



Face monoid actions and tropical hyperplane arrangements

DOI:
[10.1142/S0218196717500461](https://doi.org/10.1142/S0218196717500461)

Document Version
Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):
Johnson, M., & Kambites, M. (2017). Face monoid actions and tropical hyperplane arrangements. *International Journal of Algebra and Computation*. <https://doi.org/10.1142/S0218196717500461>

Published in:
International Journal of Algebra and Computation

Citing this paper
Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights
Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy
If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.



FACE MONOID ACTIONS AND TROPICAL HYPERPLANE ARRANGEMENTS

MARIANNE JOHNSON¹ and MARK KAMBITES²

School of Mathematics, University of Manchester,
Manchester M13 9PL, England.

ABSTRACT. We study the combinatorics of tropical hyperplane arrangements, and their relationship to (classical) hyperplane face monoids. We show that the *refinement* operation on the faces of a tropical hyperplane arrangement, introduced by Ardila and Develin in their definition of a *tropical oriented matroid*, induces an action of the hyperplane face monoid of the classical braid arrangement on the arrangement, and hence on a number of interesting related structures. Along the way, we introduce a new characterization of the *types* (in the sense of Develin and Sturmfels) of points with respect to a tropical hyperplane arrangement, in terms of *partial bijections* which attain permanents of submatrices of a matrix which naturally encodes the arrangement.

1. INTRODUCTION

The set of faces of a (central) hyperplane arrangement, and more generally the set of covectors of an oriented matroid, provide an important source of examples in the class of finite monoids known as *left regular bands*. Such monoids have been the focus of much research interest, following the celebrated work of Bidigare, Hanlon and Rockmore [7] who showed that a number of well-known Markov chains, including the Tsetlin library and the riffle shuffle, are random walks on the faces of a hyperplane arrangement, and that the representation theory of the hyperplane face monoid could be used to analyze these Markov chains. This observation created substantial interest in random walks on hyperplane face monoids and related semigroups [5, 8, 11, 25] and the theory of left regular band walks more generally [12, 13, 28], and in various aspects of the representation theory of hyperplane face monoids [27, 29], general finite left regular bands [26] and \mathcal{R} -trivial monoids [4, 6], as well as the study of several related (quasi)varieties of semigroups and monoids [2, 22]. The representation theory of hyperplane face monoids is also closely connected to the Solomon descent algebra (see [13, 29, 30, 31] for further details) and hyperplane face monoids themselves also arise in connection with combinatorial Hopf algebras [1].

The study of the tropical complex generated by the columns of a real $n \times d$ matrix M (see Section 3.2) was initiated by Develin and Sturmfels [18]. It has been observed [3] that such a tropical complex can be viewed

¹Email Marianne.Johnson@maths.manchester.ac.uk.

²Email Mark.Kambites@manchester.ac.uk.

[†]*Key words:* hyperplane face monoids; tropical convexity; tropical oriented matroids; tropical matrix permanents. *MSC:* 15A80; 52B12; 20M30

as the analogue of the abstract simplicial complex of faces resulting from a real hyperplane arrangement, by associating to each column of M a min-plus linear form specifying a min-plus hyperplane (this will be explained in detail in Section 3).

In the case of real hyperplane arrangements, the corresponding face poset is a monoid with a natural geometrically (or combinatorially) defined product. There is no obvious analogous monoid structure on the face poset of a tropical hyperplane arrangement, but we shall show instead that the hyperplane face monoid \mathcal{P}_n of the classical *braid arrangement* (see Section 2.2) acts upon the set of faces determined by d min-plus tropical hyperplanes in \mathbb{R}^n .

The faces of the tropical complex generated by M are labelled by *types* which are the tropical analogue of the *sign sequences* of a real hyperplane arrangement. We identify the collection of all types with the face poset $\mathcal{F}(M)$ of the tropical complex — this is the analogue of the oriented matroid of covectors of a real hyperplane arrangement and hence provided the inspiration for the notion of a *tropical oriented matroid* introduced by Ardila and Develin [3]. We shall show that their definition induced a natural right action of the monoid \mathcal{P}_n on the tropical face poset $\mathcal{F}(M)$ and there is a geometric interpretation of this action when viewing the elements of $\mathcal{F}(M)$ as faces. By viewing $\mathcal{F}(M)$ instead as a poset of Boolean matrices (the types), this action is just the restriction of the action of \mathcal{P}_n on the set of all $n \times d$ Boolean matrices described in Section 2.3 below.

Without the geometric intuition of the action of \mathcal{P}_n it is not so easy to see why each element of \mathcal{P}_n should act upon $\mathcal{F}(M)$; indeed, one needs to be able to determine whether a given $n \times d$ Boolean matrix is in fact a type labelling a face of $\mathcal{F}(M)$. In this paper we give a purely combinatorial characterization of the types of $\mathcal{F}(M)$, and use this to give a combinatorial argument showing that \mathcal{P}_n acts on the set of all types. To this end, we introduce another combinatorial object $\mathcal{P}(M)$ called the *permanent structure* of M , and show that this provides an equivalent encoding of the combinatorial data of the tropical complex generated by M (Corollary 6.2). This may be of independent interest.

The paper is organised as follows. In Section 2 we recall some facts about real hyperplane arrangements and the hyperplane face monoid \mathcal{P}_n , including its natural action on the power set of $\{1, \dots, n\}$. In Section 3 we recall the necessary background on tropical hyperplane arrangements. In Section 4 we begin by making a distinction between Boolean matrices whose corresponding order theoretic conditions can be *satisfied* by a point and those which cannot. We show that a Boolean matrix is *satisfiable* if and only if every partial bijection contained within it (as a submatrix) is satisfiable. In Section 5 we show that the satisfiable partial bijections are precisely the max-plus permanent-attaining ones (and hence a matrix is satisfiable if and only if every partial bijection contained within it attains the permanent). In Section 6 we give a precise characterization of those matrices which correspond to types, in terms of permanent-attaining partial bijections. As an application, we obtain a combinatorial understanding of the action of \mathcal{P}_n on $\mathcal{F}(M)$.

Notation. We denote by $[n]$ the set $\{1, \dots, n\}$. A key object of study will be the set of d -tuples of subsets of $[n]$. It is convenient to identify each such tuple (S_1, \dots, S_d) with the $n \times d$ zero-one (Boolean) matrix M such that $M_{ij} = 1$ if and only if $i \in S_j$. We write $\mathbb{B}^{n \times d}$ for the set of all $n \times d$ Boolean matrices. We write $M_{i\star}$ and $M_{\star j}$ for the i th row and j th column respectively of a (Boolean or other) matrix M . For a Boolean matrix M we will also view each row $M_{i\star}$ (respectively, column $M_{\star j}$) as the corresponding subset of $[d]$ (respectively $[n]$) when convenient.

Notice that the *dual* n -tuple (T_1, \dots, T_n) of subsets of $[d]$, where $i \in T_j$ if and only if $j \in S_i$ corresponds to the transpose matrix, so with our identification, we have the natural notation $(S_1, \dots, S_d)^\top = (T_1, \dots, T_n)$.

There is an obvious partial order on $\mathbb{B}^{n \times d}$ given by setting $M \preceq N$ if and only if $M_{ij} \leq N_{ij}$ for all i and j . Equivalently in terms of tuples, $(S_1, \dots, S_d) \preceq (T_1, \dots, T_d)$ if and only if $S_i \subseteq T_i$ for each i .

By a *partial bijection* $\sigma : [d] \dashrightarrow [n]$ we mean a bijection $\sigma : J \rightarrow I$ between subsets $J \subseteq [d]$ and $I \subseteq [n]$. We may identify such a partial bijection with the matrix $\Sigma \in \mathbb{B}^{n \times d}$ defined by $\Sigma_{i,j} = 1$ if and only if $i = \sigma(j)$; thus the partial bijections correspond to the elements of $\mathbb{B}^{n \times d}$ having at most one non-zero entry in each row and in each column.

2. FACE MONOID ACTIONS

In this section we recall some necessary background on hyperplane face monoids (introduced in Section 2.1). Our main interest is in the hyperplane face monoid of the braid arrangement (Section 2.2) and a particular action of this monoid on the set of Boolean matrices (Section 2.3).

2.1. The face monoid of a hyperplane arrangement. Consider a finite collection $\mathcal{A} = \{\mathcal{H}_1^0, \dots, \mathcal{H}_d^0\}$ of (classical) *linear hyperplanes* containing the origin in \mathbb{R}^n , each of which may be written in the form

$$\mathcal{H}_i^0 = \{x \in \mathbb{R}^n : a_i \cdot x = 0\}$$

for some fixed choice of $a_i \in \mathbb{R}^n$. The complement of a hyperplane arrangement in \mathbb{R}^n is a collection of open subsets of \mathbb{R}^n called *chambers*. More generally, a *face* of the arrangement is defined to be a non-empty set which is an intersection of the form $\bigcap_{\mathcal{H}_i \in \mathcal{A}} \mathcal{H}_i^{\sigma_i}$, where $\sigma_i \in \{-, 0, +\}$ and

$$\mathcal{H}_i^+ = \{x \in \mathbb{R}^n : a_i \cdot x > 0\}, \quad \mathcal{H}_i^- = \{x \in \mathbb{R}^n : a_i \cdot x < 0\}.$$

Note that the sequence of signs (σ_i) therefore specifies on which ‘side’ of each hyperplane $\mathcal{H}_i^0 \in \mathcal{A}$ the given face lies. For each $x \in \mathbb{R}^n$ and each hyperplane \mathcal{H}_i^0 , we let $\sigma_i(x) \in \{-, 0, +\}$ denote the sign indicating the position of the point x relative to the hyperplane \mathcal{H}_i^0 , so that $x \in \mathcal{H}_i^{\sigma_i(x)}$. It is then clear that the map $\sigma : \mathbb{R}^n \rightarrow \{-, 0, +\}^{|\mathcal{A}|}$ defined by $\sigma(x) = (\sigma_1(x), \dots, \sigma_d(x))$ is constant on faces and that distinct faces have distinct images under σ . Thus we may identify the set of all faces \mathcal{F} with the set of sign sequences $\sigma(\mathbb{R}^n)$. Under this identification, the chambers are those elements of the image in which each component is non-zero, and the face Z corresponding to the common intersection of all hyperplanes in \mathcal{A} is identified with the all-zero sequence.

It is well known (see, for example, [11]) that the set of all faces \mathcal{F} of any finite linear hyperplane arrangement can be endowed with the structure of a monoid, whose identity element is Z , and whose product can be defined geometrically as follows. For faces F and G of a given hyperplane arrangement, choose points $f \in F$ and $g \in G$. The product $F * G$ is the very first face one encounters after leaving point f when walking in a straight line towards g (which could be the face F itself); this definition does not depend upon the particular points selected, and it can be shown that this product is associative. The resulting finite monoid is called the *hyperplane face monoid* of the arrangement and it is well-known and straightforward to check that each such monoid is a *left regular band* (that is, $F * F = F$ and $F * G * F = F * G$ for all $F, G \in \mathcal{F}$) and that each chamber is a *left zero* of this monoid (that is, for any chamber $C \in \mathcal{F}$, we have $C * F = C$ for all $F \in \mathcal{F}$). We refer the reader to [12] and [27] for general reference on properties of hyperplane face monoids.

The set of signs $L = \{-, 0, +\}$ can be considered as a left regular band consisting of an identity element 0 and two left zeroes, $-$ and $+$. It turns out that σ induces a monoid embedding of \mathcal{F} in the direct power $L^{|\mathcal{A}|}$ (see [13] for example). In other words, upon identifying \mathcal{F} with the set of possible sign sequences $\sigma(\mathbb{R}^n)$, for each $\mathcal{H}_i \in \mathcal{A}$ and all $F, G \in \mathcal{F}$ we have

$$(F * G)_i = \begin{cases} F_i & \text{if } F_i \neq 0 \\ G_i & \text{if } F_i = 0. \end{cases}$$

In fact, as observed in [11, §3E], the elements of $\mathcal{F} = \sigma(\mathbb{R}^n) \subseteq L^{|\mathcal{A}|}$ satisfy the covector axioms of an *oriented matroid* (we refer the reader to [9] for general reference on oriented matroids and also [23, Section 2.3.1] for a comparison between the terminology of hyperplane face monoids and oriented matroids). In [3, Definition 3.5] Ardila and Develin suggested a definition of *tropical oriented matroid*, based on their geometric understanding of both tropical hyperplane arrangements and the oriented matroids arising from (classical) hyperplane arrangements. We shall show that their definition induces an *action* of a certain classical hyperplane face monoid on the faces of a tropical hyperplane arrangement.

2.2. The face monoid of the braid arrangement. In the case of Coxeter arrangements, the geometric description of the product of faces $F * G$ described above can be viewed as the Tits projection of G on F , denoted $\text{proj}_F G$ [31]. One well-studied example is the linear hyperplane arrangement

$$\mathcal{H}_{i,j}^0 = \{x \in \mathbb{R}^n : x_i = x_j\} \text{ for } 1 \leq i < j \leq n,$$

known as the *braid arrangement* \mathcal{B}_n , consisting of $\binom{n}{2}$ hyperplanes in \mathbb{R}^n intersecting in the one-dimensional real vector space $\mathbb{R}(1, \dots, 1)$. The corresponding hyperplane face monoid has received considerable attention in connection with shuffling schemes, random walks and the Solomon descent algebra of the symmetric group [7, 11, 12, 13, 29].

It is clear that the braid arrangement yields $n!$ chambers of the form

$$C_\tau = \{(x_1, \dots, x_n) : x_{\tau(1)} > \dots > x_{\tau(n)}\}.$$

for τ in the symmetric group $\text{Sym}(n)$. Moreover, since a face of this arrangement is defined to be a non-empty intersection of the form $\cap \mathcal{H}_{i,j}^{\sigma_{i,j}}$, where $\sigma_{i,j} \in \{-, 0, +\}$,

$$\mathcal{H}_{i,j}^+ = \{x \in \mathbb{R}^n : x_i > x_j\}, \text{ and } \mathcal{H}_{i,j}^- = \{x \in \mathbb{R}^n : x_i < x_j\},$$

it is easy to see that the faces of the braid arrangement can be identified with ordered partitions of $[n]$: for example, the face

$$\{x \in \mathbb{R}^7 : x_1 = x_3 = x_4 > x_6 > x_2 = x_7 > x_5\}$$

is identified with the ordered partition $(\{1, 3, 4\}, \{6\}, \{2, 7\}, \{5\})$. Let \mathcal{P}_n denote the set of all ordered set-partitions of $[n]$. Viewing faces F and G of the braid arrangement as ordered partitions $F = (F_1, \dots, F_l)$ and $G = (G_1, \dots, G_r)$ it is known (see [13] or [29] for example) that their product $F * G \in \mathcal{P}_n$ may be written concretely as

$$F * G = (F_1 \cap G_1, \dots, F_1 \cap G_r, F_2 \cap G_1, \dots, F_2 \cap G_r, \dots, F_l \cap G_r)^\sharp, \quad (2.1)$$

where $(\cdot)^\sharp$ denotes the operation of deleting any empty sets. It is straightforward to check that the partition $(\{1, \dots, n\})$ acts as an identity on both sides, and that the chambers C_τ of the braid arrangement correspond to ordered partitions $(\{\tau(1)\}, \dots, \{\tau(n)\})$, each of which is a left zero of the monoid.

Since \mathcal{P}_n is a monoid it acts (by multiplication) on both the left and the right of \mathcal{P}_n itself. It is clear that each chamber is a fixed point of the *right action* of faces, since each chamber is a left zero element of the monoid. We will show that the right action of the monoid \mathcal{P}_n on chambers can be extended in a natural way to give a right action on a number of interesting combinatorial structures resulting from a tropical hyperplane arrangement in \mathbb{R}^n .

2.3. An action on Boolean matrices. We begin by showing that \mathcal{P}_n acts on the right of the power set of $[n]$ as follows: given any subset $I \subseteq [n]$ and any $F \in \mathcal{P}_n$ we define

$$I \circ F := \begin{cases} \emptyset & \text{if } I = \emptyset \\ I \cap F_j & \text{if } I \cap F_j \neq \emptyset \text{ and } I \cap F_k = \emptyset \text{ for all } k > j. \end{cases} \quad (2.2)$$

That is, $I \circ F$ is the right-most non-empty intersection of I with a block of the partition F . Note that this operation is well-defined, since if $I \neq \emptyset$, then I must have non-empty intersection with *some* component of the ordered set partition F of $[n]$. The fact that this does indeed give a right action on the power set of $[n]$ can be seen by considering the equivalence relation \sim on \mathcal{P}_n in which two partitions are equivalent if and only if they have the same right-most blocks (e.g. $(\{1, 2, 3\}, \{4, 5\}) \sim (\{3\}, \{1\}, \{2\}, \{4, 5\})$). The equivalence classes are clearly labelled by the non-empty subsets of $[n]$ (e.g. the two partitions of the previous example belong to the class $\{4, 5\}$) and it is easily checked that \sim is a right congruence on \mathcal{P}_n . The quotient action can therefore be viewed as a right action of \mathcal{P}_n upon the non-empty subsets of $[n]$. By including the empty set as a global fixed point, we arrive at the action described above.

Let $P = (P_1, \dots, P_l), Q = (Q_1, \dots, Q_k)$ be distinct elements of \mathcal{P}_n and suppose that $P_{l-i} \neq Q_{k-i}$ is the first place that these two tuples differ when reading from *right to left*. Since P_{l-i} is fixed by the action of P , but cannot be fixed by the action of Q , we see that the action is faithful.

Proposition 2.1. *The above operation gives a faithful right action of the face monoid \mathcal{P}_n on the power set of $[n]$. Moreover, $I \circ F \subseteq I$ for all $I \subseteq [n]$ and $F \in \mathcal{P}_n$.*

The last part of Proposition 2.1 should not be misinterpreted as saying that the action of \mathcal{P}_n is monotonically decreasing with respect to the containment order: if $I \subseteq J \subseteq [n]$ then we do *not* necessarily have $I \circ F \subseteq J \circ F$; what is true is that either $I \circ F \subseteq J \circ F$ or $I \circ F \cap J \circ F = \emptyset$. We recall that submonoids of the monoid of all ‘extensive’ (that is, non-increasing) functions acting on the right of $\{1, \dots, n\}$ are precisely the finite \mathcal{R} -trivial monoids (see for example, [24, Chapter 4, Theorem 3.6]), and so whilst it is natural to see a representation of the braid monoid by extensive maps, in general one should not expect the representation to be by order preserving functions.

For any d the action of \mathcal{P}_n on the power set of $[n]$ obviously extends (componentwise) to an action on the set of d -tuples of subsets of $[n]$ and hence (through the identification described in the Introduction) on the set $\mathbb{B}^{n \times d}$ of $n \times d$ Boolean matrices. It follows from Proposition 2.1 that $M \circ F \preceq M$ for all $M \in \mathbb{B}^{n \times d}$ and all $F \in \mathcal{P}_n$. In particular, the action of \mathcal{P}_n restricts to an action on any subset of $\mathbb{B}^{n \times d}$ which is downward-closed under the partial order.

3. TROPICAL HYPERPLANE ARRANGEMENTS

3.1. Tropical linear forms and hyperplanes. For $a, b \in \mathbb{R}$ write $a \oplus b := \max(a, b)$, $a \boxplus b := \min(a, b)$ and $a \otimes b := a + b$. The operations \oplus and \otimes give \mathbb{R} the structure of a *semiring* (without zero element), called the *max-plus semiring* (see [14, 16] for further reference). The min-plus semiring is defined similarly, and negation provides an isomorphism between these two structures, allowing one to easily interpret min-plus results in the max-plus setting and vice versa. The word *tropical* is used throughout the literature to describe a number of algebraic and geometric constructions involving either the max-plus or min-plus semiring. For example, it is clear that these operations extend naturally to give notions of addition and scalar multiplication of vectors, hence giving \mathbb{R}^n the structure of a (semi)module over the max-plus or min-plus semiring. A subset of \mathbb{R}^n closed under max-plus (or min-plus) operations will be referred to as a max-plus (respectively min-plus) submodule of \mathbb{R}^n ; the finitely generated submodules have been termed max-plus (min-plus) ‘tropical polytopes’ in the literature. The theory of tropical semimodules and convex sets has been developed by several authors, including Cohen, Gaubert and Quadrat [15], and Develin and Sturmfels [18].

For $n > 1$ (by analogy with the case of real hyperplanes) each $a \in \mathbb{R}^n$ can be used to define a max-plus or min-plus hyperplane, by considering the

linear forms

$$(a_1 \otimes x_1) \oplus (a_2 \otimes x_2) \oplus \cdots \oplus (a_n \otimes x_n) \quad (3.1)$$

$$(a_1 \otimes x_1) \boxplus (a_2 \otimes x_2) \boxplus \cdots \boxplus (a_n \otimes x_n), \quad (3.2)$$

in co-ordinate variables x_1, \dots, x_n . The max-plus (respectively, min-plus) hyperplane is then defined to be the set of all points $x \in \mathbb{R}^n$ for which the maximum (respectively minimum) is *attained at least twice* in expression (3.1) (respectively expression (3.2)) above. In contrast to hyperplanes in usual Euclidean space (which split the containing space into two half-spaces), a tropical hyperplane divides the containing space \mathbb{R}^n into n tropical “half”-spaces or *sectors*, according to which co-ordinate is maximal (respectively, minimal). It is clear that the point $x = -a$ lies on the hyperplane defined by the linear form specified by a , and that this point is adjacent to each of the sectors, in the sense that one can move to any given sector by making some small perturbation to the co-ordinates of a . We therefore define a (max-plus/min-plus) tropical hyperplane with *apex* a to be the (max-plus/min-plus) tropical hyperplane defined by the linear form specified by $-a$. Note that this tropical hyperplane does not have a unique apex, however the apexes are precisely all tropical scalings $\lambda \otimes a$.

There is a natural partial order on \mathbb{R}^n given by $x \leq y$ if and only if $x_i \leq y_i$ for all i . For $x, y \in \mathbb{R}^n$ we write

$$\langle x \mid y \rangle = \max\{\lambda \in \mathbb{R} : \lambda \otimes x \leq y\} = \min_k\{y_k - x_k\}. \quad (3.3)$$

This operation is a *residuation* operator (see [10]) and plays an important role in tropical mathematics (see for example [15]). We say that x *dominates* y in position i if $\langle x \mid y \rangle = y_i - x_i$. It is straight-forward to verify that domination is scale invariant, that is, if x dominates y in position i , then $\lambda \otimes x$ dominates $\mu \otimes y$ in position i for all $\lambda, \mu \in \mathbb{R}$. The set


$$\text{Dom}_i(x) = \left\{ y \in \mathbb{R}^n : y_i - x_i = \min_k\{y_k - x_k\} \right\} \quad (3.4)$$

consisting of all elements of \mathbb{R}^n dominated by x in position i can be seen to be a max-plus (and min-plus) submodule of \mathbb{R}^n as well as a closed convex set when considered as a subset of Euclidean space \mathbb{R}^n (see [20, Lemma 1.1] for example). If H is the min-plus hyperplane with apex $a \in \mathbb{R}^n$, then the set $\text{Dom}_i(a)$ is the closed sector consisting of points x such that $x_i - a_i \leq x_k - a_k$ for all k .

Since domination is scale invariant, it follows that we may identify each hyperplane H and the corresponding domination sets $\text{Dom}_i(a)$ with their respective images in the (classical) quotient vector space $\mathbb{R}^n/\mathbb{R}(1, \dots, 1)$, termed *tropical projective space*, as convenient. It will also be convenient (for the purpose of drawing pictures) to identify $\mathbb{R}^n/\mathbb{R}(1, \dots, 1)$ with \mathbb{R}^{n-1} , via the linear map

$$(v_1, \dots, v_n) \mapsto (v_1 - v_n, \dots, v_{n-1} - v_n). \quad (3.5)$$

Figure 1 illustrates a single min-plus tropical hyperplane, viewed in tropical projective space $\mathbb{R}^3/\mathbb{R}(1, 1, 1)$ which is identified with \mathbb{R}^2 (consisting of three half rays corresponding to the three ways in which the minimum can be attained at least twice), together with a labelling of the sectors following the usual co-ordinate axis conventions. Note that by setting $a = (0, 0, 0)$

FIGURE 1. The min-plus tropical hyperplane with apex $a \in \mathbb{R}^3$.

and extending the half-rays to give full real hyperplanes, one obtains the braid arrangement \mathcal{B}_3 .

3.2. The faces of a tropical hyperplane arrangement. Let M be a real $n \times d$ matrix. The columns of M determine a min-plus *tropical hyperplane arrangement* $\mathcal{H}_1, \dots, \mathcal{H}_d$ in \mathbb{R}^n , where \mathcal{H}_j is the min-plus tropical hyperplane with apex $M_{\star j}$, or in other words, the set of all $y \in \mathbb{R}^n$ such that the minimum over k of $y_k - M_{k,j}$ is attained twice.

The *type* of a point $x \in \mathbb{R}^n$ with respect to the hyperplane arrangement specified by M is an $n \times d$ Boolean matrix (or equivalently, an d -tuple of subsets of $[n]$) recording in which of the closed sectors relative to each \mathcal{H}_j the point x lies. Specifically, the type of x with respect to M is the Boolean matrix T where $T_{ij} = 1$ if and only if $x \in \text{Dom}_i(M_{\star j})$.

Types may be viewed as the tropical analogue of the sign sequences associated with a real hyperplane arrangement; the collection of all types (whilst not a matroid in the usual sense) is clearly the analogue of the oriented matroid of covectors of a real hyperplane arrangement, and motivated the notion of *tropical oriented matroids* studied in [3]. Develin and Sturmfels [18, Theorem 15] have shown that the min-plus hyperplane arrangement given by (the columns of) M induces a polyhedral cell complex structure on \mathbb{R}^n , called the *tropical complex* generated by M , whose face poset $\mathcal{F}(M)$ is naturally labelled by the types.

Since domination is scale invariant we note that the type of $\lambda \otimes x$ is the same as the type of x for all $\lambda \in \mathbb{R}$, $x \in \mathbb{R}^n$. It follows that we may identify the hyperplanes \mathcal{H}_j and the faces of the resulting tropical complex with their respective images in tropical projective space $\mathbb{R}^n/\mathbb{R}(1, \dots, 1)$, and hence by (3.5) \mathbb{R}^{n-1} , when convenient. Under this latter identification the faces of the max-plus tropical polytope are precisely the bounded faces in the tropical complex [18, Theorem 15]. By the *dimension* of a face we shall mean the dimension of the corresponding polyhedron in \mathbb{R}^{n-1} .

Example 3.1. Consider the matrix

$$M = \begin{pmatrix} -8 & 10 & 15 & 0 \\ 10 & 10 & 5 & -10 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

The min-plus tropical hyperplane arrangement specified by the columns of M is shown in Figure 2. As viewed in projective space, the resulting tropical

complex has 12 cells of dimension 2 (such as the bounded cell F and the unbounded cell E), 18 cells of dimension 1 (such as the bounded cell G , which is a face of both E and F) and 7 cells of dimension 0 (such as the point marked H , which is a face of each of the other labelled cells). The bounded cells have been shaded, and these cells constitute the max-plus polytope generated by the columns of M (see [18, Theorem 15]). We identify the cells E, F, G and H with their corresponding types:

$$\begin{aligned} E &= \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} = (\{2\}, \{1\}, \{1\}, \{1\}), \\ F &= \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = (\{2\}, \{1\}, \{1\}, \{3\}), \\ G &= \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = (\{2\}, \{1\}, \{1\}, \{1, 3\}), \\ H &= \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = (\{2\}, \{1, 2\}, \{1\}, \{1, 3\}). \end{aligned}$$




FIGURE 2. An arrangement of four min-plus tropical hyperplanes in \mathbb{R}^3 . The bounded cells of the resulting tropical complex are shaded. The union of the bounded regions is the max-plus polytope generated by the points M_i .

Let $\mathcal{F}(M)$ denote the set of all faces with respect to the hyperplane arrangement specified by M . Since the faces are in one-to-one correspondence with the types specified by this arrangement, we shall regard $\mathcal{F}(M)$ as a subset of $\mathbb{B}^{n \times d}$.

4. SATISFIABLE BOOLEAN MATRICES

Continuing with the notation of Section 3.2, we fix a real $n \times d$ matrix M whose columns $M_{\star j}$ specify a tropical hyperplane arrangement in \mathbb{R}^n ; all types in this section are considered with respect to the columns of M . More

generally, for each $S \in \mathbb{B}^{n \times d}$ we define

$$D_S = \bigcap_{S_{i,j} \neq 0} \text{Dom}_i(M_{\star j}) \quad (4.1)$$

and say that S is *satisfiable* (with respect to M) if $D_S \neq \emptyset$. That is, S is satisfiable if there exists $x \in \mathbb{R}^n$ such that $M_{\star j}$ dominates x in position i for all i, j such that $S_{i,j} = 1$; we say that such a point x *satisfies* S . (For example, with respect to the hyperplane arrangement given in Example 3.1 above, the (column) tuple $(\{1\}, \{2\}, \{1, 2\}, \{1, 3\})$ is not satisfiable. The tuple $(\{2\}, \{2\}, \{1, 2\}, \{1, 3\})$ is satisfiable, however it is not the type of any point.) Notice that if S is satisfiable with respect to M then the set D_S , which by definition is a non-empty intersection of Euclidean convex max-plus submodules of \mathbb{R}^n of the form (3.4), is itself a Euclidean convex max-plus submodule of \mathbb{R}^n .

It is clear from the definitions above that x satisfies S if and only if $S \preceq T$ where T is the type of x , and that the cells of the tropical complex are the relative interiors of the sets D_T indexed by types T . It follows that the set $\mathcal{S}(M)$ of all satisfiable (with respect to M) Boolean matrices corresponding to the hyperplane arrangement specified by M is a downward-closed subset of $\mathbb{B}^{n \times d}$ and hence (by our observation at the end of Section 2.3) \mathcal{P}_n acts on the right of $\mathcal{S}(M)$. The set $\mathcal{F}(M)$ of all types with respect to M is clearly not downward-closed, but it turns out that the action of \mathcal{P}_n also restricts to an action on this set. In order to study this action, we begin with some foundational results concerning satisfiability which will enable us to provide a combinatorial characterization of the types.

Recall from the introduction that we identify each partial bijection $\sigma : [d] \dashrightarrow [n]$ with the Boolean matrix $\Sigma \in \mathbb{B}^{n \times d}$ defined by $\Sigma_{i,j} = 1$ if and only if $i = \sigma(j)$. We say that $S \in \mathbb{B}^{n \times d}$ *contains* the partial bijection σ if $\Sigma \preceq S$.

Our main result in this section is the following ‘‘local’’ characterization of satisfiability.

Theorem 4.1. *Let $M \in \mathbb{R}^{n \times d}$ and $S \in \mathbb{B}^{n \times d}$. Then S is satisfiable with respect to M if and only if all of the partial bijections contained in S are satisfiable with respect to M .*

To prove this result we require a number of technical lemmas.

Lemma 4.2. *Let $M \in \mathbb{R}^{n \times d}$. An element $S = (S_{1\star}, \dots, S_{n\star})^\top \in \mathbb{B}^{n \times d}$ is satisfiable with respect to M if and only if $(A, S_{2\star}, \dots, S_{n\star})^\top$ is satisfiable with respect to M for all $A \subseteq S_{1\star}$ with $|A| \leq 1$.*

Proof. As we have already observed, $\mathcal{S}(M)$ is a downward-closed subset of $\mathbb{B}^{n \times d}$. It follows immediately that if S is satisfiable, then so too is $(A, S_{2\star}, \dots, S_{n\star})^\top$, where $A \subseteq S_{1\star}$.

Suppose now that $(A, S_{2\star}, \dots, S_{n\star})^\top$ is satisfiable for each $A \subseteq S_{1\star}$ with $|A| \leq 1$. Since domination is scale invariant we note that we may assume without loss of generality (by rescaling if necessary) that each column $M_{\star s}$ has 0 in the first position. For each $s \in S_{1\star}$ choose $x(s) \in \mathbb{R}^n$ satisfying $(\{s\}, S_{2\star}, \dots, S_{n\star})^\top$. We shall show that $x = \bigoplus_{s \in S_{1\star}} x(s)$ satisfies S . First

note that we may choose each $x(s)$ so that its first co-ordinate (and hence the first co-ordinate of x) is 0.

Since $x(s)$ satisfies $(\{s\}, S_{2\star}, \dots, S_{n\star})^\top$, we have in particular that $M_{\star s}$ dominates $x(s)$ in position 1:

$$0 = x(s)_1 - M_{1,s} = \min_k \{x(s)_k - M_{k,s}\},$$

from which it follows easily that $M_{\star s} \leq x(s)$. On the other hand, by definition of x we see that $x(s) \leq x$ and hence $M_{\star s} \leq x$, and from the latter we can also deduce that $M_{\star s}$ dominates x in position 1 for all $s \in S_1$:

$$x_1 - M_{1,s} = 0 = \min_k \{x_k - M_{k,s}\}.$$

Now let $S' = (\emptyset, S_{2\star}, \dots, S_{n\star})^\top$ and consider the max-plus submodule $D_{S'}$ defined via (4.1). Since each $x(s)$ satisfies S' , by definition it is contained in $D_{S'}$. Thus the max-plus sum $x = \bigoplus_{s \in S_{1\star}} x(s)$ must be contained in $D_{S'}$ too. Now since $M_{\star s}$ dominates x in position 1 for all $s \in S_1$ and x satisfies S' we conclude from (4.1) that x satisfies S . \square

It is clear that the analogues of Lemma 4.2 concerning positions $2, \dots, n$ also hold.

Corollary 4.3. *Let $M \in \mathbb{R}^{n \times d}$ and $S = (S_{1\star}, \dots, S_{n\star})^\top \in \mathbb{B}^{n \times d}$. Then $S \in \mathcal{S}(M)$ if and only if $A = (A_{1\star}, \dots, A_{n\star})^\top \in \mathcal{S}(M)$ for all $A \preceq S$ with $|A_{i\star}| \leq 1$.*

Proof. Since $\mathcal{S}(M)$ is downward-closed, the forward implication is immediate. On the other hand, if $(A_{1\star}, \dots, A_{n\star})^\top$ is satisfiable for all $A \preceq S$ with $|A_{i\star}| \leq 1$, then by repeated application of Lemma 4.2 (and its analogues for positions $2, \dots, n$) we see that S must be satisfiable too. \square

The following lemma combined with Corollary 4.3 will enable us to prove the desired result (Theorem 4.1).

Lemma 4.4. *Let $M \in \mathbb{R}^{n \times d}$ and let $S, T \in \mathcal{S}(M)$ be tuples of the form $S = (S_{1\star}, \dots, S_{n\star})^\top$, $T = (T_{1\star}, \dots, T_{n\star})^\top$ satisfying $i \in S_{j\star} \cap T_{k\star}$. Then there exists $U = (U_{1\star}, \dots, U_{n\star})^\top \in \mathcal{S}(M)$ such that:*

- (i) $i \in U_{j\star} \cap U_{k\star}$; and
- (ii) if $l \in S_{p\star} \cap T_{p\star}$, then $l \in U_{p\star}$.

Proof. Let $x, y \in \mathbb{R}^n$ be such that x satisfies S and y satisfies T . Thus by assumption $M_{\star i}$ dominates x in position j and y in position k . We want to construct $u \in \mathbb{R}^n$ such that $M_{\star i}$ dominates u in both positions j and k , and that in position p , u is dominated by any columns $M_{\star l}$ which dominate both x and y in that position.

First of all we note that we may assume without loss of generality that $M_{j,i} = M_{k,i} = 0$, as translating the i th column of M by the vector $-M_{j,i}e_j - M_{k,i}e_k$ clearly does not alter the combinatorial structure of the resulting face poset. Moreover, since domination is scale invariant, we may also assume without loss of generality that $x_j = y_k = 0$.

Recalling the definition of the residuation operator $\langle \cdot | \cdot \rangle$ given in (3.3), we define

$$u = \bigoplus_{l=1}^d \min(\langle M_{\star l} | x \rangle, \langle M_{\star l} | y \rangle) \otimes M_{\star l}$$

and claim that u satisfies the required conditions. Notice that the i th term in the above definition of u is $\min(0, 0) \otimes M_{\star i} = M_{\star i}$, giving $u \geq M_{\star i}$ and hence $u_t - M_{t,i} \geq 0$ for all $t \in [n]$. It is also easy to see that $x \geq u$ and $y \geq u$, since for example by (3.3) we have $x \geq \langle M_{\star l} | x \rangle \otimes M_{\star l}$ for all l , and hence

$$x \geq \bigoplus_{l=1}^d \langle M_{\star l} | x \rangle \otimes M_{\star l} \geq \bigoplus_{l=1}^d \min(\langle M_{\star l} | x \rangle, \langle M_{\star l} | y \rangle) \otimes M_{\star l} = u.$$

Now

$$0 = M_{j,i} \leq u_j \leq x_j = 0,$$

from which we deduce that $u_j = 0$. By an identical argument using y instead of x , we get $u_k = 0$. It now follows that

$$u_j - M_{j,i} = u_k - M_{k,i} = 0 \leq u_t - M_{t,i}, \text{ for all } t \in [n],$$

or in other words, that u is dominated by $M_{\star i}$ in both of these positions.

Now suppose that $M_{\star l}$ dominates both x and y in position p , that is, $\langle M_{\star l} | x \rangle = x_p - M_{p,l}$ and $\langle M_{\star l} | y \rangle = y_p - M_{p,l}$. But now looking at the l th term in the definition of u gives

$$\begin{aligned} u_p &\geq \min(\langle M_{\star l} | x \rangle, \langle M_{\star l} | y \rangle) \otimes M_{p,l} \\ &\geq \min(x_p - M_{p,l}, y_p - M_{p,l}) + M_{p,l} \\ &\geq \min(x_p, y_p). \end{aligned}$$

Since we know that $u \leq x$ and $u \leq y$, giving $u_p \leq \min(x_p, y_p)$, we must have $u_p = \min(x_p, y_p)$. We are now in position to show that $M_{\star l}$ dominates u in position p . Assume without loss of generality that $u_p = x_p \leq y_p$ (the argument being so far symmetric in x and y). Then because $M_{\star l}$ dominates both x and y in position p we see that

$$\langle M_{\star l} | x \rangle = x_p - M_{p,l} \leq y_p - M_{p,l} = \langle M_{\star l} | y \rangle.$$

Thus by the definition of u we see that

$$u \geq \min(\langle M_{\star l} | x \rangle, \langle M_{\star l} | y \rangle) \otimes M_{\star l} = \langle M_{\star l} | x \rangle \otimes M_{\star l}.$$

By (3.3) this means that $\langle M_{\star l} | u \rangle \geq \langle M_{\star l} | x \rangle$, giving

$$\min_k \{u_k - M_{k,l}\} = \langle M_{\star l} | u \rangle \geq \langle M_{\star l} | x \rangle = x_p - M_{p,l} = u_p - M_{p,l},$$

or in other words, $M_{\star l}$ dominates u in position p . \square

Proof of Theorem 4.1. Since $\mathcal{S}(M)$ is downward-closed, if S is satisfiable, then so too is each $A \preceq S$. Suppose then that all partial bijections contained in S are satisfiable. We shall show that each $A \preceq S$ with $|A_{i\star}| \leq 1$ must be satisfiable (by repeated application of Lemma 4.4). Then by Corollary 4.3 we can conclude that S is satisfiable.

Given $A \preceq S$ with $|A_{i\star}| \leq 1$ we consider the collection \mathcal{B}_A of all partial bijections contained within A . We apply Lemma 4.4 to all pairs (S, T) of *maximal* elements $S, T \in \mathcal{B}_A$ which differ in a single column, add the resulting satisfiable element U to \mathcal{B}_A , and continue in this fashion. We note that at each step of this procedure \mathcal{B}_A contains only satisfiable elements that are less than or equal to A in the order. Since there are finitely many such elements, it is clear that this process must terminate. The resulting set \mathcal{B}_A then has unique maximum element A . \square

5. SATISFIABLE PARTIAL BIJECTIONS AND THE MAX-PLUS PERMANENT

The *max-plus tropical permanent* of $X \in M_k(\mathbb{R})$ is the real number

$$\text{perm}(X) = \bigoplus_{\tau \in \text{Sym}(k)} X_{\tau(1),1} \otimes \cdots \otimes X_{\tau(k),k}$$

where $\text{Sym}(k)$ denotes the symmetric group on $[k]$. The matrix X is said to be *max-plus tropically singular* if the expression for the permanent above is attained by at least two different permutations, and *max-plus tropically non-singular* otherwise.

Let $\sigma : [d] \dashrightarrow [n]$ be a partial bijection with domain of definition $J \subseteq [d]$ and image $I \subseteq [n]$. We say that σ is *permanent-attaining* (with respect to M) if σ attains the maximum in the calculation of the max-plus permanent for the corresponding $k \times k$ submatrix of M , that is, if for every bijection $\tau : J \rightarrow I$ we have

$$\bigotimes_{j \in J} M_{\sigma(j),j} \geq \bigotimes_{j \in J} M_{\tau(j),j}.$$

We define the *permanent structure* $\mathcal{P}(M)$ of M to be the set of all permanent-attaining partial bijections.

Proposition 5.1. *The permanent-structure $\mathcal{P}(M)$ is a downward-closed subset of $\mathbb{B}^{n \times d}$.*

Proof. Let $\sigma \in \mathcal{P}(M)$ and suppose $\tau \preceq \sigma$. Clearly τ is a partial bijection. Suppose for a contradiction that it is not permanent-attaining. Choose a permanent-attaining partial bijection θ with the same domain of definition and image as τ . Since $\tau \preceq \sigma$, σ maps the domain of θ onto the range of θ ; this means that we can extend θ to a partial bijection θ' with the same domain and image as σ , by setting $\theta'(i) = \sigma(i)$ for all i not in the domain of θ . But now it is easily seen that θ' makes a greater contribution to the permanent calculation than σ , contradicting the assumption that σ is permanent-attaining. \square

It follows from our observation in Section 2.3 that \mathcal{P}_n also acts on the right of $\mathcal{P}(M)$. The following theorem shows that the permanent-attaining partial bijections are precisely the satisfiable partial bijections (or in fact, the partial bijections which are satisfied by elements of the max-plus column space of M).

Theorem 5.2. *Let $M \in \mathbb{R}^{n \times d}$ and let $\sigma : [d] \dashrightarrow [n]$ be a partial bijection. Then the following are equivalent:*

- (i) σ is permanent-attaining with respect to M ;
- (ii) σ is satisfiable with respect to M ;
- (iii) there exists $y \in \text{Col}_{\oplus}(M)$ satisfying σ .

Proof. Let $J \subseteq [d]$ and $I \subseteq [n]$ be the domain and image respectively of σ .

It is clear that (iii) implies (ii). We shall show that (ii) implies (i) and (i) implies (iii). Let $\Sigma = (\Sigma_{1\star}, \dots, \Sigma_{n\star})^T \in \mathbb{B}^{n \times d}$ be the element defined by $j \in \Sigma_{i\star}$ if and only if $\sigma(j) = i$.

Suppose that (ii) holds. By definition there exists y in

$$D_\Sigma = \bigcap_{j \in J} \text{Dom}_{\sigma(j)}(M_{\star j}).$$

Thus for each $j \in J$ we have

$$y_{\sigma(j)} - M_{\sigma(j),j} = \min_k \{y_k - M_{k,j}\},$$

or in other words

$$M_{k,j} - M_{\sigma(j),j} \leq y_k - y_{\sigma(j)}, \quad \text{for all } j \in J \text{ and all } k \in [n]. \quad (5.1)$$

Suppose for contradiction that σ does not attain the permanent. Then there is another bijection $\tau : J \rightarrow I$ with $\tau \neq \sigma$ such that

$$\sum_{j \in J} M_{\sigma(j),j} < \sum_{j \in J} M_{\tau(j),j},$$

giving

$$0 < \sum_{j \in J} M_{\tau(j),j} - M_{\sigma(j),j}.$$

But it follows from (5.1) that the right hand side of this inequality is less than or equal to $\sum_{j \in J} (y_{\tau(j)} - y_{\sigma(j)})$, giving a contradiction since the latter is 0 (because σ and τ are partial bijections with the same domain and *image*).

Now suppose that (i) holds. Suppose first that $|I| = |J| = n = d$. For each $\varepsilon > 0$ let $M(\varepsilon)$ be the $n \times n$ square matrix such that

$$M(\varepsilon)_{i,j} = \begin{cases} M_{i,j} & \text{if } i = \sigma(j) \\ M_{i,j} - \varepsilon & \text{otherwise.} \end{cases}$$

Then the permanent of $M(\varepsilon)$ is attained uniquely by σ and hence $M(\varepsilon)$ tropically non-singular. It follows from [17, Theorem 4.2] and [18, Proposition 17] that there is a vector $y(\varepsilon) \in \text{Col}_\oplus(M(\varepsilon))$ whose type *with respect to the columns of $M(\varepsilon)$* is an n -tuple of singleton sets whose union is $[n]$. Moreover, it is not hard to see that $\text{type}(y(\varepsilon))$ - again, with respect to the columns of $M(\varepsilon)$ - must be Σ , since (having already shown that (ii) implies (i)) we note that any other vector of singletons will yield another permanent-attaining permutation.

For each $\varepsilon > 0$ we define

$$D_\Sigma(\varepsilon) = \bigcap_{j \in J} \text{Dom}_{\sigma(j)}(M(\varepsilon)_{\star j}).$$

Thus $D_\Sigma(\varepsilon)$ is the set of all elements $x \in \mathbb{R}^n$ having $\Sigma \preceq \text{type}(x)$. Since Σ is the type of a point in the max-plus column space of $M(\varepsilon)$, we note that by [18, Theorem 15] $D_\Sigma(\varepsilon)$ corresponds to a closed and bounded subset of \mathbb{R}^{n-1} (identified with tropical projective space $\mathbb{R}^n/\mathbb{R}(1, \dots, 1)$ as in (3.5)). Moreover, it is easy to see that as $\varepsilon \rightarrow 0$ we obtain a decreasing sequence of nested closed and bounded subsets of \mathbb{R}^{n-1} . Fixing $\kappa > 0$, the set

$$D_\Sigma(\kappa) = \bigcup_{0 < \varepsilon \leq \kappa} D_\Sigma(\varepsilon)$$

is therefore compact in projective space, with respect to the usual topology induced from \mathbb{R}^{n-1} . Since each $y(\varepsilon)$ with $\varepsilon \leq \kappa$ is contained in $D_\Sigma(\kappa)$, there is a convergent subsequence of $y(\varepsilon)$'s in projective space, the limit of which

must satisfy the limiting set of inequalities. In other words, there exists $y \in \mathbb{R}^n$ whose type T with respect to the columns of M contains Σ . By [18, Corollary 12] we note that face corresponding to T in the tropical complex generated by M must be bounded (each row of $T \in \mathbb{B}^{n \times d}$ is non-zero) and hence, by [18, Theorem 15], this face is contained in the column space of M . But now y satisfies Σ with respect to M and $y \in \text{Col}_{\oplus}(M)$.

Finally, suppose that $1 \leq k = |I| = |J| \leq n, d$. Let X denote the $k \times k$ submatrix of M whose rows are indexed by I and whose columns are indexed by J . By the argument above, there exists \hat{y} in the column space of X satisfying σ , that is:

$$\hat{y} \in \text{Col}_{\oplus}(X) \cap \bigcap_{j \in J} \text{Dom}_{\sigma(j)}(X_{\star j}) \subseteq \mathbb{R}^k.$$

Since $\hat{y} \in \text{Col}_{\oplus}(X)$ we can write $\hat{y} = \bigoplus_{j \in J} \alpha_j \otimes X_j$ for some $\alpha_j \in \mathbb{R}$, and one can deduce from definition (3.3) that

$$\hat{y} = \bigoplus_{j \in J} \langle X_{\star j} | \hat{y} \rangle \otimes X_{\star j}.$$

Now set

$$y = \bigoplus_{j \in J} \langle X_{\star j} | \hat{y} \rangle \otimes M_{\star j}.$$

Thus $y \in \text{Col}_{\oplus}(M)$ and $y_t = \hat{y}_t$ for all $t \in I$. Notice that for all $j \in J$

$$\langle X_j | \hat{y} \rangle = \hat{y}_{\sigma(j)} - X_{\sigma(j),j} = y_{\sigma(j)} - M_{\sigma(j),j}.$$

Thus for all $t \in [n]$ and all $j \in J$ we have

$$y_t = \max_{l \in J} (\langle X_{\star l} | \hat{y} \rangle + M_{t,l}) \geq \langle X_{\star j} | \hat{y} \rangle + M_{t,j} = y_{\sigma(j)} - M_{\sigma(j),j} + M_{t,j},$$

giving $y_t - M_{t,j} \geq y_{\sigma(j)} - M_{\sigma(j),j}$ for all $j \in J$ and all $t \in [n]$. In other words, M_j dominates y in position $\sigma(j)$ for all $j \in J$, showing that y satisfies Σ as required. \square

In light of the previous theorem we may rephrase Theorem 4.1 as follows:

Corollary 5.3. *Let $M \in \mathbb{R}^{n \times d}$ and let $S \in \mathbb{B}^{n \times d}$. Then S is satisfiable with respect to M if and only if every partial bijection contained in S is permanent-attaining with respect to M .*

It follows from Theorem 5.2 that one may deduce the permanent structure $\mathcal{P}(M)$ from the collection of all types $\mathcal{F}(M)$. In the next section we shall prove the converse, hence showing that $\mathcal{P}(M)$ and $\mathcal{F}(M)$ encode exactly the same combinatorial data about the tropical complex. We conclude this section with the following crucial observation.

Lemma 5.4. *Let $M \in \mathbb{R}^{n \times d}$ and let $T \in \mathbb{B}^{n \times d}$ be a type with respect to the columns of M . If T contains a partial bijection $\sigma : [d] \dashrightarrow [n]$, then T contains all of the permanent-attaining (with respect to M) partial bijections that have the same domain of definition and image as σ .*

Proof. Since T is a type, there exists $x \in \mathbb{R}^n$ having type T . By definition, x satisfies $A \in \mathbb{B}^{n \times d}$ only if $A \preceq T$, and so in particular x satisfies σ . Now if τ is any permanent-attaining partial bijection with the same domain of

definition and image as σ it follows from Theorem 5.2 that τ is satisfiable. We show that x also satisfies τ , from which we conclude that T contains τ .

Since σ and τ are satisfiable it follows from Theorem 5.2 that these partial bijections attain the permanent P in the submatrix of M with rows indexed by I and columns indexed by J :

$$P = \sum_{j \in J} M_{\sigma(j),j} = \sum_{j \in J} M_{\tau(j),j}.$$

Suppose that x satisfies σ . In other words, suppose that

$$x_{\sigma(j)} - M_{\sigma(j),j} \leq x_t - M_{t,j} \text{ for all } j \in J \text{ and all } t \in [n].$$

We note that since the image of τ is a subset of $[n]$, we have in particular

$$x_{\sigma(j)} - M_{\sigma(j),j} \leq x_{\tau(j)} - M_{\tau(j),j} \text{ for all } j \in J.$$

We claim that each of the last inequalities has to be an equality. Indeed, suppose not. Then summing them gives

$$\sum_{j \in J} x_{\sigma(j)} - M_{\sigma(j),j} < \sum_{j \in J} x_{\tau(j)} - M_{\tau(j),j}.$$

But since σ and τ are bijections with the same *image* I , we see that this becomes

$$\sum_{i \in I} x_i - \sum_{j \in J} M_{\sigma(j),j} < \sum_{i \in I} x_i - \sum_{j \in J} M_{\tau(j),j}.$$

We note that the first summations on either side are identical, whilst each of the remaining summations is equal to P by assumption, giving a contradiction and proving the claim. Thus we must have

$$x_{\tau(j)} - M_{\tau(j),j} = x_{\sigma(j)} - M_{\sigma(j),j} \leq x_t - M_{t,j}$$

for all $j \in J$ and all $t \in [n]$. Or in other words, x satisfies τ . \square

6. CHARACTERIZING TYPES

Given a matrix M defining a tropical complex and an $n \times d$ Boolean matrix (or equivalently, a d -tuple of subsets of $[n]$) S , it is natural to ask whether S is a type of the complex. Our results so far suffice to establish some elementary properties which S must possess if it is to be a type; our aim in this section is to establish a collection of conditions which are necessary and sufficient for S to be a type:

Theorem 6.1. *Let $M \in \mathbb{R}^{n \times d}$ and $S \in \mathbb{B}^{n \times d}$. Then S is a type of the tropical complex corresponding to M if and only if:*

- (T1) every column of S is non-empty;
- (T2) every partial bijection contained in S is permanent-attaining with respect to M ; and
- (T3) if S contains a partial bijection σ , then S contains all of the permanent-attaining (with respect to M) partial bijections with the same domain of definition and image as σ .

Theorem 6.1 is related to a recent result of Joswig and Loho [21, Proposition 38], obtained independently and published first, which gives a corresponding condition in the language of perfect matchings of graphs.

Notice that the three conditions (T1), (T2) and (T3) depend only on the matrix S and the permanent structure of M ; there is no direct dependence on the matrix M or its associated tropical complex. It follows that one can determine whether a given Boolean matrix is a type by reference only to the permanent structure. Conversely, by Theorem 5.2 a partial bijection is permanent-attaining if and only if it is contained in a type, so the permanent structure can be also be recovered from the types. This justifies the following claim which we made earlier:

Corollary 6.2. *The types of the tropical complex associated to a matrix M are completely determined by the permanent structure of M , and vice versa.*

Corollary 6.2 is closely related to a theorem of Fink and Rincón [19, Theorem 4.4], who showed that the ‘matching multifield’ of a matrix M (in our language, the partial bijections attaining the min-plus permanent in a maximal dimensional submatrix of M) depends only upon the combinatorial type of the hyperplane arrangement defined by M , or in other words, the collection of all types of the tropical complex associated to M .

We now turn our attention to the proof of Theorem 6.1. We begin with the direct implication, which is straightforward modulo what we have already shown:

Proof of the direct implication of Theorem 6.1. Suppose S is a type. Choose a point $x \in \mathbb{R}^n$ of which it is the type. Then:

- (T1) Since the point x must lie in at least one of the sectors relative to the hyperplane specified by a column of M , we see that each column of M dominates x in some position and thus each column of S must be non-empty.
- (T2) The point x satisfies every T such that $T \preceq S$, and hence in particular every partial bijection contained in S . Thus, by Theorem 5.2, every such partial bijection must be permanent-attaining.
- (T3) This is Lemma 5.4 above.

□

To establish the converse implication of Theorem 6.1 we shall need some further technical lemmas:

Lemma 6.3. *Let $S = (S_{1\star}, \dots, S_{n\star})^\top \in \mathcal{S}(M)$ satisfying conditions (T1) and (T3) above and choose a point x satisfying S . Let $T = (T_{1\star}, \dots, T_{n\star})^\top$ be the type of x and define a binary relation \sim on $[n]$ by $p \sim q$ if $S_{p\star} \cap T_{q\star} \neq \emptyset$.*

- (i) *If $p \sim q$ then $x_p - x_q = M_{p,g} - M_{q,g}$ for all $g \in S_{p\star} \cap T_{q\star}$.*
- (ii) *If $p \sim q$, y satisfies S and $y_p = x_p + \varepsilon$ for some $\varepsilon > 0$, then $y_q \geq x_q + \varepsilon$.*
- (iii) *Suppose that $I \subseteq [n]$ with the property that if $p \in I$ and $p \sim q$ then $q \in I$. Let $\varepsilon > 0$ and let $y(\varepsilon)$ be the element obtained from x by adding ε to the entries indexed by I . Then there exists $\kappa > 0$ such that $y(\varepsilon)$ satisfies S for all $0 < \varepsilon < \kappa$.*

Proof. (i) Suppose that $p \sim q$ and let $g \in S_{p^\star} \cap T_{q^\star}$. Since $g \in S_{p^\star}$ and x satisfies S , we have $x_p - M_{p,g} \leq x_q - M_{q,g}$. On the other hand, since $g \in T_{q^\star}$ and x satisfies T , we have $x_q - M_{q,g} \leq x_p - M_p$. It follows immediately from these two inequalities that $x_p - x_q = (M_g)_p - (M_g)_q$.

(ii) Suppose that $p \sim q$ and choose $g \in S_{p^\star} \cap T_{q^\star}$. By part (i) we know that $x_p - x_q = M_{p,g} - M_{q,g}$. Now, since y satisfies S we have

$$y_p - M_{p,g} \leq y_q - M_{q,g}.$$

Rearranging this last inequality and using the fact that $y_p = x_p + \varepsilon$ gives

$$x_p + \varepsilon = y_p \leq y_q + M_{p,g} - M_{q,g} = y_q + x_p - x_q,$$

and hence $x_q + \varepsilon \leq y_q$ as required.

(iii) For any $\varepsilon > 0$ we have that $x \leq y(\varepsilon)$ with equality in positions not indexed by I . It follows that $y(\varepsilon)$ must satisfy the conditions specified by the subsets S_{k^\star} with $k \notin I$; specifically, if $j \in S_{k^\star}$ for some $k \notin I$, then

$$y(\varepsilon)_k - M_{k,j} = x_k - M_{k,j} \leq x_t - M_{t,j} \leq y(\varepsilon)_t - M_{t,j}, \text{ for all } t \in [n].$$

Moreover, $j \in S_{k^\star}$ for some $k \in I$, then

$$y(\varepsilon)_k - M_{k,j} = x_k + \varepsilon - M_{k,j} \leq x_t + \varepsilon - M_{t,j}, \text{ for all } t \in [n].$$

In particular, this gives

$$y(\varepsilon)_k - M_{k,j} \leq y(\varepsilon)_t - M_{t,j}, \text{ for all } t \in I.$$

It remains to show that we can choose ε small enough so that $y(\varepsilon)$ satisfies the remaining inequalities of S too. Thus we need to be able to choose $\varepsilon > 0$ to simultaneously satisfy the inequalities

$$x_k + \varepsilon - M_{k,j} \leq x_q - M_{q,j}, \text{ where } k \in I, j \in S_{k^\star} \text{ and } q \notin I.$$

Or in other words, we want to show that there exists ε satisfying

$$0 < \varepsilon \leq \min\{x_q - x_k + M_{k,j} - M_{q,j}\},$$

where the minimum takes place over all $k \in I$, $j \in S_{k^\star}$ and $q \notin I$. Suppose not. Then, since there are only finitely many terms in the minimum, we must have

$$x_q - x_k + M_{k,j} - M_{q,j} \leq 0 \text{ for some } k \in I, j \in S_{k^\star} \text{ and } q \notin I.$$

Rearranging this gives

$$x_q - M_{q,j} \leq x_k - M_{k,j}.$$

Since $j \in S_{k^\star}$ we know that x is dominated by M_{k^\star} in position k , and so

$$x_k - M_{k,j} = \min_t\{x_t - M_{t,j}\} \leq x_q - M_{q,j}.$$

But then the last two inequalities combine to give

$$x_q - M_{q,j} = x_k - M_{k,j} = \min_t\{x_t - M_{t,j}\},$$

showing that $j \in S_{k^\star} \cap T_{q^\star}$, and hence $k \sim q$. Since $k \in I$, by assumption we must have that $q \in I$ too, contradicting $q \notin I$. \square

We shall use the previous lemma to give an inductive argument for our characterization of types.

Lemma 6.4. *Let $S \in \mathcal{S}(M)$ satisfying conditions (T1) and (T3) above, and choose a point $x \in \mathbb{R}^n$ satisfying S . Let T be the type of x and suppose that $T \neq S$. Then there exists $y \in \mathbb{R}^n$ such that $S \preceq \text{type}(y) \preceq T$ with $\text{type}(y) \neq T$.*

Proof. We shall construct y by increasing x_i by a small amount ε for all positions i contained in some carefully chosen subset $I \subseteq [n]$. It is straightforward to check that given any choice of I we can choose ε small enough so that $\text{type}(y) \preceq T$.

Suppose that I has the property that if $p \in I$ and $p \sim q$ then $q \in I$. Then Lemma 6.3 (iii) tells us that we can choose ε so that S is still satisfied. Thus for any such a choice of I we may choose ε small enough so that $S \preceq \text{type}(y) \preceq T$. It remains to show that some choice of I gives $\text{type}(y) \neq T$.

Let $S = (S_{1\star}, \dots, S_{n\star})^\top$ and $T = (T_{1\star}, \dots, T_{n\star})^\top$. Since $S \neq T$, we have $j \in T_{i\star}$, but $j \notin S_{i\star}$ for some $j \in [d]$ and $i \in [n]$. Since S satisfies (T1), j must occur somewhere in S . Suppose that $j \in S_{m\star}$. We first show that choosing I so that $i \in I$ and $m \notin I$ yields $j \notin \text{type}(y)_{i\star}$ and hence $\text{type}(y) \neq T$. Suppose for contradiction that $j \in \text{type}(y)_{i\star}$. Then $M_{\star j}$ dominates y in position i we have in particular

$$y_i - M_{i,j} \leq y_m - M_{m,j}$$

and since $i \in I$ and $m \notin I$ this gives

$$x_i + \varepsilon - M_{i,j} \leq x_m - M_{m,j}.$$

But since $j \in S_{m\star}$ and x satisfies S , we also know that $M_{\star j}$ dominates x in position m , giving in particular that

$$x_m - M_{m,j} \leq x_i - M_{i,j}.$$

It is easy to see that the last two inequalities combined contradict $\varepsilon > 0$.

To complete the proof we claim that there is a subset $I \subseteq [n]$ such that

- (a) if $p \in I$ and $p \sim q$ then $q \in I$; and
- (b) $i \in I$ and $m \notin I$.

Suppose for contradiction that any subset I satisfying condition (a) and containing i must also contain m . It follows that there must be a sequence

$$i \sim p_2 \sim \dots \sim p_{r-1} \sim m.$$

By relabelling co-ordinates as necessary we may assume that $i = 1$, $p_2 = 2, \dots, p_{r-1} = r - 1$ and $m = r$, so that $j \in T_{1\star} \cap S_{r\star}$, but $j \notin S_{1\star}$ and

$$1 \sim 2 \sim \dots \sim r - 1 \sim r.$$

Thus for $s = 1, \dots, r - 1$ we may choose $j_s \in S_{s\star} \cap T_{s+1\star}$, giving:

$$\begin{aligned} Q &= (j_1, j_2, \dots, j_{r-1}, j, \emptyset, \dots, \emptyset)^\top \preceq S \\ R &= (j, j_1, j_2, \dots, j_{r-1}, \emptyset, \dots, \emptyset)^\top \preceq T. \end{aligned}$$

It is clear that any partial bijection contained in Q or R will be satisfiable (by Theorem 4.1 it will be satisfied by x), and hence must attain the permanent in the corresponding submatrix of M . We claim that there are partial bijections $Q' \preceq Q$ and $R' \preceq R$ such that Q' and R' have the same domain of definition and the same image, with $Q'_1 = Q_1$ and $R'_1 = R_1$. We first note that this will give the desired contradiction. Since S satisfies condition (T3)

we know that the set of all partial bijections contained in S contains either all or none of the permanent-attaining bijections on any given submatrix. Thus it follows from the fact that Q' is a partial bijection contained in S that R' (a partial bijection with the same domain and image) must be contained in S too. But then $\{j\} = R'_1 \subseteq S_1$, giving a contradiction.

The partial bijections Q' and R' are constructed as follows. We know that $j_1 \in Q_{1\star}$, $j \in R_{1\star}$, $j_1 \neq j$ and j occurs in at least one row of Q . Let $i_1 \in \{2, \dots, r\}$ be the first row of Q in which j occurs and let $\rho(j)$ denote the element contained in $R_{i_1\star}$. Since R is just Q shifted one place to the right it is clear that $\rho(j) \neq j$. If $\rho(j) = j_1$, then we will have partial bijections of the form

$$\begin{aligned} Q' &= (j_1, \emptyset, \dots, \emptyset, j, \emptyset, \dots, \emptyset)^\top \preceq S \\ R' &= (j, \emptyset, \dots, \emptyset, j_1, \emptyset, \dots, \emptyset)^\top \preceq T. \end{aligned}$$

Otherwise, $\rho(j) \neq j_1$. Let i_2 be the first position of Q in which $\rho(j)$ occurs. Since R shifts Q one place to the right it is clear that $i_2 < i_1$. Let $\rho^2(j)$ denote the element contained in R_{i_2} . Again, since R shifts Q one place to the right it is clear that $\rho^2(j) \neq \rho(j)$. If $\rho^2(j) = j_1$, then we will have partial bijections of the form

$$\begin{aligned} Q' &= (j_1, \emptyset, \dots, \emptyset, \rho(j), \emptyset, \dots, \emptyset, j, \emptyset, \dots, \emptyset)^\top \preceq S \\ R' &= (j, \emptyset, \dots, \emptyset, j_1, \emptyset, \dots, \emptyset, \rho(j), \emptyset, \dots, \emptyset)^\top \preceq T. \end{aligned}$$

Otherwise $\rho^2(j) \neq \rho(j), j_1$. Continuing in this way will yield a cycle of the required form. \square

We are now ready to complete the proof of Theorem 6.1:

Proof of the converse implication of Theorem 6.1. Suppose then that S satisfies the conditions (T1)-(T3). Since S satisfies (T2), it follows from Corollary 5.3 that we may choose $x \in \mathbb{R}^n$ such that x satisfies S . Choose x so that the type of x is minimal with respect to the partial order \preceq , and denote this type by T . Thus $S \preceq T$ and if $T = S$, then we are done. Otherwise, we may apply Lemma 6.4 to find $y \in \mathbb{R}^n$ such that $S \preceq \text{type}(y) \preceq T$ with $\text{type}(y) \neq T$, contradicting our choice of x with T minimal. \square

Our characterization of types gives a combinatorial way to understand the action of \mathcal{P}_n on $\mathcal{F}(M)$ via the action on the set $\mathcal{P}(M)$.

Corollary 6.5. *The right action of the hyperplane face monoid \mathcal{P}_n of the braid arrangement on $\mathbb{B}^{n \times d}$ restricts to an action on $\mathcal{F}(M)$.*

Proof. Let $T = (T_{\star 1}, \dots, T_{\star d}) \in \mathcal{F}(M)$ and $P \in \mathcal{P}_n$. Using our characterization of types, we show that $T \circ P \in \mathcal{F}(M)$, by showing that it satisfies the three conditions given by Theorem 6.1.

(T1) By definition we see that the j th column of $T \circ P$ is $T_{\star j} \circ P$. Since T is a type, we know that $\emptyset \neq T_{\star j} \subseteq [n]$ and since P is an ordered set partition of $[n]$ it is therefore clear that $T_{\star j}$ has non-empty intersection with at least one component of P . Thus by (2.2) we conclude that $T_{\star j} \circ P$ is non-empty.

(T2) By Proposition 2.1 we have $T \circ P \preceq T$, so any partial bijection contained in $T \circ P$ must also be contained in T , and so by Theorem 6.1 is permanent-attaining.

(T3) Suppose that σ and τ are partial bijections contained in T with the same domain $J \subseteq [d]$ and image $I \subseteq [n]$. By our previous remarks, it suffices to show that if σ is contained in $T \circ P$ then τ is contained in $T \circ P$ too. Assume by relabelling if necessary that $I = J = [k]$, and that $\sigma(j) = j$ for all $j \in [k]$, with τ a non-identity permutation of $[k]$. Since σ (the identity map on $[k]$) and τ are both contained in T , by definition we know that:

$$j, \tau(j) \in T_{\star j} \text{ for all } j = 1, \dots, k. \quad (6.1)$$

For each $j \in [k]$ we define

$$m(j) := \max\{m : T_{\star j} \cap P_m \neq \emptyset\}, \quad (6.2)$$

so that by (2.2) we have $T_{\star j} \circ P = P_{m(j)}$. Since σ (the identity) is contained in $T \circ P$ it follows that

$$j \in P_{\star m(j)} \text{ for all } j \in [k]. \quad (6.3)$$

Now fix j and look at the orbit, $\{\tau(j), \dots, \tau^p(j) = j\}$ say, of j under τ . By applying (6.1) and (6.3) to each of the elements $\tau^i(j) \in [k]$ in turn, we obtain:

$$\begin{aligned} \tau(j) &\in T_{\star \tau(j)}, T_{\star j}, P_{\star m(\tau(j))} \\ \tau^2(j) &\in T_{\star \tau^2(j)}, T_{\star \tau(j)}, P_{\star m(\tau^2(j))} \\ &\vdots \\ \tau^{p-1}(j) &\in T_{\star \tau^{p-1}(j)}, T_{\star \tau^{p-2}(j)}, P_{\star m(\tau^{p-1}(j))} \\ j = \tau^p(j) &\in T_{\star j}, T_{\star \tau^{p-1}(j)}, P_{\star m(j)} \end{aligned}$$

Note that if $m(j) < m(\tau(j))$ then by definition (6.2) we see that $T_{\star j} \cap P_{\star m(\tau(j))} = \emptyset$. Since $\tau(j)$ is contained in this intersection, we must have $m(j) \geq m(\tau(j))$. Continuing in this way we note that

$$m(j) \geq m(\tau(j)) \geq \dots \geq m(\tau^p(j)) = m(j),$$

from which we conclude that these indices all coincide and hence by combining the information above we see that the set $P_{\star m(j)}$ contains the entire orbit of j under τ . In particular, $\tau(j) \in P_{\star m(j)}$ for all $j = 1, \dots, k$, hence showing that τ is contained in $T \circ P$ as required. \square

Since $A \circ P \preceq A$ for all $A \in \mathbb{B}^{n \times d}$ it follows that if a type T is equal to $T' \circ P$ for some other type T' , then T' must be a face of T . Moreover, it is easy to see that each maximal dimensional cell contained in $\mathcal{F}(M)$ is a fixed point of the right action of \mathcal{P}_n . Thus (just as in the case of usual hyperplane arrangements), this right action cannot be used to generate a random walk on the maximal dimensional cells (“chambers”) of the tropical hyperplane arrangement. We also note that although \mathcal{P}_n acts upon $\mathcal{F}(M)$, this action does not immediately restrict to give an action on the max-plus polytope generated by the columns of M (that is, the collection of bounded cells). For example, acting upon the face H of the tropical hyperplane arrangement given in Example 3.1 by the element $(\{3\}, \{2\}, \{1\}) \in \mathcal{P}_3$ gives the unbounded face E , which is not contained in the max-plus column space of M .

A very natural question is whether there is a natural multiplication on the faces of a tropical hyperplane arrangement, analogous to the well-studied *left regular band* structure of a real hyperplane arrangement. Indeed, our work was partly motivated by a desire to construct such a product, in order to generate random walks through the faces of a tropical arrangement. However, we have not been able to find such a structure in the tropical case.

ACKNOWLEDGEMENTS

We are thankful to the referee for their many helpful comments and suggestions, and also to Georg Loho for making us aware of some connections to the literature.

REFERENCES

- [1] M. Aguiar, S. Mahajan, Coxeter groups and Hopf algebras, Fields Institute Monographs, 23. American Mathematical Society, Providence, RI, 2006.
- [2] J. Almeida, S. Margolis, B. Steinberg, M. Volkov, *Representation theory of finite semigroups, semigroup radicals and formal language theory*, Trans. Amer. Math. Soc., 361, 1429–1461, 2009.
- [3] F. Ardila, M. Develin, *Tropical hyperplane arrangements and oriented matroids*, Math. Z., 262, 795–816, 2009.
- [4] A. Ayer, A. Schilling, B. Steinberg, N. M. Thiery, *Markov chains, R-trivial monoids and representation theory*, Internat. J. Algebra Comput., 25, no. 1-2, 169–231, 2015.
- [5] C. A. Athanasiadis, P. Diaconis, *Functions of random walks on hyperplane arrangements*, Adv. in Appl. Math., 45, no. 3, 410–437, 2010.
- [6] C. Berg, N. Bergeron, S. Bhargava, F. Saliola, *Primitive orthogonal idempotents for R-trivial monoids*, J. Algebra, 348, 446–461 2011.
- [7] P. Bidigare, P. Hanlon, D. Rockmore, *A combinatorial description of the spectrum for the Tsetlin library and its generalization to hyperplane arrangements*, Duke Math. J., 99, no. 1, 135–174, 1999.
- [8] A. Björner, *Random walks, arrangements, cell complexes, greedoids, and self-organizing libraries*, Building bridges, 165–203, Bolyai Soc. Math. Stud., 19, Springer, Berlin, 2008.
- [9] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, G. M. Ziegler, *Oriented matroids*, Second edition, Encyclopedia of Mathematics and its Applications, 46. Cambridge University Press, Cambridge, 1999.
- [10] T. S. Blyth, M. F. Janowitz. *Residuation theory*. International Series of Monographs in Pure and Applied Mathematics, vol. 102, Pergamon Press, Oxford, 1972.
- [11] K. S. Brown, P. Diaconis, *Random walks and hyperplane arrangements*, Ann. Probab., 26, no. 4, 1813–1854, 1998.
- [12] K. S. Brown, *Semigroups, rings, and Markov chains*, J. Theoret. Probab., 13, no. 3, 871–938, 2000.
- [13] K. S. Brown, *Semigroup and ring theoretical methods in probability*, in Representations of finite dimensional algebras and related topics in Lie theory and geometry, 3–26, Fields Inst. Commun., 40, Amer. Math. Soc., Providence, RI, 2004.
- [14] P. Butkovic, *Max-linear systems: theory and algorithms*. Springer Monographs in Mathematics. Springer-Verlag London Ltd., London, 2010.
- [15] G. Cohen, S. Gaubert, J.-P. Quadrat, *Duality and separation theorems in idempotent semimodules*, Linear Algebra Appl., 379, 395–422, 2004.
- [16] R. Cuninghame-Green, *Minimax algebra*. Lecture Notes in Economics and Mathematical Systems 166. Springer-Verlag, Berlin-New York, 1979.
- [17] M. Develin, F. Santos, B. Sturmfels. *On the rank of a tropical matrix*. In *Combinatorial and computational geometry*, volume 52 of *Math. Sci. Res. Inst. Publ.*, pages 213–242. Cambridge Univ. Press, Cambridge, 2005.
- [18] M. Develin, B. Sturmfels, *Tropical convexity*, Doc. Math., 9, 1–27 (electronic), 2004.

- [19] A. Fink, F. Rincón, *Stiefel tropical linear spaces*, J. Combin. Theory Ser. A, 135, 291–331, 2015.
- [20] M. Johnson, M. Kambites, *Convexity of tropical polytopes*, Linear Algebra Appl., 485, 531–544, 2015.
- [21] M. Joswig, G. Loho, *Weighted digraphs and tropical cones*, Linear Algebra Appl., 501, 304–343, 2016.
- [22] S. Margolis, F. Saliola, B. Steinberg, *Semigroups embeddable in hyperplane face monoids*, Semigroup Forum, 89, no. 1, 236–248, 2014.
- [23] S. Margolis, F. Saliola, B. Steinberg, *Cell complexes, poset topology and the representation theory of algebras arising in algebraic combinatorics and discrete geometry*, preprint 2015.
- [24] J.-E. Pin, Varieties of Formal Languages. Foundations of Computer Science, Plenum Publishing Corp., New York, 1986. (Translated from the French by A. Howie.)
- [25] V. Reiner, F. Saliola, V. Welker, Spectra of symmetrized shuffling operators, Mem. Amer. Math. Soc., 228, no. 1072, 2014
- [26] F. V. Saliola, *The quiver of the semigroup algebra of a left regular band*, Internat. J. Algebra Comput., 17, no. 8, 1593–1610, 2007.
- [27] F. Saliola, *The face semigroup algebra of a hyperplane arrangement*, Canad. J. Math., 61, no. 4, 904–929, 2009.
- [28] F. Saliola, *Eigenvectors for a random walk on a left-regular band*, Adv. in Appl. Math., 48, no. 2, 306–311, 2012.
- [29] M. Schocker, *The module structure of the Solomon-Tits algebra of the symmetric group*, J. Algebra, 301, no. 2, 554–586, 2006.
- [30] L. Solomon, *A Mackey formula in the group ring of a Coxeter group*, J. Algebra, 41, no. 2, 255–264, 1976.
- [31] J. Tits, *Two properties of Coxeter complexes*, J. Algebra, 41, no. 2, 265–268, 1976 (appendix to L. Solomon’s article *A Mackey formula in the group ring of a Coxeter group*, J. Algebra, 41, no. 2, 255–264, 1976).