}_G(r, c) \in \mathbb{R}^2 \mid c \in C\}$$

with straight line segments as connecting pieces. Beyond the interpolated points, L_r is extended vertically. By the previous observation, the point set contains exactly one point for every y-coordinate γ in C. Analogously, the path for column c is the x-monotone path L^c interpolating

$$\{\mathrm{ri}_G(r, c) \in \mathbb{R}^2 \mid r \in R\}$$

with horizontal extensions.

Figure 7(a) depicts a unique sink grid orientation (transitive arrows omitted), along with refined index values for row 0 and column 3. In Fig. 7(b) the resulting curves L_0 and L^3 are shown. Figure 7(c) shows the complete arrangement of four red and four blue curves.

Our next goal is to prove that any red curve crosses any blue curve exactly once, if G is a unique sink orientation. The following statement (which is known to hold in a more general framework [10, Lemma 2.14]) is the key ingredient. We provide an independent proof from which we also derive an important corollary.

Lemma 3.1. *If G is a unique sink grid orientation, then the refined index ri_G is a bijection.*

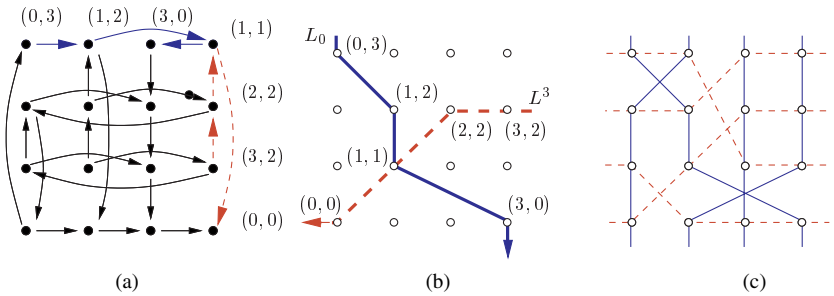

Fig. 7. Constructing curves from grid rows and columns.

Proof. It suffices to show that any pair $(\rho, \gamma) \in R \times C$ appears as a refined index value. We first prove this for pairs of the form $(\rho, 0)$, where the case $\rho = 0$ follows from the existence of a global sink. For $\rho > 0$, we inductively assume that there are vertices (r_i, c_i) with $\mathrm{ri}_G(r_i, c_i) = (i, 0)$, for all i with $0 \le i < \rho$. We actually work with the following stronger induction hypothesis: $\mathrm{Out}^{\downarrow}_G(r_i, c_i) = \{r_0, \ldots, r_{i-1}\}$. Now delete rows $r_0, \ldots, r_{\rho-1}$ from G and consider the unique sink (r_ρ, c_ρ) of the remaining subgrid G'. We want to prove that $r_i \in \mathrm{Out}^{\downarrow}_G(r_\rho, c_\rho)$, for $0 \le i < \rho$. If $c_\rho = c_i$, this immediately follows from the (stronger) induction hypothesis, otherwise it is implied by the unique sink property of the $(2, 2)$-subgrid spanned by the two row sinks (r_ρ, c_ρ) and (r_i, c_i). Because (r_ρ, c_ρ) has no outgoing edges to rows in G', $\mathrm{Out}^{\downarrow}_G(r_\rho, c_\rho) = \{r_0, \ldots, r_{\rho-1}\}$ follows.

To prove the occurrence of refined index (ρ, γ), we proceed in a similar fashion, now using induction on γ. The case $\gamma = 0$ has just been established, and for $\gamma > 0$, we inductively assume that there are vertices (r_j, c_j) with $\mathrm{ri}_G(r_j, c_j) = (\rho, j)$ and $\mathrm{Out}^{\rightarrow}_G(r_j, c_j) = \{c_0, \ldots, c_{j-1}\}$, $0 \le j < \gamma$. Delete columns $c_0, \ldots, c_{\gamma-1}$ from G to obtain G' and choose (r_γ, c_γ) with $\mathrm{ri}_{G'}(r_\gamma, c_\gamma) = (\rho, 0)$. As before, we need to prove $c_j \in \mathrm{Out}^{\rightarrow}_G(r_\gamma, c_\gamma)$ for $0 \le j < \gamma$ to arrive at the desired conclusion $\mathrm{Out}^{\rightarrow}_G(r_\gamma, c_\gamma) = \{c_0, \ldots, c_{\gamma-1}\}$.

If $r_j = r_\gamma$, this follows from the inductive hypothesis. Otherwise, suppose for a contradiction that $c_j \notin \mathrm{Out}^{\rightarrow}_G(r_\gamma, c_\gamma)$. Assume that $r_j \in \mathrm{Out}^{\downarrow}_G(r_\gamma, c_\gamma)$; the other case is symmetric. Then the situation is as in Fig. 8 (solid edges). By the unique sink property of $(2, 2)$-subgrids, we must have $r_\gamma \notin \mathrm{Out}^{\downarrow}_G(r_j, c_j)$ (edge 1 in Fig. 8). Now, because $|\mathrm{Out}^{\downarrow}_G(r_\gamma, c_\gamma)| = |\mathrm{Out}^{\downarrow}_G(r_j, c_j)| = \rho$, there must be some index $r' \in \mathrm{Out}^{\downarrow}_G(r_j, c_j) \setminus \mathrm{Out}^{\downarrow}_G(r_\gamma, c_\gamma)$, with $r' \ne r_j, r_\gamma$ (edges 2, 3). Acyclicity of columns c_j, c_γ implies the orientations of edges 4, 5. The contradiction arises, because both orientations of edge 6 lead to a $(2, 2)$-subgrid violating the unique sink property. □

The induction hypothesis in the second part of the proof yields part (i) of the following corollary, and the argument around Fig. 8 gives (ii).

Corollary 3.1.

(i) *Fix $\rho \in R$ and $\gamma, \gamma' \in C$ with $\gamma < \gamma'$. Let c, c' be the columns containing the vertices p, p' with refined indices (ρ, γ) and (ρ, γ'), respectively (note that*

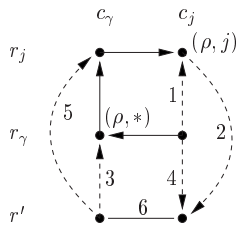

Fig. 8. Proof of Lemma 3.1.

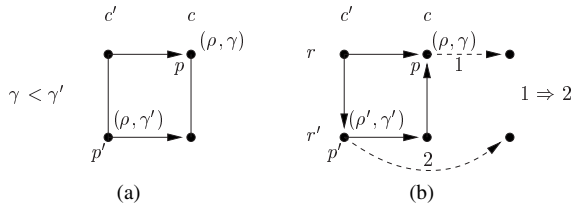

Fig. 9. Illustration of Corollary 3.1.

$c \neq c'$). *Then*

$$c \in \mathrm{Out}^{\rightarrow}_G(p'), \qquad c' \notin \mathrm{Out}^{\rightarrow}_G(p), \tag{4}$$

see Fig. 9(a).

(ii) *Fix vertices* $p = (r, c)$, $p' = (r', c')$ *in columns* $c \neq c'$ *with refined index values* (ρ, γ) *and* (ρ', γ'), *respectively, and such that* (4) *holds. If*

$$|(\mathrm{Out}^{\downarrow}_G(p) \cup \mathrm{Out}^{\downarrow}_G(p')) \cap \{r, r'\}| \in \{0, 2\},$$

then

(a) $\mathrm{Out}^{\rightarrow}_G(p) \cup \{c\} \subseteq \mathrm{Out}^{\rightarrow}_G(p')$, *in particular*

(b) $\gamma < \gamma'$,

see Fig. 9(b).

Returning to the construction of the red curves L_r, $r \in R$, and the blue curves L^c, $c \in C$ (see Fig. 7), we are now prepared to prove:

Theorem 3.2. *Any red curve* L_r *crosses any blue curve* L^c *exactly once.*

This implies that G can be read off the arrangement in the desired way (see Fig. 3): L_r crosses L^c at the point $\mathrm{ri}_G(r, c)$; hence, L_r crosses L^c before $L^{c'}$ (equivalently, at a higher y-coordinate) if and only if (r, c) has more outgoing edges within row r than (r, c'). This is the case if and only if $(r, c) \rightarrow_G (r, c')$. Symmetrically, L^c crosses L_r before $L_{r'}$ if and only if $(r, c) \rightarrow_G (r', c)$.

Proof. If we can show that $\mathrm{ri}_G(r, c)$ is the only point of intersection of L_r and L^c, the two curves must cross there, because L_r passes from top to bottom, while L^c passes from right to left. The refined index bijection shows that any grid point is used by exactly one red and exactly one blue curve. This already implies that L_r and L^c can have no grid point other than $\mathrm{ri}_G(r, c)$ in common.

Now suppose there is a point not in $R \times C$ where L_r and L^c intersect. We choose q to be such an intersection point with the smallest y-coordinate. Suppose q is on the segment of L_r between the horizontal lines $y = j + 1$ and $y = j$, and on the segment of L^c between the vertical lines $x = i + 1$ and $x = i$, see Fig. 10(a).

Let S be the square $[i, i + 1] \times [j, j + 1] \subseteq \mathbb{R}^2$. Since the order of the red curves along any line $y = t \in \mathbb{N}$ corresponds to a permutation of R, it follows that if L_r is

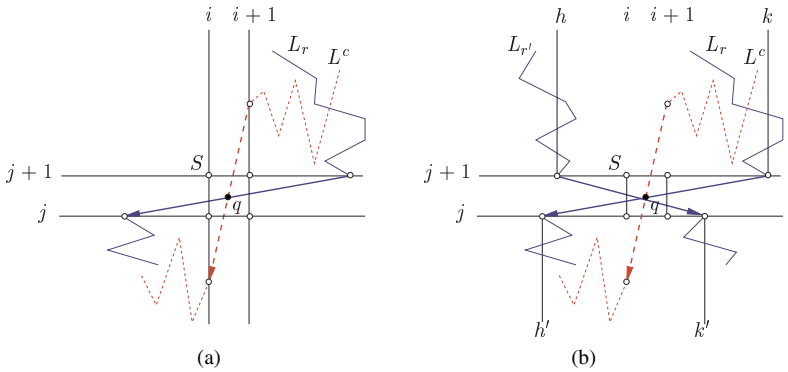

Fig. 10. Proof of Theorem 3.2: a red–blue intersection point which is not in $R \times C$.

crossing S from left to right, then there has to be a curve $L_{r'}$ crossing S from right to left. Similarly, if L^c is passing through S from bottom to top, then there is another curve $L^{c'}$ moving from top to bottom through S.

Therefore, we may assume that L_r and L^c behave as shown in the figure: L_r is right-to-left and L_c is top-to-bottom with respect to S. L_r is moving down with each segment and will leave the construction on the bottom, while L^c is moving left and leaves to the left. Hence, the two curves have to intersect again at a point below $y = j$. Due to the choice of q, this intersection is at a point of $R \times C$.

Recall that there is a red curve $L_{r'}$ crossing through the square S from left to right, see Fig. 10(b). Then there are indices $h, h' \leq i < k, k'$ such that row r contains vertices with refined indices $(k, j+1)$ and (h', j), while the refined index values $(h, j+1)$ and (k', j) appear on r'. With the roles of R and C interchanged, we now apply Corollary 3.1(i) to the pair of vertices with refined indices (h', j) (row r) and (k', j) (row r') to deduce that the situation at the vertex p with $\mathrm{ri}_G(p) = (h', j)$ is as in Fig. 11(a). A similar argument involving refined indices $(h, j+1)$ and $(k, j+1)$ yields the indicated edge orientation at the vertex p' with refined index $\mathrm{ri}_G(p') = (h, j+1)$. Now, part (b) of Corollary 3.1(ii) implies that the orientations of the two other solid edges in Fig. 11(b) have to be as depicted in the right part: assuming they both go from right to left (the only alternative) yields $j + 1 < j$.

We already know that L_r meets L^c at a (unique) point in $R \times C$ below $y = j$, meaning that $\mathrm{ri}_G(r, c) = (*, g)$, $g < j$. This equivalently means $c \in \mathrm{Out}_G^{\rightarrow}(p)$, and using part (a) of Corollary 3.1(ii), we obtain $c \in \mathrm{Out}_G^{\rightarrow}(p')$, see Fig. 11(b). It follows that $L_{r'}$ intersects

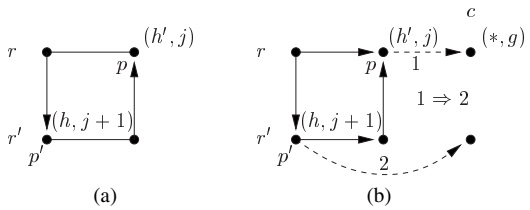

Fig. 11. Application of Corollary 3.1.

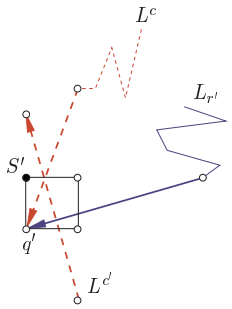

Fig. 12. Proof of Theorem 3.2: the crossing of $L_{r'}$ and $L^{c'}$ yields a contradiction.

L^c at a point $q' \in R \times C$ below $y = j + 1$. By construction of $L_{r'}$ and L^c, q' is even below $y = j$, and by the choice of q, q' is the first intersection point of $L_{r'}$ and L^c below $y = j$.

The situation at q' is therefore as shown in Fig. 12: L^c is approaching $L_{r'}$ from the top, so L^c is top-to-bottom with respect to the square S' attached to q', while $L_{r'}$ is right-to-left. As before, we can argue that there is another blue line $L^{c'}$ crossing S' from bottom to top. This line $L^{c'}$ has a crossing with $L_{r'}$ inside S'. This crossing is not at a point of $R \times C$, and it has smaller y-coordinate than point q, a contradiction to the choice of q. □

To derive our main theorem, Theorem 3.1, the representation theorem for admissible grid orientations, it remains to prove that the arrangement of the L_r, $r \in R$, and the L^c, $c \in C$, is an arrangement of pseudolines whenever G is admissible. From this, it is easy to obtain the desired *simple* arrangement in which every pair of pseudolines crosses exactly once: we slightly perturb the grid points in $R \times C$ to remove crossings of higher multiplicity, and we add the missing red–red and blue–blue crossings outside of the frame determined by $R \times C$. To obtain a drawing in the spirit of Fig. 3, we may finally rotate the construction, see Fig. 13.

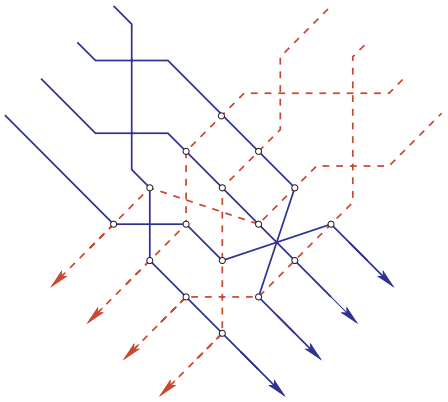

Fig. 13. A completed arrangement of pseudolines corresponding to the admissible grid orientation in Fig. 7.

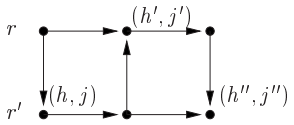

Fig. 14. Proof of Lemma 3.2: a double crossing between L_r and $L_{r'}$ yields a double-twist.

Lemma 3.2. *G is admissible if and only if every pair L_r, $L_{r'}$ of red curves has at most one crossing and every pair L^c, $L^{c'}$ of blue curves has at most one crossing.*

Proof. If we get a drawing without multiple crossings between curves, we can extend it to a simple arrangement of pseudolines as described above. By Lemma 2.2, G is admissible. Now assume we have a pair of curves, L_r, $L_{r'}$, say, with more than one crossing. Then there are y-coordinates $j > j' > j''$ such that L_r is without loss of generality right of $L_{r'}$ at heights j, j'' and left of $L_{r'}$ at height j'. Let h, h', h'' denote the smaller x-coordinates of the two curves at the respective heights. Like in the proof of Theorem 3.2, we can use Corollary 3.1 (i) to deduce the configuration of Fig. 14. This configuration is a double-twist and, therefore, G is not admissible. □

This lemma completes the proof of Theorem 3.1. We conclude this section with an example that illustrates the nonadmissible case. Figure 15(a) depicts the *snake*, a unique sink grid orientation with many induced double-twists. In the resulting arrangement (Fig. 15(b)), any pair of blue curves has the maximum possible number of intersections.

4. Grid Orientations Scrutinized

Let us reconsider what we have achieved so far: we have introduced admissible grid orientations (Definition 2.1) as combinatorial models for the graphs of $(d, d+2)$-polytopes, oriented by generic linear functions. We have seen that such polytope digraphs can equivalently be obtained from the two-dimensional scenarios of *One line and N points*, or,

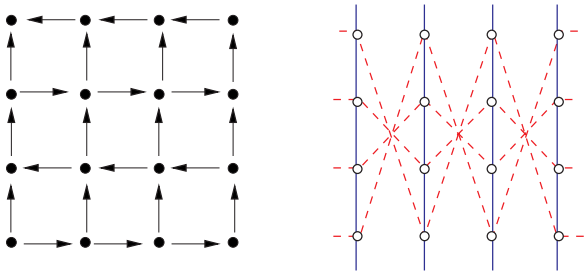

(a) (b)

Fig. 15. The nonadmissible case: the resulting arrangement will have multiple crossings between curves of the same color.

dually, from *No point and N lines*, see Section 2.2. As our main result, we have proved that admissible grid orientations are exactly the orientations arising from the generalized scenario of *No point and N pseudolines*, see Theorem 3.1. For this, we have shown something stronger: every unique sink grid orientation (Definition 2.1) can geometrically be realized by an arrangement of curves with favorable intersection properties (Theorem 3.2).

The goal of this section is to derive further structural and algorithmic results about admissible and general unique sink grid orientations, using the previous material. Some of them are already known, in which case the emphasis is on the ease with which they are obtained in the admissible case from our representation Theorem 3.1. We start with two simple properties we have already proved.

Proposition 4.1. *Let G be a unique sink grid orientation with no induced double-twist. Then*

(i) *G is acyclic, and*
(ii) *every subgrid of G has a unique source.*

Proof. (i) Going through the proof of Theorem 3.1 again reveals that acyclicity is never used. Therefore, we can associate to G an arrangement $\mathcal{A}$ of red and blue pseudolines inducing the orientation G. By Lemma 2.2, this orientation is acyclic.

(ii) Consider arrangement $\mathcal{A}$ from (i). Arguing as in the proof of Lemma 2.2, we see that the unique source of G corresponds to the unique red–blue crossing on the boundary of the *upper* envelope of $\mathcal{A}$. Subarrangements of $\mathcal{A}$ correspond to subgrids of G, meaning that all subgrids have the unique source property as well. □

The representation of unique sink grid orientation provided by Theorem 3.2 can be used to derive the (known) result that the proposition holds even under the presence of double-twists. This proof, however, would require some work, and the result is more easily established through other techniques [10].

The fact that unique sinks imply unique sources in all subgrids can be used to establish yet another characterization of admissible grid orientations that we have already anticipated in Section 2: admissible grid orientations are exactly the ones that satisfy the Holt–Klee condition [18].

Lemma 4.1. *Let G be a unique sink grid orientation. Then the following statements are equivalent*:

(i) *G is admissible.*
(ii) *Every (a, b)-subgrid has $a + b - 2$ vertex-disjoint paths from its unique source to its unique sink.*

Proof. The implication (ii) $\Rightarrow$ (i) is obvious, because (ii) excludes induced double-twists already for $(a, b) \in \{(2, 3), (3, 2)\}$. For the other direction, fix a subgrid G' of size $a \times b$ and let s be the unique source and t the unique sink of G'.

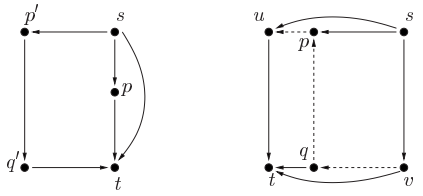

Fig. 16. Proof of Lemma 4.1: illustration of implication (i) ⇒ (ii).

Suppose first that s and t belong to the same column, see Fig. 16(a). There are $a-1$ paths from s to t in the column itself, one of the form $s \to t$, and $a-2$ of the form $s \to p \to t$, for $p \neq s, t$. By the unique sink property of $(2,2)$-subgrids, there is an additional path of the form $s \to p' \to q' \to t$ along each of the $b-1$ remaining columns. Together this gives $a+b-2$ vertex-disjoint paths from s to t. The same is true if s and t share their row.

Now suppose that s and t do not share a row or column, see Fig. 16(b). There are two paths of the form $s \to p \to t$, with p being one of the common neighbors u, v of s and t. For any other neighbor p of s, let q be the unique common neighbor of t and p not adjacent to s. We claim that $s \to p \to q \to t$, which gives us the remaining $(a-2)+(b-2)$ paths. Again, all $a+b-2$ paths are vertex-disjoint. To justify the claim, assume that $q \to p$ would hold. Then the unique sink property of $(2,2)$-subgrids implies that s, t, u, v, p, q induce a double-twist which is excluded by (i). □

Here is another interesting property of admissible grid orientations that can easily be deduced from our main result.

Proposition 4.2. *Let G be an admissible grid orientation, and let u, v be vertices such that there is a directed path between u and v in G. Then there is also a directed path between u and v that has no more than three edges.*

Proof. As before, we associate to G an arrangement of pseudolines, but now consider its dual pseudoconfiguration $(\mathcal{A}, S)$ of points, see Section 2.2 and Fig. 6. In this dual view, u and v are represented by two pseudolines ℓ_u, ℓ_v, each of which passes through a blue point to the left of the vertical line ℓ and a red point to the right. Without loss of generality, assume that ℓ_u and ℓ_v intersect to the right of ℓ. Because there is a directed path from u to v, ℓ intersects ℓ_u above ℓ_v. If ℓ_u and ℓ_v actually intersect in a point of S, u and v are neighbors in G. Otherwise, we will show that there can only be a directed path from u to v if there is some red point below ℓ_u in the "triangle" spanned by ℓ, ℓ_u, ℓ_v. This red point can be used to construct a directed path of length 2 or 3, see Fig. 17. To complete the argument, assume the indicated triangle contains no red point. In particular, $u \not\to v$. Consider a vertex u' such that $u \to u'$ in G. The replacement of ℓ_u by $\ell_{u'}$ leads to a situation in which the triangle either disappears ($\ell_{u'}$ intersects ℓ below ℓ_v), or becomes strictly smaller. Again, $u' \not\to v$ holds, and, inductively, it follows that there can be no directed path from u to v in G. □

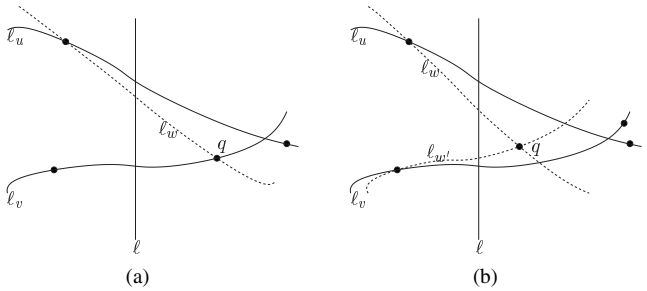

Fig. 17. A path $u \to w \to v$ of length 2 (a), and a path $u \to w \to w' \to v$ of length 3 (b) in the orientation G represented by $(\mathcal{A}, S)$.

Proposition 4.2 requires the absence of induced double-twists: there is a unique sink orientation of the (n, m)-grid where the shortest directed path between two specific vertices has $2\min(n, m) - 1$ edges, and this is indeed the maximum possible length of a shortest directed path [26].

The interpretation of admissible grid orientations in terms of pseudoconfigurations can also be used to establish the following result, a generalization of a theorem that was originally shown for *One line and N points* [11]. The arguments in the proof—based on above/below relations involving points and ℓ-edges—smoothly extend to the case of pseudoconfigurations.

Proposition 4.3. *Let G be an admissible orientation of the (n, m)-grid, and let v be any vertex. Consider the random walk that starts in v and at a generic vertex u proceeds along an outgoing edge chosen uniformly at random from all outgoing edges leaving u. Then for all v, the expected number of steps until the walk reaches the global sink of G is bounded by*

$$O(\log(n + 1)\log(m + 1)).$$

In the special case of *One line and N points*, or $(d, d + 2)$-polytopes with linear functions, this result establishes a tight upper bound for the expected performance of the simplex method with the *Random Edge* pivot rule [11]. Despite being the simplest and most natural randomized pivot rule, *Random Edge* has turned out to be very difficult to analyze. On general polytopes, only a small improvement over the trivial upper bound (number of vertices) is known [9], and no superpolynomial lower bounds exist. Only recently, it was shown that an abstract objective function of the n-cube graph exists on which *Random Edge* may require an expected superpolynomial number of steps to find the sink [22].

It is open whether the bound of Proposition 4.3 holds for general unique sink grid orientations. We already know that such orientations are acyclic, so the random walk will be simple. However, nontrivial upper bounds on its expected length are not known.

We continue with another statement whose validity for admissible grid orientations is easy to see from our representation theorem. Let G and G' be grid orientations. We say that G and G' are related by an *edge-flip* if the two orientations differ in the orientation of exactly one edge, see Fig. 18.

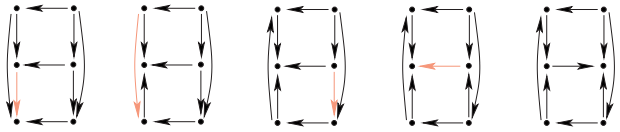

Fig. 18. A sequence of edge-flips.

Proposition 4.4. *Any two admissible orientations G and G' of the (n, m)-grid can be transformed into each other by a sequence of edge-flips, with all intermediate orientations being admissible as well.*

Proof. Let T be a triangular face of an arrangement $\mathcal{A}$. The local transformation replacing T by a triangle with the opposite orientation (see Fig. 19), is called a *triangular-flip* at triangle T.

Let $\mathcal{A}$ and $\mathcal{A}'$ be the arrangements of $m+n$ pseudolines associated with G and G'. It is known (see, e.g., [8]) that $\mathcal{A}$ can be transformed into $\mathcal{A}'$ by a sequence of triangular-flips. In our context we carry the red–blue partition of the pseudolines through. Consider two adjacent arrangements and their associated grid orientations. If the triangular-flip involves only pseudolines of one color, then the two grids are the same. If the flip involves two red and one blue pseudolines L_r, $L_{r'}$, and L_c, then the grids are related by an edge-flip at the edge connecting vertices (r, c) and (r', c). If the triangular flip involves two blue and one red pseudolines (as in Fig. 19), the situation is symmetric. □

The most interesting (and maybe surprising) new result of this section is that Proposition 4.4 generalizes to unique sink orientations. This is a nontrivial application of our general arrangement construction in Section 3.

Theorem 4.1. *Any two unique sink orientations G and G' of the (n, m)-grid can be transformed into each other by a sequence of edge-flips, with all intermediate orientations being unique sink orientations as well.*

In the proof of Proposition 4.4, we have noted that edge-flips in the grid correspond to triangular flips involving pseudolines of both colors in the arrangement. This fact remains true in the more general case of a unique sink orientation G and the associated arrangement $\mathcal{A}$ of multicrossing curves. In our proof of Theorem 4.1 we take advantage of this: we will show that the arrangements $\mathcal{A}$ and $\mathcal{A}'$ associated with G and G' can be transformed into each other by a sequence of triangular-flips. Caution! The previous

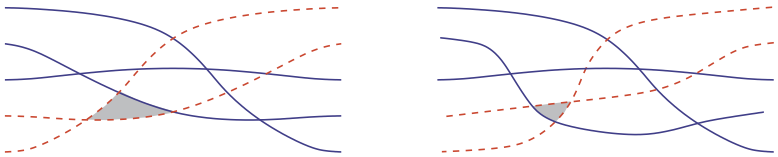

Fig. 19. A triangular-flip at the gray triangle of $\mathcal{A}$.

sentence cannnot be completely true. A triangular-flip keeps the total number of line crossings invariant but there are pairs $\mathcal{A}$, $\mathcal{A}'$ of arrangements associated with unique sink orientations of (n, m)-grids which have different numbers of crossings. Let a *pair-flip* be the addition or removal of a crossing between curves of the same color such that the crossing belongs to a triangle involving the line at infinity. The corrected version of the above statement is: arrangements $\mathcal{A}$ and $\mathcal{A}'$ associated with G and G' can be transformed into each other by a sequence of triangular-flips and pair-flips. Since pair-flips keep the associated grid orientation unaffected we confine our attention to triangular-flips.

Let $\mathcal{A}$ be the arrangement associated with G through the construction from Section 3. The crossings of blue and red curves occur at integer points (i, j) with $0 \leq i < n$ and $0 \leq j < m$. We say the arrangement is *combed* if the rectangle with corners $(0, 0)$ and $(n-1, m-1)$ only contains the $n \cdot m$ two-colored crossings, i.e., if all red and blue curves pass the grid as straight lines.

For $0 \leq i < n$ and $0 \leq j < m$ let $S_{(i,j)}$ be the open unit square $(i, i+1) \times (j, j+1)$. In the proof of Theorem 3.2 we have noted that a square which is intersected by a red and a blue curve would have to contain a red–blue crossing which is impossible. Therefore, each square $S_{(i,j)}$ of $\mathcal{A}$ is either empty, or it is a *red square*, i.e., a square intersected by a red curve, or a *blue square*, i.e., a square intersected by a blue curve.

We join sets of consecutive blue squares in the same column to *blue blocks* and sets of consecutive red squares in the same row to *red blocks*. Figure 20 shows an example.

The idea is to transform the arrangement $\mathcal{A}$ into a combed one by evacuating the blocks one by one. A blue block is *movable* if it can be shifted all the way to the left, without bumping into another block, and without ever touching another blue block. Two blocks *touch* if they are incident to a common square edge. Similarly, a red block is *movable* if it can be shifted all the way down without bumping into another block, and without ever touching another red block. In the example from Fig. 20, only the bottommost red block is movable. One of the main ingredients to the proof is the following lemma.

Lemma 4.2. *Let $\mathcal{A}$ be a two-colored arrangement of multicrossing curves corresponding to a unique sink grid orientation G. If $\mathcal{A}$ is not combed, then there is a movable block in $\mathcal{A}$.*

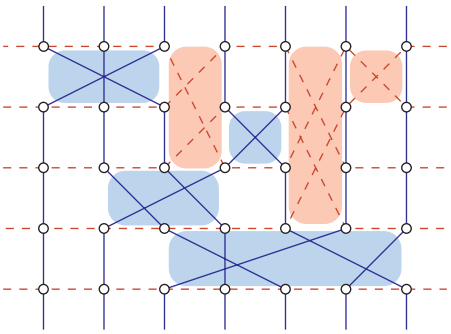

Fig. 20. An arrangement $\mathcal{A}$ with three (vertical) blue and four (horizontal) red blocks.

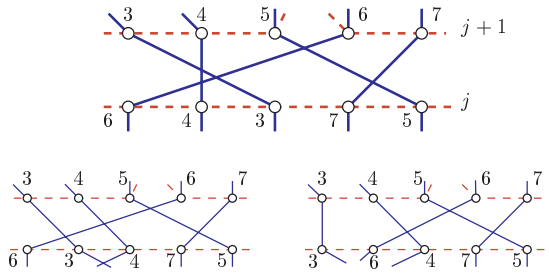

Fig. 21. The permutation induced by a movable block before (top) and after application of the sorting transpositions $(4, 3) \to (3, 4)$ and $(6, 3) \to (3, 6)$.

Before proving the lemma, we complete the proof of Theorem 4.1. Knowing that there is a movable block, we shift it out of the picture. Consider a movable red block first. The red line segments in this block induce a permutation between the points on the upper side of the block and the points on the lower side. The movable block from the example induces the permutation $(6, 4, 3, 7, 5)$, see Fig. 21. The permutation π induced by a block can be sorted by a sequence of adjacent transpositions. These transpositions can be realized through triangular-flips on $\mathcal{A}$. The flip corresponding to the transposition $(\pi(i), \pi(i+1)) \to (\pi(i+1), \pi(i))$ involves the red curves through points $(\pi(i), j)$ and $(\pi(i+1), j)$ and the blue curve connecting the two points. Note that the two points are really contained in the same blue curve, otherwise, a blue block would have to be incident to the edge connecting $(\pi(i), j)$ and $(\pi(i+1), j)$; because the red block is movable, the latter cannot happen. It may also be necessary to perform some triangular-flips involving only red curves before the described two-colored triangular-flip becomes admissible; this is the case for the transposition $(4, 3) \to (3, 4)$ in Fig. 21.

The sequence of transpositions described so far will push the movable red block one unit further down. Iterating this step, the block will eventually disappear from the picture, i.e., pushed across the x-axis out of the rectangle with corners $(0, 0)$ and $(n-1, m-1)$.

Since the situation with movable blue blocks is symmetric, we can summarize: as long as $\mathcal{A}$ is not combed, there is a movable block (Lemma 4.2) which can be evacuated from the picture. Since no new blocks are created during this procedure, we reach a state where $\mathcal{A}$ is combed. The remaining red–red and blue–blue crossings outside of the rectangle can be removed through pair-flips, resulting in an arrangement where all curves are proper lines.

Such an arrangement $\mathcal{A}$ is completely described by two permutations encoding the order of labels of the red and the blue curves on the x- and the y-axis. The order of adjacent red curves can be changed by pushing a crossing of them (a red block of size 1×1) through the picture. A sequence of such transpositions will bring the red curves into any desired order. By symmetry, the same is true for the order of blue curves. Hence, starting from arbitrary arrangements $\mathcal{A}$ and $\mathcal{A}'$, we can transform both into the *same* combed one by triangular flips and pair-flips. This completes the proof of Theorem 4.1.

It remains to supplement the proof of Lemma 4.2. To begin, we extend the concept of movability to points. A point $p = (p_x, p_y) \in \mathbb{R}^2$ is *movable* if its lower left quadrant

$L(p) = \{q\colon q_x \le p_x \text{ and } q_y \le p_y\}$ has empty intersection with the union of all blocks. If p is movable and $q \in L(p)$, then q is also movable. The *jump-line* of $\mathcal{A}$ is the boundary of the set of all movable points. If $\mathcal{A}$ is not combed, the jump-line consists of a sequence of vertical and horizontal line segments.

Think of the jump-line as being oriented from northwest to southeast. Then the first segment is vertical, with its x-coordinate determined by the leftmost block. The last segment is horizontal, and its y-coordinate is determined by the bottommost block.

Corners of the jump-line where a vertical segment is followed by a horizontal segment are called *concave corners*; it follows that there is at least one concave corner. It is easy to see that every concave corner of the jump-line is a corner of a block, and we let the concave corner inherit the color of this block.

If the first corner of the jump-line (which must be concave) is blue, then the blue block corresponding to this corner is movable: indeed, it is a leftmost block, so by shifting it to the left, it can neither bump into nor touch any other block. Similarly, if the last (concave) corner is red, the corresponding red block is movable. Now suppose the first corner is red and the last corner is blue. Then there is a pair of consecutive concave corners v and v' such that v is red and v' is blue. Let b and b' be the blocks corresponding to v and v'. Starting from v, the jump-line has a horizontal segment e along the lower border of b. If b does not extend beyond e, then the red block b is movable: by definition of the jump-line, it will not bump into another block when shifted down, and the only block it can possibly touch is the blue block b'.

If, on the other hand, b does extend beyond e, then the left border of b' cannot extend beyond the vertical segment e' ending in c'. In this case the blue block b' is movable. This shows how to find a movable block in every arrangement $\mathcal{A}$ which has blocks and hence a jump-line.

5. Realizability of Admissible Grids Orientations

In this final part we show that not all admissible grid orientations are "properly" realizable, and we provide a minimal such orientation. We will further see that the structure of $(2, m)$-grids is simple enough to guarantee realizability, for all positive integers m.

The fact that there are nonrealizable admissible grid orientations already follows from our main result via counting arguments, as follows. It is known that the number $T(N)$ of simple arrangements of $N = n + m$ pseudolines satisfies the inequalities

$$2^{cN^2} < T(N) < 2^{CN^2},$$

for suitable constants $c, C > 0$ with $C/c < 2$ [12]. Here, an arrangement is identified by a $(+/-)$-string of length $\binom{N}{3}$ that assigns an orientation to each triple of pseudolines. Because there are only N ways of specifying the red–blue partition, we may assume that $T(N)$ actually counts the red–blue arrangements. Different arrangements induce the same grid orientation G if the orientations of all two-colored triples of pseudolines agree. Therefore, any red–blue arrangement that induces G is determined by the orientations of the monochromatic triples, equivalently by its blue and its red subarrangement. If

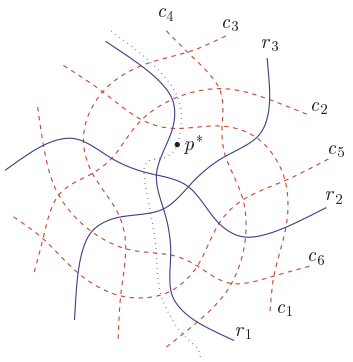

Fig. 22. Ringel's nonstretchable arrangement augmented by a designated line-at-infinity.

$n = m = N/2$, it follows that G is induced by at most

$$2^{2Cn^2} = 2^{(C/2)N^2} = 2^{(c-\delta)N^2}$$

arrangements, for some constant $\delta > 0$, meaning that at least $2^{\delta N^2}$ different admissible grid orientations can be obtained. On the other hand, there are only

$$2^{O(N \log N)}$$

red–blue arrangements of N lines [14], and via *No point and N lines*, only this many *realizable* grid orientations can be induced.

A Minimal Nonrealizable Admissible Grid Orientation. It is a well-known fact that every arrangement of fewer than nine pseudolines is *stretchable*, meaning that it is combinatorially equivalent to an arrangement of lines [4]. Based on the Pappus-configuration, Ringel [25] gave an example of a nonstretchable simple arrangement of nine pseudolines. Up to isomorphism in the sense of oriented matroids this is the only nonstretchable simple arrangement of nine pseudolines. Figure 22 shows a sketch of Ringel's arrangement. We have added a dotted pseudoline; taking this dotted line as the line at infinity and the point p^* on this line as the north pole, i.e., as the point corresponding to the vertical direction, the arrangement can be deformed as shown in Fig. 23(a). Figure 23(b) shows the corresponding admissible grid orientation.

Proposition 5.1. *The admissible orientation of the* $(3, 6)$ *grid shown in Fig.* 23(b) *is not realizable.*

Proof. The data provided by the grid only defines a partial arrangement. We have to verify that all completions of this partial arrangement correspond to the isomorphism class of Ringel's example.

The freedom left by the partial arrangement consists in the addition of those blue–blue and red–red crossings which are not determined by the grid. In the given case these are the crossings (c_1, c_4), (c_2, c_4), (c_3, c_4), (c_3, c_5), (c_3, c_6) and (r_1, r_2), (r_1, r_3), (r_2, r_3).

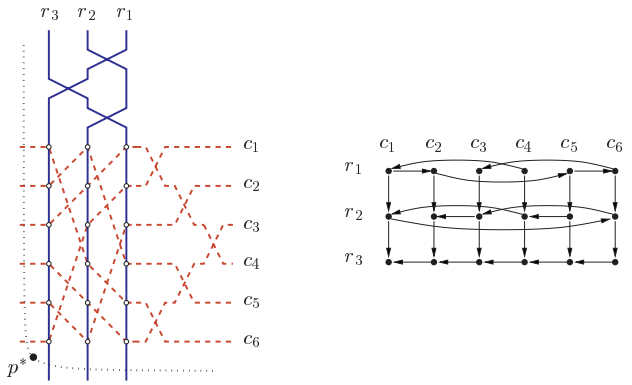

Fig. 23. Ringel's example redrawn and the corresponding grid orientation.

Consider the blue lines. The crossings on each of these lines come in two consecutive blocks: the block of crossings determined by the grid and the second block of "free" blue–blue crossings. The blue–blue crossings have to be arranged such that they permute the blue lines from $c_1, c_2, c_3, c_4, c_5, c_6$ to $c_4, c_1, c_2, c_5, c_6, c_3$. The first of these orders is given by the crossing order of blue lines on the line-at-infinity, the second is the crossing order on r_3. The grid together with these requirements completely determine the order of crossings on each of the blue lines.

With the red–red crossings there is slightly more choice. Assume that the line-at-infinity is such that on each red line the two leftmost crossings are with the other red lines, as in Fig. 23. There are two configurations for the red–red crossings which are related by the triangular-flip at the red triangle. In general such a flip will change the isomorphism class of the arrangement. In the given case, however, the two arrangements are mapped onto each other by an isomorphism: The vertical reflection in Fig. 23, followed by a change of the line-at-infinity. □

Realizability of Admissible $(2, m)$*-Grids.* With the previous proposition we gave an example of a nonrealizable admissible grid orientation with three rows. This example is row-minimal as we show next.

Theorem 5.1. *Let G be an admissible orientation of the $(2, m)$-grid, where m is any positive integer. Then G is realizable.*

Proof. By picture. We know from Theorem 3.1 that G can be represented by an arrangement of pseudolines, using a grid with m rows and two columns, see Fig. 24(a). Relevant to the fact that the arrangement represents G is the order in which crossings occur on the red lines, including the crossing of the red lines. This order of crossings can be preserved in a stretched arrangement, see Fig. 24(b). □

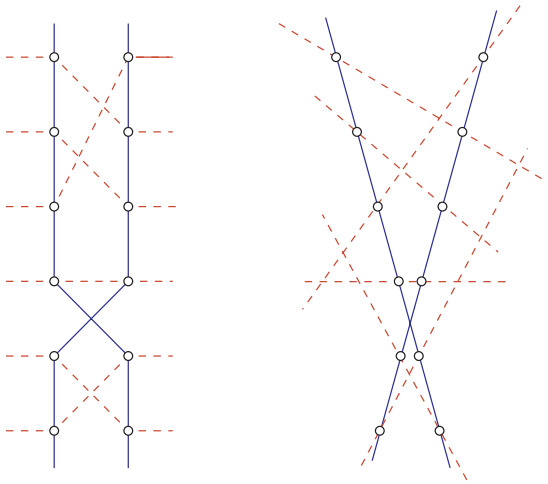

(a) (b)

Fig. 24. Stretching an arrangement of two red and m blue pseudolines.

6. Discussion

We have shown in this paper that the known combinatorial properties of linear functions on simple $(d, d+2)$-polytopes are captured by admissible grid orientations. For simple $(d, d+3)$-polytopes, the situation is completely unclear. Only some of these polytopes—the products of simplices—have graphs that are three-dimensional grids coming from the Cartesian product of three sets. Even in this restricted case, it is unknown how well Holt–Klee functions "simulate" linear functions. One positive result is that three-dimensional grid orientations induced by Holt–Klee functions can be characterized in terms of finitely many (actually, only three) forbidden subgrid orientations [10].

Note that on $(d, d+3)$-polytopes, the Holt–Klee condition does not rule out cyclic orientations, see Fig. 25 for the smallest such example with $d = 3$.

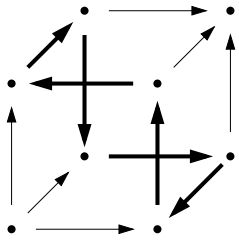

Fig. 25. The smallest example for a cyclic unique sink orientation of a $(d, d+3)$-polytope.

Completability

This article has its origins in an older manuscript by the second and the third authors. The title was *"Partial chirotopes and the digraph of* $(d, d+2)$*-polytopes"*. We give a brief description of the completability problem hidden in the notion of partial chirotopes, and of results and an open problem related to it.

A uniform rank 3 chirotope on a set X of n elements is a mapping from all triples of X to $\{+1, -1\}$ obeying certain properties (axioms). The completability problem is to decide whether a set of specified values on some triples can be completed to a full system of values that satisfies the axioms. Chirotopes are one of the many cryptomorphic axiomatizations for oriented matroids.

Simple arrangements of pseudolines are a universal model for uniform rank 3 chirotopes, i.e., for uniform oriented rank 3 matroids. Admissible grids, as discussed in this article, can be interpreted as partial chirotopes and our main representation result (Theorem 3.1) translates to a completability statement: *partial chirotopes given by admissible grids are completable*. The thesis of Tschirschnitz [26] contains a proof that in general the completability problem for partial uniform rank 3 chirotopes in NP-complete. Baier [3] shows that this result is already implied by NP-completeness proofs provided by Knuth in Section 6 of his monograph *Axioms and Hulls* [21].[3]

There still remain interesting completability problems in this area whose complexity has not been settled. We mention two of them:

Signotopes are a combinatorial encoding for simple Euclidean arrangements of pseudolines (see [7]). A 3-signotope on $[n] = \{1, \ldots, n\}$ is a mapping from all 3-element subsets of $[n]$ to $\{+1, -1\}$ obeying certain properties (generalized transitivity). The completability problem for partial 3-signotopes is open.

Hyperline sequences, or local sequences, are an alternative model for uniform rank 3 oriented matroids, equivalent to simple projective arrangements of pseudolines. The hyperline associated to line i is a circular sequence containing j or $\bar{j}$ for each $j \in X \backslash i$, where X is the set of all lines; this sequence records order and orientation of crossings along line i. A partial hyperline for i is a circular sequence containing j or $\bar{j}$ for each $j \in Y_i$, where Y_i is some subset of $X \backslash i$.

The completability problem for partial hyperline sequences is open.

References

1. Ilan Adler and Romesh Saigal. Long monotone paths in abstract polytopes. *Mathematics of Operations Research*, 1(1):89–95, 1976.
2. Pankaj K. Agarwal and Micha Sharir. Pseudo-line arrangements: duality, algorithms, and applications. In *Proc*. 13*th ACM–SIAM Symposium on Discrete Algorithms*, pages 800–809, 2002.
3. Patrick Baier. NP-completeness of partial chirotope extendibility. Manuscript (http://arxiv.org/abs/math.CO/0504430), 2005.

[3] Independently, this connection was pointed out by a referee of the earlier manuscript.

4. Anders Björner, Michel Las Vergnas, Neil White, Bernd Sturmfels, and Günter M. Ziegler. *Oriented Matroids*. Cambridge University Press, Cambridge, 1993.
5. Roswitha Blind and Peter Mani-Levitska. On puzzles and polytope isomorphisms. *Aequationes Mathematicae*, 34:287–297, 1987.
6. Mike Develin. On LP-orientations of cubes and crosspolytopes. Preprint, 2002.
7. Stefan Felsner. *Geometric Graphs and Arrangements*. Vieweg Verlag, Braunschweig, 2004.
8. Stefan Felsner and Helmut Weil. Sweeps, arrangements and signotopes. *Discrete Applied Mathematics*, 18:257–267, 2001.
9. Bernd Gärtner and Volker Kaibel. Two new bounds for the Random Edge pivot rule. Manuscript (http://arxiv.org/abs/math.CO/0502025), 2004.
10. Bernd Gärtner, Walter D. Morris, and Leonard Rüst. Unique sink orientations of grids. In *Proc*. 11*th Conference on Integer Programming and Combinatorial Optimization* (*IPCO*), pages 210–224. Volume 3509 of Lecture Notes in Computer Science. Springer-Verlag, Berlin, to appear.
11. Bernd Gärtner, József Solymosi, Falk Tschirschnitz, Pavel Valtr, and Emo Welzl. One line and n points. *Random Structures & Algorithms*, 23(4):453–471, 2003.
12. Jacob E. Goodman and Richard Pollack. Multidimensional sorting. *SIAM Journal on Computing*, 12(3):484–507, August 1983.
13. Jacob E. Goodman and Richard Pollack. Semispaces of configurations, cell complexes of arrangements. *Journal of Combinatorial Theory, Series A*, 37:257–293, 1984.
14. Jacob E. Goodman and Richard Pollack. Upper bounds for configurations and polytopes in $\Re^d$. *Discrete & Computational Geometry*, 1:219–227, 1986.
15. Branko Grünbaum. *Convex Polytopes*. Interscience, London, 1967.
16. Branko Grünbaum. *Convex Polytopes*, 2nd edition. Volume 221 of Graduate Texts in Mathematics. Springer-Verlag, Heidelberg, 2003.
17. Peter L. Hammer, Bruno Simeone, Thomas M. Liebling, and Dominique de Werra. From linear separability to unimodality: a hierarchy of pseudo-boolean functions. *SIAM Journal on Discrete Mathematics*, 1(2):174–184, 1988.
18. Fred B. Holt and Victor Klee. Many polytopes meeting the conjectured Hirsch bound. *Discrete & Computational Geometry*, 20:1–17, 1998.
19. Gil Kalai. A simple way to tell a simple polytope from its graph. *Journal of Combinatorial Theory*, 49:381–383, 1988.
20. Gil Kalai. A subexponential randomized simplex algorithm. In *Proc*. 24*th Annual ACM Symposium on Theory of Computing*, pages 475–482, 1992.
21. Donald E. Knuth. *Axioms and Hulls*. Volume 606 of Lecture Notes in Computer Science. Springer-Verlag, Berlin, 1992. `http://www-cs-faculty.stanford.edu/~knuth/aah.html`.
22. Jiři Matoušek and Tibor Szabó. RANDOM EDGE can be exponential on abstract cubes. In *Proc*. 45*th Annual IEEE Symposium on Foundations of Computer Science* (*FOCS*), pages 92–100, 2004.
23. Jed Mihalisin and Victor Klee. Convex and linear orientations of polytopal graphs. *Discrete & Computational Geometry*, 24:421–435, 2000.
24. Walter D. Morris, Jr. Distinguishing cube orientations arising from linear programs. Submitted, 2002.
25. Gerhard Ringel. Teilung der Ebene durch Geraden oder topologische Geraden. *Mathematische Zeitschrift*, 20:79–102, 1956.
26. Falk Tschirschnitz. LP-related Properties of Polytopes with Few Facets. Ph.D. thesis, ETH Zürich, 2003. `http://e-collection.ethbib.ethz.ch/show?type=diss&nr=15179`.
27. Doug Wiedemann. Unimodal set-functions. *Congressus Numerantium*, 50:165–169, 1985.
28. Kathy Williamson Hoke. Completely unimodal numberings of a simple polytope. *Discrete Applied Mathematics*, 20:69–81, 1988.

Received March 7, 2003, *and in revised form April* 26, 2005. *Online publication August* 3, 2005.