



A version of Tutte's polynomial for hypergraphs

Tamás Kálmán

Tokyo Institute of Technology, Department of Mathematics, mail code H-214, 2-12-1 Ookayama, Meguro-ku, Tokyo 152-8551, Japan

Received 3 June 2011; accepted 3 June 2013

Available online 26 June 2013

Communicated by Erwin Lutwak

Abstract

Tutte's dichromate $T(x, y)$ is a well known graph invariant. Using the original definition in terms of internal and external activities as our point of departure, we generalize the valuations $T(x, 1)$ and $T(1, y)$ to hypergraphs. Our generating functions are sums over hypertrees, i.e., instances of a certain generalization of the indicator function of the edge set of a spanning tree. We prove that hypertrees are the lattice points in a polytope which in turn is the set of bases in a polymatroid. In fact, we extend our invariants to integer polymatroids as well. Several properties are established, including a generalization of the deletion–contraction formulas. We also examine hypergraphs that can be represented by planar bipartite graphs, write their hypertree polytopes in the form of a determinant, and prove a duality property that leads to an extension of Tutte's Tree Trinity Theorem.

© 2013 Elsevier Inc. All rights reserved.

MSC: 05; 52

Keywords: Hypergraph; Bipartite graph; Tutte polynomial; Integer polytope; Planar duality; Arborescence number

1. Introduction

In this paper we investigate how much of the theory of the Tutte polynomial, which is a large and important branch of graph and matroid theory, can be generalized to hypergraphs and polymatroids. We will see that some of the basic definitions carry over in a new and interesting way. From among the many important properties of the Tutte polynomial, we find an elegant generalization of the equation relating planar dual graphs. The deletion–contraction rule also has

E-mail address: kalman@math.titech.ac.jp.

an extension to our context but, interestingly, such formulas play a lesser role here than in the classical case.

The central object of our theory is the integer polytope $Q_{\mathcal{H}}$, associated to a hypergraph \mathcal{H} , which generalizes the well known spanning tree polytope of ordinary graphs. The lattice points in $Q_{\mathcal{H}}$ are called hypertrees. A hypertree, see Definition 3.1, is essentially the valence distribution, taken at one side of the bipartition, of a spanning tree in the bipartite graph naturally associated to \mathcal{H} . It is the use of this concept that sets apart the current paper from earlier work.

From $Q_{\mathcal{H}}$, we will read off the one-variable polynomial invariants $I_{\mathcal{H}}$ and $X_{\mathcal{H}}$ of \mathcal{H} which generalize $T_G(x, 1)$ and $T_G(1, y)$, respectively. (Here T_G is the Tutte polynomial of the graph G and we think of \mathcal{H} as a generalization of G .) They are called the interior and exterior polynomials. Both have positive integer coefficients, the sum of which is the number of hypertrees. Their definitions involve a direct, yet non-obvious generalization of Tutte's idea of ordering the (hyper)edges and using the order to define their internal and external active/inactive status with respect to (hyper)trees.

The hypertree polytope is naturally embedded as the set of bases in a certain integer polymatroid. The latter can be viewed as a generalization to hypergraphs of the cycle matroid of a graph. The interior and exterior polynomials can in fact be defined for all integer (extended) polymatroids, i.e., for all integer-valued submodular set functions regardless of whether they are non-decreasing. In this paper we will merely mention this possibility as we intend to keep the focus on bipartite graphs and hypergraphs.

As the counterpart of the planar duality relation $T_{G^*}(x, y) = T_G(y, x)$ we obtain the formula $I_{\mathcal{H}^*} = X_{\mathcal{H}}$. Here the dual \mathcal{H}^* is naturally associated to \mathcal{H} when the hypergraph \mathcal{H} can be represented with a plane bipartite graph.

Our invariants have another fundamental property which does not manifest itself in the classical case. Namely, any hypergraph \mathcal{H} has an abstract dual $\overline{\mathcal{H}}$ that results from interchanging the roles of its vertices and hyperedges. For such a pair, we conjecture the identity $I_{\mathcal{H}} = I_{\overline{\mathcal{H}}}$.

A weaker statement is that \mathcal{H} and $\overline{\mathcal{H}}$ have the same number of hypertrees. This was recently proven by Alexander Postnikov. Indeed the discovery of $Q_{\mathcal{H}}$ (and its dual relationship with $Q_{\overline{\mathcal{H}}}$) should be attributed to him. In his beautiful paper [4] he puts these polytopes in several important contexts. The polynomials $I_{\mathcal{H}}$ and $X_{\mathcal{H}}$ however seem to appear for the first time in this article. I read Postnikov's work after most results presented here have been obtained. Prior to that I was only able to prove that $Q_{\mathcal{H}}$ and $Q_{\overline{\mathcal{H}}}$ had the same number of lattice points when \mathcal{H} was planar. This brings us to the second half of the paper.

We will revisit a delicate picture that William Tutte [7] introduced before his discovery of the dichromate. It consists of three plane bipartite graphs that together triangulate the sphere S^2 . Their planar duals are canonically directed so that a so-called arborescence number can be associated to each. Tutte's Tree Trinity Theorem states that in such a triple, the three arborescence numbers coincide. We will see that this value is also the number of hypertrees in any of the six hypergraphs found in Tutte's picture. The six hypertree polytopes, for which we will derive a concise determinant formula based on work of Kenneth Berman, form three centrally symmetric pairs. The total number of interior and exterior polynomials associated to the six hypergraphs is reduced from 12 to 6 by our planar duality result, and the conjecture above on abstract duality implies that the number is in fact only 3.

The notions that we are about to introduce also have a strong connection to low-dimensional topology. Indeed, the results contained in this paper are by-products of the author's investigations in knot theory. There is a well known method, called the median construction, that one may apply

to a plane graph to obtain an alternating link diagram. When the graph is bipartite, the link has a natural orientation. The interior polynomial was originally developed to describe a pattern found in the Homfly polynomials of these associated links. Then in joint work with András Juhász and Jacob Rasmussen, we discovered a manifestation of the hypertree polytope in Heegaard Floer theory. Therefore the results of this paper provide an unexpected link between the Homfly polynomial and Floer homology. But since the result linking the Homfly and interior polynomials remains a conjecture at the time of this writing, it feels prudent to save the exploration of topological aspects to future papers.

In another forthcoming paper, joint with Alex Bene, we extend the theory of the hypertree polytope and interior and exterior polynomials to topological bipartite graphs and hypergraphs, in the same way that the Bollobás–Riordan polynomial generalizes the Tutte polynomial for fatgraphs.

The paper is organized as follows. Section 2 fixes terminology and recalls Tutte’s definition of the dichromate. In Section 3, we introduce (that is to say, recall from [4]) the hypertree polytope and in Section 4 we examine its basic geometry. After all this preparation, in Section 5 we define interior and exterior polynomials. The material on polymatroids is mainly found in Sections 4.1 and 5.2. Section 6 establishes a variety of properties of the new invariants. In Section 7 we state the main unsolved problem of the paper, the abstract duality conjecture, along with some supporting evidence. We discuss planar duality in Section 8. Finally, Section 9 introduces trinitities of plane bipartite graphs and some earlier work on them, and in Section 10 we lay out our new results on that subject.

2. Preliminaries

We will use the standard notion of a *graph* as an ordered pair $G = (V, E)$ with finite vertex set V and finite edge (multi)set E . Loop edges and multiple edges are allowed. By $G' = (V', E')$ being a *subgraph* of G we will mean $V' = V$ and $E' \subset E$ (thus subgraphs of a given graph are equivalent to their edge sets). For example, if a subgraph contains no edge adjacent to some vertex v , then $\{v\}$ (more precisely, $(\{v\}, \emptyset)$) is a connected component of the subgraph.

We will write $G \setminus \{e\}$ for the graph that results from removing (one copy of) the edge e from E . The symbol $G - \{v\}$ will stand for the graph that remains after removing the vertex v and all its adjacent edges from G .

Graphs can be viewed as one-dimensional cell complexes and in that sense the *nullity* of a graph, $n(G)$, is the rank of its first homology group. (Thus we may also refer to the nullity as the *first Betti number*.)

2.1. Review of the Tutte polynomial

Tutte’s original construction [8] of the two-variable polynomial $T_G(x, y)$, associated to the graph $G = (V, E)$, proceeds as follows. Order E arbitrarily. Consider the set \mathcal{T} of *spanning trees*, that is connected and cycle-free subgraphs, for G . (In order for \mathcal{T} to be non-empty, G needs to be connected. This will almost always be assumed.)

Definition 2.1. Given a spanning tree $\Gamma \in \mathcal{T}$, call an edge $e \in \Gamma$ *internally active* with respect to Γ if, after removing e from Γ , connectedness of the subgraph cannot be restored by adding an edge to $\Gamma \setminus \{e\}$ that is smaller than e .

An edge $e \notin \Gamma$ is *externally active* if, after adding e to Γ , cycle-freeness cannot be restored by removing an edge from $\Gamma \cup \{e\}$ that is smaller than e .

In fact Tutte put ‘larger’ instead of ‘smaller’ in both cases above. But as we will see in Theorem 2.3 below, reversing the order does not affect the following notion.

Definition 2.2. Let $\iota(\Gamma)$ and $\varepsilon(\Gamma)$ denote the number of internally and externally active edges, respectively, with respect to Γ (under the given order). Then, the *Tutte polynomial* or *dichromate* of the graph G is the generating function

$$T_G(x, y) = \sum_{\Gamma \in \mathcal{T}} x^{\iota(\Gamma)} y^{\varepsilon(\Gamma)}.$$


Theorem 2.3 (Tutte [8]). *The polynomial $T_G(x, y)$ is independent of the chosen edge order.*

Many properties of $T_G(x, y)$ are known. A famous example is the deletion–contraction relation

$$\begin{aligned} T_G &= xT_{G/e} && \text{if } e \text{ is a bridge} \\ T_G &= yT_{G \setminus \{e\}} && \text{if } e \text{ is a loop} \\ T_G &= T_{G \setminus \{e\}} + T_{G/e} && \text{if } e \text{ is neither a bridge nor a loop.} \end{aligned} \tag{1}$$

(A bridge is an edge whose removal disconnects the graph. The contraction G/e will be formally defined in Section 6.3.) For an excellent survey on why the Tutte polynomial is central in graph theory, see [2].

Example 2.4. We show a graph and its Tutte polynomial.



2.2. Hypergraphs and bipartite graphs

For the purposes of this paper these structures are almost equivalent, as follows.

Definition 2.5. A *bipartite graph* is a triple $G = (V_0, V_1, E)$, where V_0 and V_1 are disjoint finite sets, called *color classes*, and E is a finite set of edges, each connecting an element of V_0 to an element of V_1 (multiple edges are not allowed). We will treat (V_0, V_1, E) , (V_1, V_0, E) , and the graph $(V_0 \cup V_1, E)$ as the same object.

Definition 2.6. A *hypergraph* is a pair $\mathcal{H} = (V, E)$, where V is a finite set and E is a finite multiset of non-empty subsets of V . Elements of V are called *vertices* and the elements of E are the *hyperedges*.

Thus, hyperedges with multiplicity (that is, several copies of the same subset of V) are allowed. On the other hand, obviously, each hyperedge contains each of its elements exactly once.

A hypergraph is both a generalization and a special case of a graph. The first point is obvious. Conversely, the sets V and E that constitute the hypergraph \mathcal{H} may be viewed as the color

classes of a bipartite graph $\text{Bip } \mathcal{H}$: we connect $v \in V$ to $e \in E$ with an edge if $v \in e$. We will refer to the result as the *bipartite graph associated to the hypergraph* \mathcal{H} .

The construction of $\text{Bip } \mathcal{H}$ above is reversible if we specify one of the two color classes in the bipartite graph $G = (V_0, V_1, E)$. We will use the notation

$$\mathcal{G}_0 = (V_1, V_0) \quad \text{and} \quad \mathcal{G}_1 = (V_0, V_1) \tag{2}$$

for the resulting pair of hypergraphs. (The index of the hypergraph is chosen to emphasize its hyperedges.) Note that $\text{Bip } \mathcal{G}_0 = \text{Bip } \mathcal{G}_1 = G$.

Definition 2.7. The bipartite graph G above is said to *induce* the hypergraphs \mathcal{G}_0 and \mathcal{G}_1 . Two hypergraphs are called *abstract duals* if they can be obtained in the form (2). In other words, the abstract dual $\overline{\mathcal{H}} = (E, V)$ of a hypergraph $\mathcal{H} = (V, E)$ is defined by interchanging the roles of its vertices and hyperedges.

3. Hypertrees

To generalize Section 2.1 to hypergraphs, first we need a notion corresponding to spanning trees. The rest of the paper is built around this concept, and almost all novelty contained herein stems from its use.

Definition 3.1. Let $\mathcal{H} = (V, E)$ be a hypergraph so that its associated bipartite graph $\text{Bip } \mathcal{H}$ is connected. By a *hypertree* in \mathcal{H} we mean a function (vector) $\mathbf{f}: E \rightarrow \mathbf{N} = \{0, 1, \dots\}$ so that a spanning tree of $\text{Bip } \mathcal{H}$ can be found which has valence $\mathbf{f}(e) + 1$ at each $e \in E$. Such a spanning tree is said to *realize* or to *induce* \mathbf{f} .

If \mathbf{f} is a hypertree, then $0 \leq \mathbf{f}(e) \leq |e| - 1$ for each $e \in E$. The condition that $\text{Bip } \mathcal{H}$ be connected is not quite essential. We could talk about ‘hyperforests’ realized by spanning forests but, partly to avoid such terminology, we will generally assume that the bipartite graphs associated to our hypergraphs are connected. (It is in fact common (cf. [6, Section 3.3]) to call \mathcal{H} connected if and only if $\text{Bip } \mathcal{H}$ is connected.)

Remark 3.2. Hypertrees generalize spanning trees of graphs: there, an edge e is in the tree if and only if $\mathbf{f}(e) = 1$ and not in the tree if and only if $\mathbf{f}(e) = 0$. In our case, we allow hyperedges to be in hypertrees only “partially.”

The number of realizations may vary from hypertree to hypertree. However this phenomenon will not be incorporated to our theory in any way. For example, if \mathbf{f} is a hypertree and $\mathbf{f}(e) = 0$, then any edge of $\text{Bip } \mathcal{H}$ adjacent to the hyperedge e can be chosen to be part of a realization of \mathbf{f} , irrespective of the rest of the construction. The next lemma claims only slightly more.

Lemma 3.3. *Let $\mathcal{H} = (V, E)$ be a hypergraph and $\mathbf{f}: E \rightarrow \mathbf{N}$ a hypertree. For any collection of edges $\{\alpha_e\}$ of $\text{Bip } \mathcal{H}$ so that α_e is adjacent to e for all $e \in E$, there is a realization of \mathbf{f} that contains α_e for all e .*

Proof. Start with an arbitrary realization Γ of \mathbf{f} . If e is such that $\alpha_e \notin \Gamma$, then add α_e to Γ . This creates a unique cycle in the subgraph which goes through e . Now if we remove the edge of the cycle that is adjacent to e but different from α_e , the result is another realization of \mathbf{f} . Carry this process out at every hyperedge. \square

It turns out that hypertrees of a given hypergraph are exactly the integer lattice points in a convex polytope. This fact was first realized by Postnikov [4]. But as the author discovered it independently and our points of view and proofs are different, we shall give a full account here. This will be followed by a sampling of Postnikov’s ideas, including a sketch of the proof of his duality theorem.

3.1. *The hypertree polytope*

In a hypergraph $\mathcal{H} = (V, E)$, let $E' \subset E$ be a non-empty subset and let $\bigcup E' \subset V$ denote the set of its neighboring points in $\text{Bip } \mathcal{H}$. Let $\text{Bip } \mathcal{H}|_{E'}$, the *bipartite graph induced by E'* , be the graph with color classes E' and $\bigcup E'$ and edges inherited from $\text{Bip } \mathcal{H}$. Let us denote the number of connected components in $\text{Bip } \mathcal{H}|_{E'}$ by $c(E')$.

Theorem 3.4. *Let $\mathcal{H} = (V, E)$ be a hypergraph so that its associated bipartite graph $\text{Bip } \mathcal{H}$ is connected. The hypertrees $\mathbf{f}: E \rightarrow \mathbf{N}$ of \mathcal{H} are exactly the integer solutions of the following system of linear inequalities in \mathbf{R}^E :*

$$\mathbf{f}(e) \geq 0 \quad \text{for all } e \in E \tag{3a}$$

$$\sum_{e \in E'} \mathbf{f}(e) \leq \left| \bigcup E' \right| - c(E') \quad \text{for all non-empty } E' \subset E \tag{3b}$$

$$\sum_{e \in E} \mathbf{f}(e) = |V| - 1. \tag{3c}$$

Remark 3.5. We get an equivalent system of inequalities if we replace (3b) with $\sum_{e \in E'} \mathbf{f}(e) \leq \left| \bigcup E' \right| - 1$, still required for all non-empty $E' \subset E$. This is because if $\text{Bip } \mathcal{H}|_{E'}$ is not connected, then its connected components are also of the form $\text{Bip } \mathcal{H}|_{E''}$ and we get the stronger version of the inequality by summing the inequalities associated to these smaller subsets E'' .

Remark 3.6. The conditions (3a) follow from (3b) and (3c): this is obvious if $E = \{e\}$ and otherwise, with $E' = E \setminus \{e\}$, we have

$$\mathbf{f}(e) = |V| - 1 - \sum_{e' \in E'} \mathbf{f}(e') \geq \left(|V| - \left| \bigcup E' \right| \right) + (c(E') - 1),$$

where the right hand side is the sum of two non-negative quantities.

Proof of Theorem 3.4. If \mathbf{f} is a hypertree, the conditions (3) hold because of the well known facts that

- (i) the number of edges in a spanning forest (maximal cycle-free subgraph) of a graph is the number of vertices minus the number of connected components of the graph and
- (ii) any cycle-free subgraph is part of a spanning forest.

Indeed, for a given hypertree \mathbf{f} and non-empty subset $E' \subset E$, any spanning tree of $\text{Bip } \mathcal{H}$ that realizes \mathbf{f} has an intersection with $\text{Bip } \mathcal{H}|_{E'}$ which is a cycle-free subgraph of the latter. As such, it may have at most $|E'| + \left| \bigcup E' \right| - c(E')$ edges. Since each of those has exactly one of its ends at an element of E' , we have

$$\sum_{e \in E'} (\mathbf{f}(e) + 1) \leq |E'| + \left| \bigcup E' \right| - c(E'),$$

which is just another form of (3b). The claim (3a) is obvious and (3c) is immediate from (i).

To see why (3) is also sufficient for \mathbf{f} to be a hypertree, let us formulate a lemma that is marginally stronger than what we currently need, but the form in which we state it will be useful later. Recall that a graph has nullity zero if and only if it is cycle-free.

Lemma 3.7. *Suppose that the integer vector $\mathbf{f}: E \rightarrow \mathbf{N}$ satisfies conditions (3a) and (3b) (including for $E' = E$) but not necessarily (3c). Then there exists a cycle-free subgraph in $\text{Bip } \mathcal{H}$ that has valence $\mathbf{f}(e) + 1$ at all elements $e \in E$.*

Let the vector \mathbf{f} satisfy the conditions and choose an arbitrary subgraph P of $\text{Bip } \mathcal{H}$ whose valence at each $e \in E$ is $\mathbf{f}(e) + 1$. This is possible because (3b) applied to $E' = \{e\}$ says $\mathbf{f}(e) \leq |e| - 1$. If P is cycle-free, we are done. Assume it is not.

It suffices to show that there is another subgraph of $\text{Bip } \mathcal{H}$ that has the same valences at elements of E as P (prescribed by \mathbf{f}) but whose nullity is one less than that of P . The subgraph P has a connected component C containing a cycle. Let C intersect E in the set E' . Applying (3b) to E' , we see that there is a hyperedge $e \in E'$ which is connected by an edge α of $\text{Bip } \mathcal{H}$ to a vertex which is not in C . (Else, P and $\text{Bip } \mathcal{H}|_{E'}$ would form a connected and not cycle-free intersection; since such a subgraph contains a spanning tree of $\text{Bip } \mathcal{H}|_{E'}$ as a proper subgraph, we would get a contradiction between (3b) and (i).)

Now if e is part of a cycle in C , we are done because we may remove from P an edge of that cycle (adjacent to e) and replace it with α ; as a result, valences at hyperedges remain the same but the nullity reduces by one. Otherwise, each edge of C adjacent to e leads to a different connected component of $C - \{e\}$. At least one of these (call it C') still contains a cycle. Replace the edge of P from e to C' with α . This results in a subgraph P' that has the right valences and the same nullity as P , but which has a non-tree component C' containing fewer elements of E than C did. Repeat the procedure of this paragraph to P' and C' . After a finite number of iterations, the desired reduction in the nullity will occur. This finishes the proof of the lemma.

If we apply Lemma 3.7 to a vector \mathbf{f} that satisfies (3), the resulting subgraph is not just cycle-free but, because of (3c), it is actually a spanning tree. \square

Definition 3.8. Let \mathcal{H} be a hypergraph. The *hypertree polytope* of \mathcal{H} , denoted with $Q_{\mathcal{H}}$, is the set of solutions in \mathbf{R}^E of the inequalities (3).

The set $Q_{\mathcal{H}}$ is indeed a polytope: convexity and closedness are obvious and (3a) and (3c) ensure boundedness. Theorem 3.4 also implies that the set of hypertrees in \mathcal{H} can be written as $Q_{\mathcal{H}} \cap \mathbf{Z}^E$. Furthermore, in Section 4.1 we will see that $Q_{\mathcal{H}}$ is an integer polytope, that is $Q_{\mathcal{H}} = \text{Conv}(Q_{\mathcal{H}} \cap \mathbf{Z}^E)$, where Conv denotes the usual convex hull.

The hypertree polytope generalizes the usual spanning tree polytope of a graph (cf. Remarks 3.2 and 3.5). Equation (3c) means that $Q_{\mathcal{H}}$ is part of an affine hyperplane in \mathbf{R}^E . Thus, $Q_{\mathcal{H}}$ is situated in the lattice cut out from \mathbf{Z}^E by the hyperplane.

Example 3.9. The graph used in Example 2.4 is in fact bipartite. Let us denote it with G and label its color classes with $V_0 = \{a, b, c\}$ and $V_1 = \{p, q, r, s\}$, as in Fig. 1. There are $T_G(1, 1) = 50$ spanning trees in G . They give rise to seven hypertrees in \mathcal{G}_0 , and the same number in \mathcal{G}_1 , cf. Theorem 3.11. These are shown, with a concrete spanning tree realizing each, in Fig. 6. Because of (3c), it is natural to view the polytopes $Q_{\mathcal{G}_0}$ and $Q_{\mathcal{G}_1}$ in barycentric coordinate systems with basepoints labeled with

$$\mathbf{a}(3, 0, 0), \quad \mathbf{b}(0, 3, 0), \quad \mathbf{c}(0, 0, 3)$$




Fig. 1. A bipartite graph.




Fig. 2. Hypertree polytopes of a pair of abstract dual hypergraphs.

and

$$\mathbf{p}(2, 0, 0, 0), \quad \mathbf{q}(0, 2, 0, 0), \quad \mathbf{r}(0, 0, 2, 0), \quad \mathbf{s}(0, 0, 0, 2),$$

respectively. In Fig. 2, we indicated individual hypertrees by dots and the two hypertree polytopes by thickened edges.

The hypertree polytope is of course closed under convex linear combinations. For its integer lattice points, that is the hypertrees themselves, convexity translates to the following obvious (from Theorem 3.4) but useful property.

Lemma 3.10. *Let $\{\mathbf{f}_i\}$ be some set of hypertrees in the hypergraph $\mathcal{H} = (V, E)$ and let $\mathbf{g}: E \rightarrow \mathbf{N}$ be a function so that $\mathbf{g}(e) \geq 0$ for all $e \in E$ and $\sum_{e \in E} \mathbf{g}(e) = |V| - 1$. If it is also true that for all non-empty $E' \subset E$, there is an index i so that $\sum_{e \in E'} \mathbf{g}(e) \leq \sum_{e \in E'} \mathbf{f}_i(e)$, then \mathbf{g} is a hypertree.*

3.2. Postnikov’s approach

In the terminology of [4], a hypertree is a draconian sequence [4, Definition 9.2] and the hypertree polytope is a trimmed generalized permutohedron [4, Definition 11.2]. Let us now summarize Corollary 11.8, Theorem 12.9, and some other surrounding material from [4].

Theorem 3.11 (Postnikov). *Let $G = (V_0, V_1, E)$ be a connected bipartite graph with associated hypergraphs \mathcal{G}_0 and \mathcal{G}_1 as in (2). Then $|Q_{\mathcal{G}_0} \cap \mathbf{Z}^{V_0}| = |Q_{\mathcal{G}_1} \cap \mathbf{Z}^{V_1}|$. In other words, abstract dual hypergraphs have the same number of hypertrees.*

It turns out that $Q_{\mathcal{G}_0}$ and $Q_{\mathcal{G}_1}$ have even more in common. Conjecture 7.1 below proposes a generalization of Postnikov’s theorem.

Without going too deep into the details of an elegant but complicated proof, we will quote some of Postnikov’s observations. Of course, an equivalent statement is

$$|Q_{\mathcal{H}} \cap \mathbf{Z}^E| = |Q_{\overline{\mathcal{H}}} \cap \mathbf{Z}^V|,$$

where $\mathcal{H} = (V, E)$ is an arbitrary hypergraph and $\overline{\mathcal{H}} = (E, V)$ is its abstract dual. We will use this latter formulation. Denote the standard generators of \mathbf{R}^E with $\mathbf{i}_{\{e\}}$, $e \in E$. Recall that the *Minkowski sum* of two subsets A, B of a vector space is the set of vectors one can obtain as $\mathbf{a} + \mathbf{b}$ with $\mathbf{a} \in A$ and $\mathbf{b} \in B$. The *Minkowski difference* $A - B$ is defined as the set of vectors \mathbf{x} so that $\mathbf{x} + B \subset A$. It can also be thought of as the set of all translates of B that are contained in A . See [5, Chapter 3] for a thorough introduction to these operations.

Next, let us quote [4, Lemma 11.7].

Proposition 3.12 (Postnikov). *The hypertree polytope $Q_{\mathcal{H}} \subset \mathbf{R}^E$ of the hypergraph $\mathcal{H} = (V, E)$ can be written as*

$$Q_{\mathcal{H}} = \left(\sum_{v \in V} \Delta_v \right) - \Delta_E,$$

where Σ denotes the Minkowski sum of the simplices $\Delta_v = \text{Conv}\{\mathbf{i}_{\{e\}} \mid v \in e\}$, $\Delta_E = \text{Conv}\{\mathbf{i}_{\{e\}} \mid e \in E\}$ is the unit simplex, and the right hand side is a Minkowski difference.

Of course, if we denote the standard basis for \mathbf{R}^V with $\mathbf{i}_{\{v\}}$, $v \in V$ and define the simplices $\Delta_e = \text{Conv}\{\mathbf{i}_{\{v\}} \mid v \in e\} \subset \mathbf{R}^V$ for all $e \in E$, as well as $\Delta_V = \text{Conv}\{\mathbf{i}_{\{v\}} \mid v \in V\}$, then

$$Q_{\overline{\mathcal{H}}} = \left(\sum_{e \in E} \Delta_e \right) - \Delta_V.$$

The Minkowski sums (in Postnikov’s terminology, the (untrimmed) generalized permutohedra)

$$Q_{\mathcal{H}}^+ = \sum_{v \in V} \Delta_v \quad \text{and} \quad Q_{\overline{\mathcal{H}}}^+ = \sum_{e \in E} \Delta_e$$

are related through the so-called *root polytope*

$$Q = \text{Conv}\{\mathbf{i}_{\{e\}} + \mathbf{i}_{\{v\}} \mid v \in e\} \subset \mathbf{R}^E \oplus \mathbf{R}^V$$

(an $(|E| + |V| - 2)$ -dimensional polytope whose vertices are indexed by the edges of $\text{Bip } \mathcal{H}$) as described below.

Proposition 3.13. *If we define two affine subspaces of $\mathbf{R}^E \oplus \mathbf{R}^V$ with*

$$S_E = \pi_V^{-1} \left(\frac{1}{|V|} \sum_{v \in V} \mathbf{i}_{\{v\}} \right) \quad \text{and} \quad S_V = \pi_E^{-1} \left(\frac{1}{|E|} \sum_{e \in E} \mathbf{i}_{\{e\}} \right),$$

where $\pi_V: \mathbf{R}^E \oplus \mathbf{R}^V \rightarrow \mathbf{R}^V$ and $\pi_E: \mathbf{R}^E \oplus \mathbf{R}^V \rightarrow \mathbf{R}^E$ are the standard projections, then, up to translation,

$$Q_{\mathcal{H}}^+ = |V|(Q \cap S_E) \quad \text{and} \quad Q_{\overline{\mathcal{H}}}^+ = |E|(Q \cap S_V).$$

Then, the key idea is to study triangulations of Q . Postnikov observes that a set of vertices of Q is affinely independent if and only if the corresponding edges form a forest in $\text{Bip } \mathcal{H}$.

The maximal simplices T_Γ that arise this way from spanning trees Γ of $\text{Bip } \mathcal{H}$ have the same volume.

Most importantly, if we form the intersections $T_\Gamma \cap S_E$ and $T_\Gamma \cap S_V$ and identify them with subsets of $Q_{\mathcal{H}}^+$ and $Q_{\mathcal{H}^c}^+$, respectively, as described in Proposition 3.13, then those subsets contain unique translates of Δ_E and Δ_V , respectively, by integer vectors. (Also, $T_\Gamma \cap S_E$ and $T_\Gamma \cap S_V$ are essentially disjoint from other integer translates of those unit simplices.) In fact, those two integer vectors are the two hypertrees induced by Γ .

Thus Postnikov finds that any triangulation of Q gives a bijection between the integer lattice points of the Minkowski differences which are the hypertree polytopes. We also see that the volume of the root polytope associated to a bipartite graph is proportional to the number of hypertrees in each of its induced hypergraphs.

4. The geometry of the hypertree polytope

4.1. Polymatroids

It turns out that for our central concepts, namely the interior and exterior polynomials which will be introduced in the next section, the right level of generality is that of a polymatroid. Moreover, basic submodular function techniques will be useful to simplify our arguments, most notably the proof that our polynomials are well defined, even in the hypergraph case. We will recall some elements of this theory here, using [6, Volume B, Chapter 44] as basic reference.

Definition 4.1. Let S be a finite ground set and $\mu: \mathcal{P}(S) \rightarrow \mathbf{R}$ a set function, i.e., a function defined on all subsets of S . We say that μ is *submodular* if

$$\mu(U) + \mu(V) \geq \mu(U \cap V) + \mu(U \cup V)$$

holds for all subsets U, V of S . The set function μ is called *non-decreasing* if $U \subset V \subset S$ implies $\mu(U) \leq \mu(V)$.

For an arbitrary set function $v: \mathcal{P}(S) \rightarrow \mathbf{R}$, we define the polyhedra

$$P_v = \{ \mathbf{x} \in \mathbf{R}^S \mid \mathbf{x} \geq \mathbf{0}, \mathbf{x} \cdot \mathbf{i}_U \leq v(U) \text{ for all } U \subset S \}$$

$$\text{and } EP_v = \{ \mathbf{x} \in \mathbf{R}^S \mid \mathbf{x} \cdot \mathbf{i}_U \leq v(U) \text{ for all } U \subset S \}.$$

Here $\mathbf{x} \geq \mathbf{0}$ means that all entries in \mathbf{x} are non-negative and \mathbf{i}_U is the indicator function of the subset U so that the dot product $\mathbf{x} \cdot \mathbf{i}_U$ becomes the sum of entries in \mathbf{x} corresponding to elements of U . We also define

$$B_v = \{ \mathbf{x} \in EP_v \mid \mathbf{x} \cdot \mathbf{i}_S = v(S) \}$$

and note

Lemma 4.2. For any set function $v: S \rightarrow \mathbf{R}$, the set B_v is a polytope. If v is non-decreasing, then we have $B_v \subset P_v$.

Proof. It is obvious that B_v is convex. For boundedness, it suffices to show that for any $s \in S$, the s -components of elements of B_v are bounded below. (From that, the condition $\mathbf{x} \cdot \mathbf{i}_S = v(S)$ implies upper bounds.) Let $\mathbf{x} \in B_v$, $s \in S$, and $U = S \setminus \{s\}$. Then indeed, $\mathbf{x}(s) = v(S) - \mathbf{x} \cdot \mathbf{i}_U \geq v(S) - v(U)$.

If v is non-decreasing, then the previous lower bound satisfies $v(S) - v(U) \geq v(S) - v(S) = 0$ as well. \square

Definition 4.3. If $\mu: \mathcal{P}(S) \rightarrow \mathbf{R}$ is a submodular set function, then P_μ and EP_μ are called the *polymatroid* and *extended polymatroid*, respectively, of μ . A vector $\mathbf{x} \in B_\mu$ is called a *base* and B_μ itself is referred to as the *base polytope*.

The notions of submodular set function and extended polymatroid are essentially equivalent as described in [6, Section 44.4]. By the same token, non-decreasing submodular set functions are equivalent to polymatroids. Integer-valued submodular functions correspond to (extended) polymatroids that are *integer*, meaning that they are the convex hulls of their integer lattice points.

The following proposition expresses a basic property of polymatroids that will have an important role in our treatment.

Definition 4.4. Let $\nu: \mathcal{P}(S) \rightarrow \mathbf{R}$ be a set function and $\mathbf{x} \in EP_\nu$. We say that the subset $U \subset S$ is *tight* at \mathbf{x} if $\mathbf{x} \cdot \mathbf{i}_U = \nu(U)$.

The set function ν is called *tight* if for all $\mathbf{x} \in EP_\nu$, the set of subsets of S that are tight at \mathbf{x} is closed under taking unions and intersections.

Proposition 4.5 ([6], *Theorem 44.2*). *Submodular set functions are tight.*

Our next claim is probably well known to experts on polymatroids. It concerns certain 3-dimensional cross-sections of the base polytope.

Lemma 4.6. *Let $\mu: \mathcal{P}(S) \rightarrow \mathbf{R}$ be a submodular set function and let $S' = S \setminus \{p, q, r, s\}$, where p, q, r , and s are distinct elements of S . Fix the value of $\mathbf{x} \in B_\mu$ at $\nu(t)$ for all elements $t \in S'$ and denote the resulting subset of B_μ with B' . Then, assuming $B' \neq \emptyset$, the face of B' along which the sum of the p and q components takes its maximum is a (possibly degenerate) rectangle with sides parallel to the vectors $\mathbf{i}_{\{p\}} - \mathbf{i}_{\{q\}}$ and $\mathbf{i}_{\{r\}} - \mathbf{i}_{\{s\}}$.*

Proof. Let $\sigma = \mu(S) - \sum_{t \in S'} \nu(t)$ and let the 3-dimensional affine subspace $A \subset \mathbf{R}^S$ be defined by $A = \{\mathbf{x} \in \mathbf{R}^S \mid \mathbf{x} \cdot \mathbf{i}_S = \mu(S) \text{ and } \mathbf{x}(t) = \nu(t) \text{ for all } t \in S'\}$.

The set B' is cut out of A as the intersection of fourteen half-spaces corresponding to the non-trivial subsets of $\{p, q, r, s\}$. Six of these have the normal vectors $\mathbf{i}_{\{p\}} + \mathbf{i}_{\{q\}} - \mathbf{i}_{\{r\}} - \mathbf{i}_{\{s\}}$ and so on so that their intersection is a (possibly degenerate) cuboid C ; see Fig. 3. The other eight half-spaces have normal vectors such as $\pm(\mathbf{i}_{\{p\}} + \mathbf{i}_{\{q\}} + \mathbf{i}_{\{r\}} - 3\mathbf{i}_{\{s\}})$, and they cut off pieces of C near its vertices so that eventually B' is obtained. Our task is to show that in B' there cannot remain a segment of positive length from any edge of C .

If the opposite was the case and an interior point \mathbf{x} remained on, say, the edge L of C that is thickened in Fig. 3, then \mathbf{x} would be a base in B_μ at which

- the sets $\{p, q\}$ and $\{p, r\}$ are tight (this is because \mathbf{x} lies along both the top face and the right-side face of C , cf. the two upper panels of Fig. 3) but
- neither their union $\{p, q, r\}$ nor their intersection $\{p\}$ is tight (this is apparent from the fact that if we move along L away from the vertex \mathbf{p}' of C , the sum of the p, q , and r -coordinates increases, whereas if we move toward \mathbf{p}' , the p -coordinate increases).

As this is impossible, we see that the correct form of B' is as shown in the lower panel of Fig. 3. \square

Submodular functions are relevant in this paper due to the following fact, pointed out to the author by László Lovász.




Fig. 3. Upper left: the portion of the simplex $\mu(S)\Delta_S$ that lies in the affine subspace A . It is a superset of B' and it is pictured here to establish the coordinate directions. Writing them in the order p, q, r, s , the four non-fixed components of the vertices shown are $\mathbf{p}(\sigma, 0, 0, 0)$ and so on. Upper right: the cuboid C , the normal vectors of its faces, as well as an edge L of C and two more supporting planes of B' in an arrangement that we wish to rule out. Bottom: the polytope $B' = B_\mu \cap A$ and the direction vectors of two of its edges.

Proposition 4.7. Let $\mathcal{H} = (V, E)$ be a hypergraph. The function

$$\mu(E') = \left| \bigcup E' \right| - c(E') \tag{4}$$

of Theorem 3.4, extended to the empty set as $\mu(\emptyset) = 0$, is non-decreasing and submodular on the set E of hyperedges.

Proof. Let $E' \subset E$ and let $e \notin E'$ be a hyperedge. If in $\text{Bip } \mathcal{H}$, the vertex e is connected to m components of $\text{Bip } \mathcal{H}|_{E'}$ and also to n elements of V that are not in $\bigcup E'$, then we have

$$\left| \bigcup (E' \cup \{e\}) \right| - \left| \bigcup E' \right| = n \quad \text{and} \quad c(E' \cup \{e\}) - c(E') = -(m - 1),$$

i.e., $\mu(E' \cup \{e\}) - \mu(E') = n + m - 1$. Here $n, m \geq 0$ and since e is non-empty, at least one of the two is actually positive. This implies that μ is non-decreasing.

As to the submodularity of μ , according to [6, Theorem 44.1], it suffices to prove

$$(\mu(E' \cup \{e_1\}) - \mu(E')) + (\mu(E' \cup \{e_2\}) - \mu(E')) \geq \mu(E' \cup \{e_1, e_2\}) - \mu(E') \tag{5}$$

for all $E' \subset E$ and distinct $e_1, e_2 \in E \setminus E'$. Assume that from the connected components of $\text{Bip } \mathcal{H}|_{E'}$, the number of those that are connected to both e_1 and e_2 is m_{12} . Aside from these, let $e_i, i = 1, 2$ be connected to another m_i components of $\text{Bip } \mathcal{H}|_{E'}$. Furthermore,

let us denote the number of such common elements of e_1 and e_2 that are not in $\bigcup E'$ with n_{12} and assume that other than these, e_i contains n_i non-elements of $\bigcup E'$. Then we have

$$\begin{aligned} \mu(E' \cup \{e_1\}) - \mu(E') &= n_1 + n_{12} + m_1 + m_{12} - 1 \text{ and} \\ \mu(E' \cup \{e_2\}) - \mu(E') &= n_2 + n_{12} + m_2 + m_{12} - 1, \end{aligned}$$

whereas the right hand side of (5) is

$$\mu(E' \cup \{e_1, e_2\}) - \mu(E') = \begin{cases} n_1 + n_2 + m_1 + m_2 - 2 & \text{if } n_{12} = m_{12} = 0 \\ n_1 + n_2 + n_{12} + m_1 + m_2 + m_{12} - 1 & \text{otherwise.} \end{cases}$$

The required inequality holds in all cases (with equality if and only if $n_{12} = m_{12} = 0$ or one of the two is 1 and the other 0). \square

Definition 4.8. The polymatroid $P_{\mathcal{H}} = P_{\mu}$ associated to the non-decreasing submodular set function μ of equation (4) is called the *polymatroid of the hypergraph \mathcal{H}* .

Each graph has a so-called cycle matroid. This is a classical construction that goes back to the very beginning of matroid theory. The rank function of any matroid is submodular and therefore defines a polymatroid. (In fact, one way to define a matroid is as a polymatroid associated to a non-negative integer-valued, non-decreasing submodular set function that assigns 0 to \emptyset and 0 or 1 to singleton sets.) Thus every graph has a natural polymatroid associated to it, and Definition 4.8 generalizes this association to hypergraphs.

The next proposition is a reformulation of Theorem 3.4 (and Lemma 3.7, which was used in its proof).

Proposition 4.9. *The integer lattice points in the polymatroid $P_{\mathcal{H}}$ of the hypergraph \mathcal{H} are exactly the functions $\mathbf{f}: E \rightarrow \mathbf{N}$ so that $\text{Bip } \mathcal{H}$ has a cycle-free subgraph with valence $\mathbf{f}(e) + 1$ at each $e \in E$. The hypertree polytope $Q_{\mathcal{H}}$ coincides with the base polytope of $P_{\mathcal{H}}$.*

It is well known [6, Section 44.6c] that polymatroids defined by integer-valued non-decreasing submodular functions are integer polytopes, that is to say, convex hulls of finitely many integer lattice points. In particular, their base polytopes are integer as well. Hence we have

Corollary 4.10. *For any hypergraph \mathcal{H} , the hypertree polytope $Q_{\mathcal{H}}$ is integer.*

The next subsection, among other things, contains a description of the vertices of the hypertree polytope.

4.2. Hyperedges in sequence

We are going to need the following elementary observations later.

Lemma 4.11. *Let the spanning tree $\Gamma \subset \text{Bip } \mathcal{H}$ induce the hypertree $\mathbf{f}: E \rightarrow \mathbf{N}$. The inequality (3b) is sharp for $E' \subset E$ and \mathbf{f} (in other words, E' is tight at \mathbf{f}) if and only if $\Gamma \cap (\text{Bip } \mathcal{H}|_{E'})$ is a spanning forest in $\text{Bip } \mathcal{H}|_{E'}$. Given such a Γ and E' , if we remove $\Gamma \cap (\text{Bip } \mathcal{H}|_{E'})$ from Γ and replace it with another spanning forest F of $\text{Bip } \mathcal{H}|_{E'}$, then the result $\tilde{\Gamma}$ is another spanning tree of $\text{Bip } \mathcal{H}$. It realizes a hypertree which agrees with \mathbf{f} on the set $E \setminus E'$.*

Proof. The first assertion is trivial from the proof of Theorem 3.4. For the second, notice that paths in Γ and $\tilde{\Gamma}$ are in a one-to-one correspondence, as follows. Any path $\varphi \subset \Gamma$ has maximal

segments that fall within $\text{Bip } \mathcal{H}|_{E'}$ and obviously each such segment is in one connected component of $\text{Bip } \mathcal{H}|_{E'}$. Hence, we get a new path $\tilde{\varphi} \subset \tilde{T}$ by replacing each segment with the unique connection that exists between its endpoints in F . This correspondence of paths has an inverse constructed in the analogous way. As φ and $\tilde{\varphi}$ share the same endpoints, we see that \tilde{T} is connected and cycle-free. The final claim of the lemma is of course just stating the obvious. \square

Now suppose that the set E of hyperedges in $\mathcal{H} = (E, V)$ has been ordered and labeled so that

$$e_1 < e_2 < \dots < e_{|E|}.$$

Think of \mathcal{H} (and of $\text{Bip } \mathcal{H}$) as being built step-by-step by adding one hyperedge at a time in the prescribed order. We will use the notation

$$G_k = \text{Bip } \mathcal{H}|_{\{e_1, \dots, e_k\}}$$

for the bipartite graphs of the intermediate stages. In each step, the nullity of the graph may increase. We record this by introducing the *nullity-jump function* of the chosen order,

$$nj(e_k) = n(G_k) - n(G_{k-1}),$$

where $k = 1, 2, \dots, |E|$. Obviously, $nj(e) \geq 0$ for all e and $\sum_{e \in E} nj(e) = n(\text{Bip } \mathcal{H})$.

Lemma 4.12. *For any order of the hyperedges, the vector $\mathbf{g}(e) = |e| - 1 - nj(e)$ is a hypertree. It is the unique hypertree that makes the inequality (3b) true with an equality sign for all subsets $\{e_1, \dots, e_k\}$, $k = 1, 2, \dots, |E|$.*

Proof. Using the order, it is easy to construct a spanning tree that realizes \mathbf{g} . The k 'th stage of the construction will be a spanning forest F_k of the bipartite graph G_k . All edges of $\text{Bip } \mathcal{H}$ that are adjacent to e_1 will be part of F_1 (note that $nj(e_1) = 0$). Suppose that a forest $F_{k-1} \subset G_{k-1}$, with the valences $\mathbf{g}(e_i) + 1$ at e_1, \dots, e_{k-1} , respectively, has been defined.

The graph $G_k - \{e_k\}$ consists of G_{k-1} and some isolated points. Suppose that e_k is joined by an edge to c connected components of $G_k - \{e_k\}$. It is easy to see that $nj(e_k) = |e_k| - c$, since the ‘first edge’ of G_k to connect e_k to one of the c components does not increase the nullity, whereas all others after the first increase it by 1. We define F_k by adding to F_{k-1} a collection of $c = |e_k| - nj(e_k) = \mathbf{g}(e_k) + 1$ edges, all adjacent to e_k and leading to different components of $G_k - \{e_k\}$. This is a spanning forest of G_k . After the last stage, $F_{|E|}$ is a spanning tree of $G_{|E|} = \text{Bip } \mathcal{H}$ which realizes \mathbf{g} .

From the fact that F_k is a spanning forest of G_k and the first claim in Lemma 4.11, it is immediate that

$$\text{the inequality (3b) is sharp for } \mathbf{g} \text{ and each set } E' = \{e_1, \dots, e_k\}. \tag{6}$$

No other hypertree can have that property simply because

$$\mathbf{g}(e_k) = \mu(\{e_1, \dots, e_k\}) - \mu(\{e_1, \dots, e_{k-1}\}) \tag{7}$$

is uniquely determined as the difference of two consecutive right-hand sides. \square

Remark 4.13. From (6), we see that the greedy algorithm offers another way of defining the hypertree \mathbf{g} of the previous lemma: if for all $k = 1, \dots, |E|$, we choose $\mathbf{f}(e_k)$ to be the maximal value so that the valences $\mathbf{f}(e_i) + 1, i = 1, \dots, k$, can be realized by a cycle-free subgraph of G_k , then $\mathbf{f} = \mathbf{g}$.

We will also refer to the hypertree \mathbf{g} of Lemma 4.12 as the *greedy hypertree* of the given order.

Proposition 4.14. *The set of greedy hypertrees for all possible orders of S is exactly the vertex set of \mathcal{QH} .*

Proof. Compare (7) to [6, Equation (44.49)]. \square

The next notion will be fundamental in the rest of the paper.

Definition 4.15. Let $\mathcal{H} = (V, E)$ be a hypergraph, $\mathbf{f}: E \rightarrow \mathbb{N}$ a hypertree and $a, b \in E$ hyperedges. We say that \mathbf{f} is such that a *transfer of valence* is possible from a to b if the function obtained from \mathbf{f} by reducing $\mathbf{f}(a)$ by 1 and increasing $\mathbf{f}(b)$ by 1 is also a hypertree.

Lemma 4.16. *For any non-empty collection $E' \subset E$ of hyperedges, there exists a hypertree $\mathbf{f}: E \rightarrow \mathbb{N}$ so that (3b) is sharp for \mathbf{f} and E' . If E' and \mathbf{f} are so that (3b) is not sharp, then \mathbf{f} is such that it is possible to transfer valence from some element of $E \setminus E'$ to some element of E' .*

Proof. Choose any order in which the elements of E' are the smallest and construct its greedy hypertree $\mathbf{g}: E \rightarrow \mathbb{N}$. By Lemma 4.12, \mathbf{g} has the required property.

To prove the second claim, we give an indirect argument using Propositions 4.7 and 4.5. Let \mathbf{f} be a hypertree and $E' \subset E$ a subset so that no transfer of valence is possible from an element of $E \setminus E'$ to an element of E' . By Theorem 3.4, this implies that for any $a \in E'$ and $b \in E \setminus E'$, there exists a subset $U_{a,b} \subset E$ of hyperedges so that with the set function μ of (4), we have

$$a \in U_{a,b}, \quad b \notin U_{a,b}, \quad \text{and} \quad \sum_{x \in U_{a,b}} \mathbf{f}(x) = \mu(U_{a,b}).$$

In other words, $U_{a,b}$ is tight at \mathbf{f} . Then so is $E' = \bigcup_{a \in E'} \left(\bigcap_{b \in E \setminus E'} U_{a,b} \right)$, i.e., (3b) is sharp for E' and \mathbf{f} . \square

Lemma 4.17. *Let $\mathcal{H} = (V, E)$ be a hypergraph with its hyperedges ordered as above. Let $\mathbf{f}: E \rightarrow \mathbb{N}$ be a hypertree and $e \in E$ a hyperedge so that $\mathbf{f}(e) > \mathbf{g}(e)$, where \mathbf{g} is the greedy hypertree of the order. Then there exists a hyperedge $x < e$ so that \mathbf{f} is such that a transfer of valence is possible from e to x .*

Proof. Assume the contrary. Then for all $x < e$, there exists a set $U_{e,x}$ of hyperedges that is tight at \mathbf{f} , contains x , and does not contain e . Let $U = \bigcup_{x < e} U_{e,x}$. Then we have $e \notin U$ and $\{x \mid x < e\} \subset U$ while, by Proposition 4.5, U is also tight at \mathbf{f} .

Now define a new order $<'$ on E by letting the elements of U be smallest while keeping the original order among them. (It does not matter how the rest of E gets ordered.) Let \mathbf{g}' be the greedy hypertree of $<'$. By Lemma 4.12, both $\{x \mid x < e\}$ and U are tight at \mathbf{g}' . Fix realizations Γ' of \mathbf{g}' and Γ of \mathbf{f} . By Lemma 4.11, these are such that both $\Gamma \cap (\text{Bip } \mathcal{H}|_U)$ and $\Gamma' \cap (\text{Bip } \mathcal{H}|_U)$ are spanning forests in $\text{Bip } \mathcal{H}|_U$. If in Γ we replace the former with the latter, the result, again by Lemma 4.11, is a spanning tree in $\text{Bip } \mathcal{H}$. Now for the hypertree \mathbf{f}' realized by this spanning tree, we have

$$\begin{aligned} \sum_{x \leq e} \mathbf{f}'(x) &= \mathbf{f}'(e) + \sum_{x < e} \mathbf{f}'(x) = \mathbf{f}(e) + \sum_{x < e} \mathbf{g}'(x) = \mathbf{f}(e) + \sum_{x < e} \mathbf{g}(x) \\ &> \mathbf{g}(e) + \sum_{x < e} \mathbf{g}(x) = \sum_{x \leq e} \mathbf{g}(x) = \mu(\{x \mid x \leq e\}), \end{aligned}$$

which contradicts Theorem 3.4. \square




Fig. 4. A rhombus of hypertrees.

5. Interior and exterior polynomials

5.1. Activities

Let $\mathcal{H} = (V, E)$ be a hypergraph so that $\text{Bip } \mathcal{H}$ is connected. Just as in Sections 2.1 and 4.2, we order the set E arbitrarily. With regard to a fixed hypertree \mathbf{f} , we make the following definitions.

Definition 5.1. A hyperedge $e \in E$ is *internally active* with respect to the hypertree \mathbf{f} if it is not possible to decrease $\mathbf{f}(e)$ by 1 and increase \mathbf{f} at a hyperedge smaller than e by 1 so that another hypertree results.

A hyperedge $e \in E$ is *externally active* with respect to \mathbf{f} if it is not possible to increase $\mathbf{f}(e)$ by 1 and to decrease \mathbf{f} at a smaller hyperedge by 1 so that another hypertree results.

Recall that in Definition 4.15, the operation on hypertrees used above was called a transfer of valence. Hence a hyperedge is internally active with respect to a hypertree if it cannot transfer valence to smaller hyperedges, and it is externally active if valence cannot be transferred to it from smaller hyperedges. Regarding transfers of valence, we will need the following lemma later.

Lemma 5.2. *Let $a, b,$ and c be hyperedges in $\mathcal{H} = (V, E)$ and \mathbf{f} such a hypertree that a can transfer valence to b and b can transfer valence to c . Then \mathbf{f} is also such that a can transfer valence to c . In other words, regarding the rhombus of Fig. 4 (explained below), if the three lattice points indicated by full dots are hypertrees then so is the one marked by the hollow dot.*

Proof. This is immediate from Lemma 3.10 but we think it worthwhile to examine the picture. In \mathbf{R}^E , the three-dimensional affine subspace through \mathbf{f} spanned by $\{\mathbf{i}_{\{a\}}, \mathbf{i}_{\{b\}}, \mathbf{i}_{\{c\}}\}$ forms a two-dimensional intersection Q_0 with $Q_{\mathcal{H}}$. The triangle we labeled with \mathbf{abc} in Fig. 4 is in fact spanned by the vectors whose $a, b,$ and c -components are equal to $(\mathbf{f}(a) + \mathbf{f}(b) + \mathbf{f}(c), 0, 0), (0, \mathbf{f}(a) + \mathbf{f}(b) + \mathbf{f}(c), 0),$ and $(0, 0, \mathbf{f}(a) + \mathbf{f}(b) + \mathbf{f}(c)),$ respectively, whereas their other components are identical to those of \mathbf{f} .

The assumptions in the lemma are that $\mathbf{f}, \mathbf{f}_1,$ and \mathbf{f}_2 are in $Q_{\mathcal{H}}$ and the conclusion is that so is $\hat{\mathbf{f}}$. Now this is obvious from the fact (Theorem 3.4) that Q_0 is cut out from the plane by lines parallel to the sides of the triangle. \square

It is easy to see that the notions of Definition 5.1 generalize the ones found in Definition 2.1. It will be slightly more convenient for us to work with their negations in what follows.




Fig. 5. Two ‘staples’ formed by hypertrees.

Definition 5.3. Let $\mathcal{H} = (V, E)$ be a hypergraph so that $\text{Bip } \mathcal{H}$ is connected. With respect to a given hypertree $\mathbf{f}: E \rightarrow \mathbf{N}$ and some order of the hyperedges, let the number of internally inactive hyperedges of \mathcal{H} be denoted with $\bar{i}(\mathbf{f})$ and the number of externally inactive hyperedges be denoted with $\bar{e}(\mathbf{f})$. Then, let the *interior polynomial* and the *exterior polynomial* of \mathcal{H} be defined as

$$I_{\mathcal{H}}(\xi) = \sum_{\mathbf{f}} \xi^{\bar{i}(\mathbf{f})} \quad \text{and} \quad X_{\mathcal{H}}(\eta) = \sum_{\mathbf{f}} \eta^{\bar{e}(\mathbf{f})}, \tag{8}$$

respectively, where both summations are over all hypertrees \mathbf{f} in \mathcal{H} .

These notions generalize the valuations

$$\xi^{|V|-1} T_G(1/\xi, 1) \quad \text{and} \quad \eta^{|E|-|V|+1} T_G(1, 1/\eta), \tag{9}$$

respectively, of the classical Tutte polynomial T_G of the graph $G = (V, E)$.

There are two ways of treating hypergraphs when their associated bipartite graphs are disconnected. One is to assign to them the product of the polynomials associated to their connected components. The other is to extend the definition verbatim, which results in the value 0 for both polynomials because a disconnected graph has no spanning trees and therefore the set of hypertrees is empty. In this paper we take the latter approach. This will hardly matter since we almost always assume $\text{Bip } \mathcal{H}$ to be connected.

Our first order of business, however, is to address well definedness.

Theorem 5.4. *The formulas (8) for the interior and exterior polynomials do not depend on the chosen order of the hyperedges.*

We will need the following statement.

Lemma 5.5. *Let $e_1 \neq e_2$ be hyperedges in the hypergraph $\mathcal{H} = (V, E)$ and let $E_0 = E \setminus \{e_1, e_2\}$. Fix a function $\mathbf{f}_0: E_0 \rightarrow \mathbf{N}$ and consider its extensions to E that are hypertrees in \mathcal{H} . Among these, let \mathbf{f}_1 and \mathbf{f}_2 be such that $\mathbf{f}_1(e_1) > \mathbf{f}_2(e_1)$ and let also $x \in E_0$ be a hyperedge.*

1. *If \mathbf{f}_1 is such that valence can be transferred from e_1 to x then \mathbf{f}_2 is such that valence can be transferred from e_2 to x .*
2. *If \mathbf{f}_1 is such that valence can be transferred from x to e_2 then \mathbf{f}_2 is such that valence can be transferred from x to e_1 .*




Fig. 6. Hypertrees with associated $(\bar{l}, \bar{\varepsilon})$ values in a pair of abstract dual hypergraphs. Hollow shapes are internally inactive and full ones are internally active; circles are externally inactive while squares are externally active.

In other words, regarding both ‘staples’ of Fig. 5, if the three lattice points indicated by full dots are hypertrees then so is the fourth one represented by the hollow dot.

Proof. The argument is similar to the proof of Lemma 5.2. From Fig. 5, it is easy to see that if $\mathbf{f}_1, \mathbf{f}_2,$ and \mathbf{g}_1 satisfy all constraints of Theorem 3.4, then so does $\hat{\mathbf{g}}_1$. Likewise, if $\mathbf{f}_1, \mathbf{f}_2,$ and \mathbf{g}_2 are in $Q_{\mathcal{H}}$ then so is $\hat{\mathbf{g}}_2$. \square

Proof of Theorem 5.4. We will emulate Tutte’s original proof, i.e., analyze the effect of changing the relative position in the order of two adjacent hyperedges. Let $a, b \in E$ and assume that $<_1$ is an order on E so that $a <_1 b$ with no other hyperedge in between, whereas the order $<_2$ only differs from $<_1$ in that $b <_2 a$. The rest of the hyperedges are split in two sets so that in both orders, elements of E_- are smaller than a and b while elements of E_+ are larger than both.

Let $\mathbf{f}: E \rightarrow \mathbf{N}$ be a hypertree. Our goal is to compare the values $\bar{l}_1(\mathbf{f})$ and $\bar{l}_2(\mathbf{f})$ (as well as $\bar{\varepsilon}_1(\mathbf{f})$ and $\bar{\varepsilon}_2(\mathbf{f})$) resulting from the two orders. If a hyperedge differs from both a and b , then it is easy to check that (in the internal as well as in the external sense) it is active with respect to \mathbf{f} in $<_1$ if and only if the same holds in $<_2$. Let us now separate three cases.

- I. If \mathbf{f} is such that no transfer of valence is possible between a and b , then the activity statuses of a and b are also unaffected by the change in the order. Hence in such cases $\bar{l}_1(\mathbf{f}) = \bar{l}_2(\mathbf{f})$ and $\bar{\varepsilon}_1(\mathbf{f}) = \bar{\varepsilon}_2(\mathbf{f})$.
- II. Next, assume that \mathbf{f} is such that valence can be transferred between a and b in both directions. Then in the order $<_1$ the hyperedge b is not active whereas with respect to $<_2$, the hyperedge a is not active. (This holds in both the internal and external senses.) Now according to Lemma 5.2, if a is not active in $<_1$ (i.e., there is some $x \in E_-$ so that valence can be transferred from a to x (internal case) or from x to a (external case)), then b is not active in $<_2$ and vice versa. Therefore we again have $\bar{l}_1(\mathbf{f}) = \bar{l}_2(\mathbf{f})$ and $\bar{\varepsilon}_1(\mathbf{f}) = \bar{\varepsilon}_2(\mathbf{f})$.
- III. Lastly, let \mathbf{f} be such that valence can be transferred between a and b but only in one direction. Before analyzing activities, let us establish that these kinds of hypertrees are grouped into pairs. Indeed, the line

$$\left\{ \mathbf{g}: E \rightarrow \mathbf{R} \mid \mathbf{g}|_{E \setminus \{a,b\}} = \mathbf{f}|_{E \setminus \{a,b\}}, \sum_{e \in E} \mathbf{g}(e) = |V| - 1 \right\} \subset \mathbf{R}^E$$

intersects the polytope $Q_{\mathcal{H}}$ in a segment. Among the lattice points of the segment, \mathbf{f} represents one extreme; let the other extreme be denoted with \mathbf{f}^* .¹ This is necessarily different from \mathbf{f} by our assumption and if a transfer of valence is possible only from a to b (and not in the other direction) at one of \mathbf{f} and \mathbf{f}^* , then the opposite transfer (and only that) is possible at the other one. It is also clear that $(\mathbf{f}^*)^* = \mathbf{f}$, so we get the desired pairs. (The lattice points between \mathbf{f} and \mathbf{f}^* were discussed in case II, whereas case I deals with the degenerate situation when $\mathbf{f} = \mathbf{f}^*$.)

We are going to study \bar{t} and \bar{e} values associated to \mathbf{f} and \mathbf{f}^* simultaneously. Without loss of generality we may assume that \mathbf{f} is such that a can transfer valence to b and the opposite transfer is possible at \mathbf{f}^* . Let us first examine the hyperedges a and b themselves. We will start with their internal activities.

- (a) Assume that $\bar{t}_1(\mathbf{f}) = \bar{t}_2(\mathbf{f})$. We will show that in this case $\bar{t}_1(\mathbf{f}^*) = \bar{t}_2(\mathbf{f}^*)$ holds, too. As the activity status of b with respect to \mathbf{f} is the same in \langle_1 as in \langle_2 , a similar property has to hold for a . Since a is not active with respect to \mathbf{f} in \langle_2 , this is only possible if there exists a hyperedge $x \in E_-$ so that \mathbf{f} is such that valence can be transferred from a to x . By Lemma 5.5, this implies that b is inactive with respect to \mathbf{f}^* in both \langle_1 and \langle_2 . Whether a is active with respect to \mathbf{f}^* depends on the same thing in both orders: namely, on whether \mathbf{f}^* is such that valence can be transferred from a to some element of E_- .
- (b) The only way $\bar{t}_1(\mathbf{f})$ and $\bar{t}_2(\mathbf{f})$ can be different is if \mathbf{f} is such that a cannot transfer valence to any element of E_- . In such cases b has the same property by Lemma 5.2 so that $\bar{t}_1(\mathbf{f}) = \bar{t}_2(\mathbf{f}) - 1$ (b is active with respect to \mathbf{f} in both orders and a is active only in \langle_1). Examining \mathbf{f}^* now, by Lemma 5.5 we see that it has to be such that b cannot transfer valence to any element of E_- . This implies, using Lemma 5.2, that a has the same property. Therefore we have $\bar{t}_1(\mathbf{f}^*) = \bar{t}_2(\mathbf{f}^*) + 1$ as a is active with respect to \mathbf{f}^* in both orders and b is active in \langle_2 only.

External activities can be handled in almost the same way.

- (a') If $\bar{e}_1(\mathbf{f}) = \bar{e}_2(\mathbf{f})$, then (since the activity status of a with respect to \mathbf{f} is independent of order) there must be some $x \in E_-$ so that valence can be transferred from x to b . Then by Lemma 5.5, \mathbf{f}^* is such that valence can be transferred from x to a , making a inactive with respect to \mathbf{f}^* in both orders. Since the activity status of b with respect to \mathbf{f}^* is the same in both orders, we obtain $\bar{e}_1(\mathbf{f}^*) = \bar{e}_2(\mathbf{f}^*)$.
- (b') If $\bar{e}_1(\mathbf{f}) \neq \bar{e}_2(\mathbf{f})$, that is if \mathbf{f} is such that no $x \in E_-$ can transfer valence to b , then the same holds for transfers of valence from x to a by Lemma 5.2. Therefore a is active with respect to \mathbf{f} in both orders, whereas b is active in \langle_2 and inactive in \langle_1 . By the preceding paragraph, a similar analysis has to apply to \mathbf{f}^* , i.e., b is active with respect to \mathbf{f}^* in both orders, whereas a is active in \langle_1 and inactive in \langle_2 . So we have obtained $\bar{e}_1(\mathbf{f}) = \bar{e}_2(\mathbf{f}) + 1$ and $\bar{e}_1(\mathbf{f}^*) = \bar{e}_2(\mathbf{f}^*) - 1$.

The only hypertrees whose \bar{t} or \bar{e} values actually changed as a result of switching from \langle_1 to \langle_2 were described in the cases (b) and (b') above. We saw that they occur in pairs of the form $\{\mathbf{f}, \mathbf{f}^*\}$. Now to complete our proof it suffices to show that for such a pair and any hyperedge y different from a and b , the activity status of y with respect to \mathbf{f} is the same as with respect to \mathbf{f}^* . (We saw that the choice between \langle_1 and \langle_2 does not matter for y , only for a and b . Switching

¹ The points \mathbf{f} and \mathbf{f}^* are actually the endpoints of the segment. This is easy to see from Theorem 3.4 and the fact that for any subset of the coordinates, the sum of those coordinates is either constant along the segment or takes integer values that change by 1 between neighboring lattice points.

to a new hypertree, on the other hand, could in principle make a big difference.) Indeed, that will imply $\bar{\iota}_1(\mathbf{f}) = \bar{\iota}_2(\mathbf{f}^*)$ and $\bar{\iota}_2(\mathbf{f}) = \bar{\iota}_1(\mathbf{f}^*)$ in case (b), as well as $\bar{\varepsilon}_1(\mathbf{f}) = \bar{\varepsilon}_2(\mathbf{f}^*)$ and $\bar{\varepsilon}_2(\mathbf{f}) = \bar{\varepsilon}_1(\mathbf{f}^*)$ in case (b') so that, regardless of order, the generating functions $I_{\mathcal{H}}$ and $X_{\mathcal{H}}$ always encode the same information.

As we are about to relate activities with respect to different hypertrees, this last part of the proof is where we will rely most heavily on our assumptions. What we need is essentially [Lemma 4.6](#) but we chose to spell the argument out without an explicit reference to it. It will be convenient to branch into several cases again.

(1) We will deal with internal activities first. Recall that we are under the assumptions of the case III(b), in particular \mathbf{f} and \mathbf{f}^* are such that neither a nor b can transfer valence to any element of E_- . Let y be a hyperedge which is internally inactive with respect to \mathbf{f} . Our goal is to show that y is also internally inactive with respect to \mathbf{f}^* .

i. Assume $y \in E_+$. If \mathbf{f} is such that y can transfer valence to a or b , then [Lemma 5.2](#) implies that it can definitely transfer to b and then [Lemma 5.5](#) says that \mathbf{f}^* is such that y can transfer valence to a . Therefore y is inactive with respect to \mathbf{f}^* .

If $y \in E_+$ but \mathbf{f} is such that y can transfer valence neither to a nor to b , then \mathbf{f}^* has to have the same property by the usual combination of [Lemmas 5.2](#) and [5.5](#). Because y is inactive with respect to \mathbf{f} , the hypertree \mathbf{f} must be such that y can transfer valence to some $x < y$ with $a \neq x \neq b$. We will show that \mathbf{f}^* is also such that y can transfer valence to x . Suppose the opposite is true. Then there are sets of hyperedges $U_a \ni a$, $U_b \ni b$, and $U_x \ni x$ so that all three are tight at \mathbf{f}^* and none of the three contains y . (Here tightness is in the sense of [Definition 4.4](#) with regards the set function μ of [Proposition 4.7](#).) By [Proposition 4.5](#), the union of the three subsets is also tight at \mathbf{f}^* . Since the union contains both a and b , the sum of the \mathbf{f}^* -values over it agrees with the sum of the \mathbf{f} -values. Therefore $U_a \cup U_b \cup U_x$ is also tight at \mathbf{f} , which contradicts our assumption that \mathbf{f} is such that y (which is not in the union) can transfer valence to x (which is).

ii. Let now $y \in E_-$. Then of course there is a hyperedge $x \in E_-$ so that \mathbf{f} is such that y can transfer valence to x . We will show that \mathbf{f}^* is also such that y can transfer valence to x . Assuming the contrary, we find sets of hyperedges $U_y \not\ni y$, $U_a \not\ni a$, and $U_b \not\ni b$ which are tight at \mathbf{f}^* and contain x . Their intersection is also tight at \mathbf{f}^* by [Proposition 4.5](#). As $U_a \cap U_b \cap U_y$ contains neither a nor b , it is also tight at \mathbf{f} . But that is a contradiction with the transfer of valence from $y \notin U_a \cap U_b \cap U_y$ to $x \in U_a \cap U_b \cap U_y$ which is possible at \mathbf{f} .

(2) In the external case, the assumptions of III(b') were that \mathbf{f} and \mathbf{f}^* are both such that no transfer of valence is possible from elements of E_- to a or to b . In such a situation, let y be externally inactive with respect to \mathbf{f} . We wish to prove that y is also externally inactive with respect to \mathbf{f}^* .

i'. If $y \in E_+$, then it may be that \mathbf{f} is such that a or b can transfer valence to y . If b can, then so can a by [Lemma 5.2](#). Thus in either case we can apply [Lemma 5.5](#) to conclude that \mathbf{f}^* is such that b can transfer valence to y .

Suppose now that $y \in E_+$ but \mathbf{f} (and consequently \mathbf{f}^*) is such that neither a nor b can transfer valence to y . Then some other hyperedge $x < y$ can and we will prove that the same transfer is also possible at \mathbf{f}^* . If not, then there are sets of hyperedges $U_a \not\ni a$, $U_b \not\ni b$, and $U_x \not\ni x$, all containing y , that are tight at \mathbf{f}^* . Their intersection has the same property by [Proposition 4.5](#) and moreover, as it contains neither a nor b , it is tight at \mathbf{f} as well. That contradicts the assumption that at \mathbf{f} , a transfer of valence is possible from outside of the set (from x) to one of its elements (namely y).

ii'. Assuming $y \in E_-$, it being externally inactive implies the existence of another hyperedge $x \in E_-$ so that \mathbf{f} is such that x can transfer valence to y . If the same transfer is not possible at \mathbf{f}^* , that means that there are sets of hyperedges $U_y \ni y, U_a \ni a$, and $U_b \ni b$ that are tight at \mathbf{f}^* and none of them contains x . Then their union has the same properties and it is also tight at \mathbf{f} , which yields the usual contradiction.

This completes the proof. \square

Example 5.6. We revisit the hypertrees that were discussed in [Example 3.9](#). They are shown in [Fig. 6](#) with an actual spanning tree realization for each. For the orders $a < b < c$ and $p < q < r < s$, respectively, we also indicated the internal and external activity status of each hyperedge with respect to each hypertree, as well as the resulting (\bar{l}, \bar{e}) values. Thus, we find that

$$I_{\mathcal{G}_0}(\xi) = 1 + 3\xi + 3\xi^2 \quad \text{and} \quad X_{\mathcal{G}_0}(\eta) = 1 + 3\eta + 3\eta^2, \tag{10}$$

whereas

$$I_{\mathcal{G}_1}(\xi) = 1 + 3\xi + 3\xi^2 \quad \text{and} \quad X_{\mathcal{G}_1}(\eta) = 1 + 2\eta + 3\eta^2 + \eta^3. \tag{11}$$

These values are in line with [Proposition 6.2](#) and [Theorem 6.3](#) below, as well as with [Conjecture 7.1](#).

Up to isomorphism there is only one order on the three-element set V_0 . The four-element set V_1 , on the other hand, has four essentially different orders depending on the position of q . The reader may check that the other three give rise to the same interior and exterior polynomials as in (11).

Note the lack of similarity between our polynomials and the one shown in [Example 2.4](#). Indeed, the author is not aware of any formula relating the Tutte polynomial $T_G(x, y)$ of a bipartite graph G and the interior and exterior polynomials of its induced hypertrees.

Remark 5.7. Unfortunately, for a hypergraph that is not a graph, the two-variable polynomial $\sum_{\mathbf{f}} \xi^{\bar{l}(\mathbf{f})} \eta^{\bar{e}(\mathbf{f})}$ (cf. [Definition 2.2](#)) does depend on the choice of order. For instance, for the hypergraph \mathcal{G}_1 of the previous example, the order that we used there gives rise to the two-variable generating function $\eta^3 + 3\xi\eta^2 + 2\xi^2\eta + \xi^2$ (this is directly read off of the bottom row of [Fig. 6](#)), whereas the order $p < r < s < q$ gives $\xi\eta^3 + \xi^2\eta^2 + \xi\eta^2 + \eta^2 + 2\xi^2\eta + \xi$. (Note how the two polynomials do coincide after setting $\xi = 1$ or $\eta = 1$.)

5.2. Extension to polymatroids

The interior and exterior polynomials are constructed from the hypertree polytope. The same procedure can be carried out after replacing $Q_{\mathcal{H}}$ with the base polytope B_μ of an integer extended polymatroid. Because of [Lemma 4.2](#) and the fact that integer polymatroids can always be defined by integer-valued non-decreasing submodular set functions, we immediately obtain invariants of integer polymatroids as well. In this subsection we outline the minimal modifications that are needed in our arguments to get the generalization.

Let S be a finite set and $\mu: \mathcal{P}(S) \rightarrow \mathbf{R}$ a submodular set function with base polytope $B_\mu \subset EP_\mu$. We say that the base $\mathbf{x} \in B_\mu \cap \mathbf{Z}^S$ is such that a *transfer* is possible from $s_1 \in S$ to $s_2 \in S$ if by decreasing the s_1 -component of \mathbf{x} by 1 and increasing its s_2 -component by 1, we get another base.

Now order S arbitrarily. Call an element $s \in S$ *internally active* with respect to the base $\mathbf{x} \in B_\mu \cap \mathbf{Z}^S$ if \mathbf{x} is such that no transfer is possible from s to a smaller element of S . We say that




Fig. 7. The geometry of the interior polynomial.

s is *externally active* with respect to \mathbf{x} if it is such that no transfer is possible to s from a smaller element of S .

For $\mathbf{x} \in B_\mu \cap \mathbf{Z}^S$, define $\bar{i}(\mathbf{x})$ to be the number of elements of S that are not internally active with respect to \mathbf{x} . Similarly, let $\bar{e}(\mathbf{x})$ denote the number of elements of S that are not externally active with respect to \mathbf{x} .

Definition 5.8. Let $\mu: S \rightarrow \mathbf{Z}$ be an integer-valued submodular set function. Define

$$I_\mu(\xi) = \sum_{\mathbf{x} \in B_\mu \cap \mathbf{Z}^S} \xi^{\bar{i}(\mathbf{x})} \quad \text{and} \quad X_\mu(\eta) = \sum_{\mathbf{x} \in B_\mu \cap \mathbf{Z}^S} \eta^{\bar{e}(\mathbf{x})}$$

and call these quantities the *interior polynomial*, respectively the *exterior polynomial* of μ .

The fact that these polynomials do not depend on the order that was used to define them can be shown just like in the proof of [Theorem 5.4](#). Indeed, the first half of that argument depended on the elementary [Lemma 5.2](#) (rhombus lemma) and [5.5](#) (staple lemma) where we only used the fact that $\mathcal{Q}_{\mathcal{H}}$ is cut out by placing upper bounds on each partial sum of the components of its elements. We also used the triviality that lines intersect bounded convex sets in segments. Then in the second half of the proof we relied on [Proposition 4.5](#) and the fact that our upper bounds are values of a tight set function.

To illustrate the geometry underlying our polynomials, we show in [Fig. 7](#) the hypothetical base polytope of a polymatroid P_μ with a ground set $S = \{a, b, c, d\}$ of four elements. It sits in a tetrahedron with vertices $\mathbf{a}(\mu(S), 0, 0, 0)$, $\mathbf{b}(0, \mu(S), 0, 0)$, $\mathbf{c}(0, 0, \mu(S), 0)$, and $\mathbf{d}(0, 0, 0, \mu(S))$. If we set the order $a < b < c < d$, then the marked vertex (which is the greedy hypertree of the order) is the only hypertree with $\bar{i} = 0$; the other lattice points along the six thickened edges are the ones with $\bar{i} = 1$; the rest of the lattice points that are visible (along a total of seven faces) in our view are those with $\bar{i} = 2$; finally the invisible lattice points have $\bar{i} = 3$.

A similar picture applies to the exterior polynomial. It would be very interesting to see whether invariants of arbitrary (i.e., not necessarily integer) extended polymatroids can be defined in a way analogous to [Definition 5.8](#), but replacing counts of lattice points along certain faces by taking the volumes of those faces.

It may also be worth it to investigate if there is a wider class of set functions (or polytopes) for which interior and exterior polynomials are well defined. Already in the set function case some assumption is certainly necessary, as the next example shows.

Example 5.9. Consider the tetrahedron

$$Q = \text{Conv}\{(1, 1, 0, 0), (1, 0, 1, 0), (1, 0, 0, 1), (0, 0, 1, 1)\} \subset \mathbf{R}_{xyzt}^4.$$

The four given points are exactly the integer solutions of the equation $x + y + z + t = 2$ and the linear inequalities

$$\begin{aligned} x \leq 1; & \quad x + y \leq 2; & \quad y + z + t \leq 2; \\ y \leq 1; & \quad x + z \leq 2; & \quad x + z + t \leq 2; \\ z \leq 1; & \quad x + t \leq 2; & \quad x + y + t \leq 2; \\ t \leq 1; & \quad y + z \leq 1; & \quad x + y + z \leq 2. \\ & \quad y + t \leq 1; \\ & \quad z + t \leq 2; \end{aligned}$$

Furthermore, each of the fourteen (fifteen) inequalities is sharp for at least one of the four points. The right hand sides *do not* give a tight set function. For example, the set of the y and z coordinates, as well as the set of the y and t coordinates is tight at the point $(1, 1, 0, 0)$. However the union of the two sets is not tight at the same point. We also see that **Lemma 4.6** is violated: the face along which the sum of the y and z (or the y and t) coordinates takes its maximum is a triangle instead of a rectangle.

From the order $x < y < z < t$, we obtain the interior and exterior polynomials $1 + 2\xi + \xi^2$ and $1 + 2\eta + \eta^2$, respectively. If we use $y < z < t < x$ instead, the interior polynomial becomes $2 + 2\xi^2$. For $x < t < z < y$, the exterior polynomial is $2 + 2\eta^2$.

6. Properties

6.1. Low-order terms

We first make an observation on the degrees of the interior and exterior polynomials and then turn our attention to some individual coefficients.

Proposition 6.1. *For a hypergraph $\mathcal{H} = (V, E)$, the degree of its interior polynomial is at most $\min\{|E|, |V|\} - 1$, while the degree of its exterior polynomial is at most $|E| - 1$.*

Proof. That $|E| - 1$ is an upper bound for both degrees is obvious from the observation that the smallest hyperedge in an order is both internally and externally active with respect to any hypertree.

As to $\deg I_{\mathcal{H}} \leq |V| - 1$, note that in order for a hyperedge e to be internally inactive with respect to a hypertree \mathbf{f} , we have to have $\mathbf{f}(e) \geq 1$. This combined with (3c) gives the result. \square

For the constant term in the exterior polynomial and the two lowest-order coefficients in the interior polynomial, the following two results offer a more direct way of seeing their order-independence.

Proposition 6.2. *Both the interior and exterior polynomials of the hypergraph $\mathcal{H} = (E, V)$ have 1 for constant term.*

Proof. Let us fix an arbitrary order on E and denote the greedy hypertree of **Lemma 4.12** with \mathbf{g} . We claim that \mathbf{g} is the unique hypertree with respect to which (and the given order) every

hyperedge is internally active. Comparing the assertions of Lemma 4.12 and Theorem 3.4 with the definition of internal activity, it is clear that \mathbf{g} is indeed one such hypertree.

The fact that there are no others follows easily either from Lemma 4.16 or from Lemma 4.17. Yet it may be interesting to see an argument that does not rely on abstract properties of submodular functions.

To this end, let $\mathbf{f}: E \rightarrow \mathbf{N}$ be a hypertree different from \mathbf{g} . We need to show that there is an internally inactive hyperedge with respect to it. We will use constructions from Section 4.2, in particular the subgraphs $G_1 \subset G_2 \subset \dots \subset G_{|E|} = \text{Bip } \mathcal{H}$ and the sequence of forests $F_1 \subset F_2 \subset \dots \subset F_{|E|}$ that was constructed in the proof of Lemma 4.12. Let e_k be the smallest hyperedge in the order so that $\mathbf{f}(e_k) \neq \mathbf{g}(e_k)$, which of course implies $\mathbf{f}(e_k) < \mathbf{g}(e_k)$. By Lemma 4.11 we may choose a realization Γ for \mathbf{f} so that $\Gamma \cap G_{k-1} = F_{k-1}$.

There is at least one connected component of $G_k - \{e_k\}$ which is not connected to e_k by any edge in Γ , even though there is an edge α of $\text{Bip } \mathcal{H}$ between e_k and this component (we can take for example the one which was selected into F_k). Adding α to Γ creates a cycle which has to go through a hyperedge e_l with $l > k$. By removing from Γ an edge of the cycle adjacent to e_l and replacing it with α , we create a new spanning tree for $\text{Bip } \mathcal{H}$. It induces a hypertree which only differs from \mathbf{f} at e_l (where it is one smaller than \mathbf{f}) and e_k (where it is one bigger). This shows that e_l is internally inactive with respect to \mathbf{f} .

In the case of the exterior polynomial, the unique hypertree without an (externally) inactive hyperedge is the greedy hypertree of the reverse order. The rest of the proof can be carried out just like above. \square

Theorem 6.3. *For any hypergraph $\mathcal{H} = (E, V)$, the coefficient of the linear term in the interior polynomial $I_{\mathcal{H}}$ is the nullity (first Betti number) $n(\text{Bip } \mathcal{H})$ of the bipartite graph $\text{Bip } \mathcal{H}$.*

Proof. We will extend the analysis carried out in the proofs of Lemma 4.12 and Proposition 6.2 a little further. After fixing an order on E , we are going to construct $n(\text{Bip } \mathcal{H})$ hypertrees. Namely, if e is a hyperedge, then we will associate to it $n_j(e)$ hypertrees as follows.

Recall the greedy hypertree $\mathbf{g}(e) = |e| - 1 - n_j(e)$ of Lemma 4.12 and fix one of its realizations. If e is a hyperedge that has $n_j(e) > 0$ with respect to the order, then add to the realization one more edge adjacent to e . This creates a cycle which goes through another hyperedge e' . Remove an edge of this cycle adjacent to e' to get a new spanning tree. Its induced hypertree is the result of a transfer of valence from e' to e . Then add another edge and make another transfer and so on until all edges adjacent to e are used up. The hyperedges that we transfer valence from are all smaller than e because of Lemma 4.12 and Theorem 3.4.

The argument in the previous paragraph shows that for all $1 \leq i \leq n_j(e)$, there is an i -element multiset of hyperedges smaller than e so that the result of reducing the \mathbf{g} -value associated to each of its elements by the multiplicity of the element, and simultaneously increasing $\mathbf{g}(e)$ to $\mathbf{g}(e) + i$, is a hypertree in \mathcal{H} . Now define $\mathbf{f}_{e,i}: E \rightarrow \mathbf{N}$ ($e \in E$, $1 \leq i \leq n_j(e)$) to be the hypertree among these so that its associated multiset $M_{e,i}$ is largest in reverse lexicographical order, i.e., if $e_1 < \dots < e_k$ are all the hyperedges smaller than e , then $M_{e,i}$ has the highest possible multiplicity at e_k ; then among those the highest possible multiplicity at e_{k-1} , and so on.

We claim that the $\mathbf{f}_{e,i}$ are exactly those hypertrees in \mathcal{H} that have a unique internally inactive hyperedge in the given order.

It is easy to see that the $\mathbf{f}_{e,i}$ do have this property. Namely, e is the unique hyperedge that is not internally active with respect to $\mathbf{f}_{e,i}$. Indeed, if $e' > e$ then e' is internally active with respect to

$\mathbf{f}_{e,i}$ because $\sum_{x < e'} \mathbf{f}_{e,i}(x) = \sum_{x < e'} \mathbf{g}(x)$ is already at the largest value allowed by (3b). If $e' < e$, then it is internally active because a downward transfer of valence from e' would contradict the way $M_{e,i}$ was chosen. Finally, e itself is not active because by Lemma 3.10 applied to $\mathbf{f}_1 = \mathbf{g}$ and $\mathbf{f}_2 = \mathbf{f}_{e,i}$, it can transfer valence to any element of $M_{e,i}$.

Let now \mathbf{f} be a hypertree in \mathcal{H} so that it has a unique internally inactive hyperedge e . Our goal is to show that \mathbf{f} is one of the $\mathbf{f}_{e,i}$.

First we claim that (3b) is sharp for \mathbf{f} and the set $E' = \{x \in E \mid x \leq e'\}$ for all $e' \geq e$: indeed if it was not, then by Lemma 4.16 there would have to be hyperedges larger than e' (and hence different from e) that are internally inactive with respect to \mathbf{f} . Because of the similar property of \mathbf{g} stated in Lemma 4.12, we see that $\mathbf{f} = \mathbf{g}$ (and hence $\mathbf{f} = \mathbf{f}_{e,i}$ for any i) for all hyperedges larger than e . It also follows that $\mathbf{f}(e) \geq \mathbf{g}(e)$ and in fact $\mathbf{f}(e) > \mathbf{g}(e)$, for otherwise e could not be inactive.

Next, we note that if e' is a hyperedge smaller than e , then $\mathbf{f}(e') \leq \mathbf{g}(e')$ because if this was not the case then Lemma 4.17 would imply that e' is internally inactive. So far we have shown that \mathbf{f} is obtained from \mathbf{g} by transferring valence to e from a (non-empty) multiset of hyperedges that are smaller than e .

Assume now that $\mathbf{f} \neq \mathbf{f}_{e,i}$, where we set $i = \mathbf{f}(e) - \mathbf{g}(e)$. Because of the way $\mathbf{f}_{e,i}$ was constructed, the largest hyperedge e' where the two hypertrees differ is so that $e' < e$ and $\mathbf{f}(e') > \mathbf{f}_{e,i}(e')$. We are going to argue that e' is internally inactive with respect to \mathbf{f} by showing that \mathbf{f} is such that e' can transfer valence to at least one element of the set $S = \{x \in E \mid \mathbf{f}(x) < \mathbf{f}_{e,i}(x)\}$. (Note that S is non-empty because $\sum_{y \in E} \mathbf{f}(y) = \sum_{y \in E} \mathbf{f}_{e,i}(y)$ and that its elements are smaller than e' .)

Suppose again that the opposite is true. Then by the usual argument based on Proposition 4.5, there exists a set U of hyperedges that is tight at \mathbf{f} , contains S , and does not contain e' . Over $E \setminus U$, the sum of \mathbf{f} -values is higher than the sum of $\mathbf{f}_{e,i}$ -values; therefore over U , the sum of $\mathbf{f}_{e,i}$ -values is higher. But this means that $\mathbf{f}_{e,i}$ and U contradict the inequality (3b) because the sum of \mathbf{f} -values over U is already $\mu(U)$, where μ is as in equation (4).

Therefore e' is internally inactive, but that contradicts our assumption on \mathbf{f} . The only possible conclusion, then, is $\mathbf{f} = \mathbf{f}_{e,i}$. \square

Remark 6.4. If we change the definition of $\mathbf{f}_{e,i}$ in the proof above to require that $M_{e,i}$ have the lowest possible multiplicity at e_1 , then among those choices the lowest possible multiplicity at e_2 and so on, then a similar argument reveals that this, too, is a hypertree in which e is the unique internally inactive hyperedge. Therefore the two descriptions define the same hypertree.

It is not hard to see (cf. [6, Section 44.6.c]) that Proposition 6.2 generalizes to arbitrary integer extended polymatroids, i.e., interior and exterior polynomials always have a constant term of 1. Theorem 6.3, in turn, offers a way of defining the nullity of an integer extended polymatroid via its interior polynomial. The resulting notion is a kind of ‘girth’ of the base polytope, cf. Fig. 7.

In the classical case, Proposition 6.2 translates to a well known property of the Tutte polynomial. Theorem 6.3, however, reveals (to the best of the author’s knowledge) previously undiscovered information.

Corollary 6.5. *Let $G = (V, E)$ be a connected graph. By summing the coefficients of its Tutte polynomial $T_G(x, y)$ in front of terms that contain $x^{|V|-2}$, we obtain the nullity of G . In other words, under any order, the number of spanning trees that contain exactly one internally inactive edge is the first Betti number of G as a one-dimensional cell complex.*

6.2. Product formulas

Let us now discuss situations when the interior and exterior polynomials behave multiplicatively. Our first claim is obvious.

Lemma 6.6. From $\mathcal{H} = (V, E)$, construct another hypergraph \mathcal{H}' by either

- (a) adding a singleton hyperedge $e' = \{v\}$ to E for some $v \in V$, or
- (b) adding a new vertex v' to V that is part of exactly one hyperedge $e \in E$.

Then $I_{\mathcal{H}'} = I_{\mathcal{H}}$ and $X_{\mathcal{H}'} = X_{\mathcal{H}}$.

Proof. In both cases, new and old hypertrees are in an obvious bijection. In the first, new hypertrees are exactly the extensions of old ones by assigning 0 to e' . The new hyperedge e' is both internally and externally active with respect to all hypertrees regardless of the order used.

In the second case, increase $\mathbf{f}(e)$ by 1 in each hypertree \mathbf{f} of \mathcal{H} to get the collection of hypertrees in \mathcal{H}' . With respect to any order, this correspondence preserves \bar{i} and \bar{e} values. \square

The main result of this subsection examines the picture when two disjoint bipartite graphs are joined by identifying one edge from each.

Theorem 6.7. Let $\mathcal{H}_1 = (V_1, E_1)$ and $\mathcal{H}_2 = (V_2, E_2)$ be hypergraphs so that $V_1 \cap V_2 = \{v\}$. Suppose that $e_1 \in E_1$ and $e_2 \in E_2$ both contain v and form the hypergraph

$$\mathcal{H} = (V, E) = (V_1 \cup V_2, (E_1 \setminus \{e_1\}) \cup (E_2 \setminus \{e_2\}) \cup \{e\})$$

by merging e_1 and e_2 into a single hyperedge $e = e_1 \cup e_2$. For hypertrees \mathbf{f}_1 in \mathcal{H}_1 and \mathbf{f}_2 in \mathcal{H}_2 , define the function $\mathbf{f}_1 \# \mathbf{f}_2: E \rightarrow \mathbf{N}$ by

$$(\mathbf{f}_1 \# \mathbf{f}_2)(e) = \mathbf{f}_1(e_1) + \mathbf{f}_2(e_2)$$

and by otherwise letting $(\mathbf{f}_1 \# \mathbf{f}_2)|_{E_i \setminus \{e_i\}} = \mathbf{f}_i, i = 1, 2$. Then $\#$ defines a bijection

$$(Q_{\mathcal{H}_1} \cap \mathbf{Z}^{E_1}) \times (Q_{\mathcal{H}_2} \cap \mathbf{Z}^{E_2}) \cong Q_{\mathcal{H}} \cap \mathbf{Z}^E.$$

Consequently,

$$I_{\mathcal{H}} = I_{\mathcal{H}_1} I_{\mathcal{H}_2} \quad \text{and} \quad X_{\mathcal{H}} = X_{\mathcal{H}_1} X_{\mathcal{H}_2}.$$

Proof. For the claim on the sets of hypertrees, choose arbitrary elements $\mathbf{f}_i \in Q_{\mathcal{H}_i} \cap \mathbf{Z}^{E_i}$ and realize them with spanning trees which contain the edges connecting e_i and v . This can be done by Lemma 3.3. After merging e_1 and e_2 , the union of the two trees becomes a spanning tree of Bip \mathcal{H} realizing $\mathbf{f}_1 \# \mathbf{f}_2$, which is therefore a hypertree.

Conversely, if \mathbf{f} is a hypertree in \mathcal{H} then we may realize it with a spanning tree that contains the edge between e and v . This can be separated into spanning trees of Bip \mathcal{H}_1 and of Bip \mathcal{H}_2 which induce hypertrees in \mathcal{H}_1 and \mathcal{H}_2 , respectively. It is easy to see that this defines an inverse to the correspondence $(\mathbf{f}_1, \mathbf{f}_2) \mapsto \mathbf{f}_1 \# \mathbf{f}_2$.

To prove the claim on the polynomials, we make the following observation. If the hypertree $\mathbf{f} = \mathbf{f}_1 \# \mathbf{f}_2$ is such that valence can be transferred from $a_1 \in E_1 \setminus \{e_1\}$ to $a_2 \in E_2 \setminus \{e_2\}$, then

- (a) \mathbf{f}_1 is such that valence can be transferred from a_1 to e_1 and
- (b) \mathbf{f}_2 is such that valence can be transferred from e_2 to a_2 .

Indeed, the result \mathbf{f}' of the assumed transfer of valence has a decomposition $\mathbf{f}' = \mathbf{f}'_1 \# \mathbf{f}'_2$ where (since the entries in \mathbf{f}_i and \mathbf{f}'_i share the same sum) \mathbf{f}'_1 is a hypertree that shows the truth of (a) and the existence of \mathbf{f}'_2 proves (b).

Let us now fix an order on E in which e is smallest and use its restrictions to order E_1 and E_2 (so that e_i becomes the smallest element in E_i). Then, we claim that for any pair of hypertrees \mathbf{f}_i in \mathcal{H}_i , we have

$$\bar{i}(\mathbf{f}_1 \# \mathbf{f}_2) = \bar{i}(\mathbf{f}_1) + \bar{i}(\mathbf{f}_2) \quad \text{and} \quad \bar{e}(\mathbf{f}_1 \# \mathbf{f}_2) = \bar{e}(\mathbf{f}_1) + \bar{e}(\mathbf{f}_2) \tag{12}$$

because, both in the internal and in the external sense, the set of inactive hyperedges for $\mathbf{f}_1 \# \mathbf{f}_2$ is the union of the sets of inactive hyperedges for \mathbf{f}_1 and for \mathbf{f}_2 . (As e, e_1 , and e_2 are smallest, they cannot be inactive with respect to any hypertree.) The backward inclusion is obvious because if, say, \mathbf{f}'_1 results from \mathbf{f}_1 by a single transfer of valence, then $\mathbf{f}'_1 \# \mathbf{f}_2$ and $\mathbf{f}_1 \# \mathbf{f}_2$ are related by essentially the same transfer. The forward inclusion follows equally easily with the help of (a) and (b) above.

Comparing (12) with the definitions of I and X completes the proof. \square

The next two statements describe the situation when two bipartite graphs are joined at a single vertex. They follow from [Theorem 6.7](#) and [Lemma 6.6](#) by first slightly enlarging one hypergraph as in the Lemma and then joining it to the other hypergraph as in the Theorem.

Corollary 6.8. *Let $\mathcal{H}_1 = (V_1, E_1)$ and $\mathcal{H}_2 = (V_2, E_2)$ be hypergraphs with disjoint vertex sets. Choose $e_1 \in E_1$ and $e_2 \in E_2$ and form the hypergraph \mathcal{H} with vertex set $V_1 \cup V_2$ and hyperedge set $(E_1 \setminus \{e_1\}) \cup (E_2 \setminus \{e_2\}) \cup \{e_1 \cup e_2\}$. Then $I_{\mathcal{H}} = I_{\mathcal{H}_1} I_{\mathcal{H}_2}$ and $X_{\mathcal{H}} = X_{\mathcal{H}_1} X_{\mathcal{H}_2}$.*

Corollary 6.9. *Let $\mathcal{H}_1 = (V_1, E_1)$ and $\mathcal{H}_2 = (V_2, E_2)$ be hypergraphs with disjoint vertex sets. Choose $v_1 \in V_1$ and $v_2 \in V_2$ and form the hypergraph \mathcal{H} by identifying them to a single vertex v . I.e., the vertex set of \mathcal{H} is $(V_1 \setminus \{v_1\}) \cup (V_2 \setminus \{v_2\}) \cup \{v\}$ and its hyperedge set is identified with $E_1 \cup E_2$ so that v is an element of a ‘new’ hyperedge if and only if either v_1 or v_2 was an element of the ‘old’ hyperedge. Then $I_{\mathcal{H}} = I_{\mathcal{H}_1} I_{\mathcal{H}_2}$ and $X_{\mathcal{H}} = X_{\mathcal{H}_1} X_{\mathcal{H}_2}$.*

6.3. Deletion and contraction

These operations are of special importance in the classical theory of the Tutte polynomial. However in the hypergraph case, so far they have played surprisingly small roles.

Definition 6.10. Let $\mathcal{H} = (V, E)$ be a hypergraph and $e \in E$ a hyperedge. *Deleting e from \mathcal{H} means passing to the hypergraph $\mathcal{H} \setminus \{e\} = (V, E \setminus \{e\})$. The result of contracting e is the hypergraph \mathcal{H}/e which is obtained from $\mathcal{H} \setminus \{e\}$ by identifying the elements of e . I.e., the hyperedge set of \mathcal{H}/e is essentially $E \setminus \{e\}$ and its vertex set is $(V \setminus e) \cup \{\bar{e}\}$, where the new vertex \bar{e} belongs to the hyperedge $e' \in E \setminus \{e\}$ if and only if $e' \cap e \neq \emptyset$.*

Regarding the hypertree polytope $Q_{\mathcal{H}}$ and the hyperedge e , let us define the *top face* of $Q_{\mathcal{H}}$ with respect to e as $\{\mathbf{f} \in Q_{\mathcal{H}} \mid \mathbf{f}(e) = |e| - 1\}$ and let the *bottom face* be $\{\mathbf{f} \in Q_{\mathcal{H}} \mid \mathbf{f}(e) = 0\}$. Because of [Corollary 4.10](#), the top and bottom faces with respect to any hyperedge coincide with the convex hulls of the hypertrees that they contain.

Proposition 6.11. *Let $\mathcal{H} = (V, E)$ be a hypergraph with hypertree polytope $Q_{\mathcal{H}}$. For any hyperedge $e \in E$, the hypertree polytopes $Q_{\mathcal{H} \setminus \{e\}}$ and $Q_{\mathcal{H}/e}$ are naturally isomorphic to the bottom and top face, respectively, of $Q_{\mathcal{H}}$ with respect to e .*

Proof. It suffices to equate the sets of hypertrees within the polytopes.

Deletion: if we take any hypertree in $\mathcal{H} \setminus \{e\}$ and an arbitrary realization, then by adding to it a single edge of $\text{Bip } \mathcal{H}$ adjacent to e , we realize a hypertree that lies along the bottom face. Conversely, any realization of a hypertree from the bottom face of $Q_{\mathcal{H}}$ has a single edge adjacent to e and by removing it we obtain a tree which realizes a hypertree in $\mathcal{H} \setminus \{e\}$. Therefore a bijection is defined by extending hypertrees in $\mathcal{H} \setminus \{e\}$ to e with the value zero. (It is possible for the bottom face of $Q_{\mathcal{H}}$ to be empty, namely when $\text{Bip } \mathcal{H} - \{e\} = \text{Bip}(\mathcal{H} \setminus \{e\})$ is disconnected. Of course in such a case $Q_{\mathcal{H} \setminus \{e\}}$ is empty as well.)

Contraction: let \bar{e} denote the new vertex in \mathcal{H}/e that results from the contraction of e . If \mathbf{f} is a hypertree in \mathcal{H}/e then, using Lemma 3.3, construct a realization Γ for it which contains all edges adjacent to \bar{e} . Based on this tree, we carry out the following construction in $\text{Bip } \mathcal{H}$. Keep all edges of Γ that are not adjacent to \bar{e} . If $e' \in E \setminus \{e\}$ is such that $e' \cap e \neq \emptyset$, then connect e' to an arbitrarily chosen element of the intersection. Finally, add all edges that are adjacent to e . The result is a spanning tree in $\text{Bip } \mathcal{H}$ so that its induced hypertree is on the one hand part of the top face of $Q_{\mathcal{H}}$, and on the other hand it agrees with \mathbf{f} on $E \setminus \{e\}$.

The inverse of this correspondence is constructed as follows. Let \mathbf{g} be a hypertree from the top face of $Q_{\mathcal{H}}$. Any of its realizations contains all edges of $\text{Bip } \mathcal{H}$ adjacent to e . Take one such tree and (viewing it as a topological space) contract the union of its edges adjacent to e to a single point. The result is another tree which naturally embeds in $\text{Bip}(\mathcal{H}/e)$ as a spanning tree so that it realizes a hypertree in \mathcal{H}/e that agrees with \mathbf{g} over $E \setminus \{e\}$. \square

We now turn to generalizations of the classical deletion–contraction formulas (1). Lemma 6.6 can be viewed as a trivial extension of the second formula to hypergraphs. (The reader may wish to consult (9) to see that there is no contradiction.) The first formula also has an easy generalization to hypergraphs as follows.

Proposition 6.12. *Let $\mathcal{H} = (V, E)$ be a hypergraph with a hyperedge e so that $\text{Bip } \mathcal{H}$ is connected but $\text{Bip } \mathcal{H} - \{e\}$ has $|e|$ connected components. Then $I_{\mathcal{H}} = I_{\mathcal{H}/e}$ and $X_{\mathcal{H}} = X_{\mathcal{H}/e}$ (and both are equal to obvious $|e|$ -fold products).*

Proof. Immediate from Lemma 6.6 and Corollaries 6.8 and 6.9. \square

Remark 6.13. For the abstract duals of the hypergraphs in the previous proposition, the same argument yields the same conclusion: $I_{\overline{\mathcal{H}}} = I_{\overline{\mathcal{H}/e}}$ and $X_{\overline{\mathcal{H}}} = X_{\overline{\mathcal{H}/e}}$, since both sides of each equation agree with the obvious $|e|$ -fold product.

Finally, let us generalize the third deletion–contraction formula from (1).

Proposition 6.14. *Let $\mathcal{H} = (V, E)$ be a hypergraph that contains a two-element hyperedge e so that $\text{Bip } \mathcal{H} - \{e\}$ is connected. Then we have*

$$I_{\mathcal{H}}(\xi) = I_{\mathcal{H} \setminus \{e\}}(\xi) + \xi I_{\mathcal{H}/e}(\xi) \quad \text{and} \quad X_{\mathcal{H}}(\eta) = \eta X_{\mathcal{H} \setminus \{e\}}(\eta) + X_{\mathcal{H}/e}(\eta). \tag{13}$$

Proof. As $|e| = 2$, each hypertree in \mathcal{H} lies either on the top (if its value at e is 1) or on the bottom (if it is 0) face of $Q_{\mathcal{H}}$ with respect to e . Let us choose an order on E in which e is the largest hyperedge.

With respect to hypertrees along the bottom face, e is internally active. With respect to those on the top face, by Lemma 4.17, e is internally inactive: indeed the lemma applies because $n_j(e) = 1$ by our assumption which means that the value of the greedy hypertree at e is $\mathbf{g}(e) = |e| - 1 - n_j(e) = 0$. With this, our first equation follows easily from Proposition 6.11.

The second is equally easy after noticing that with respect to hypertrees of the top face, e is externally active, whereas with respect to those along the bottom face, it is externally inactive. The former is obvious and the latter can be argued for example like this: if \mathbf{f} is a hypertree in \mathcal{H} so that $\mathbf{f}(e) = 0$, then take any of its realizations and add to it the ‘other’ edge adjacent to e , too. The cycle thus created goes through at least one more hyperedge e' and if we break the cycle by removing one of its edges adjacent to e' , we get a realization of a hypertree which is the result of a transfer of valence to e from the smaller hyperedge e' . \square

It would be desirable to formulate a version of the last proposition for hyperedges of arbitrary size. Such formulas, however, are currently lacking.

7. Abstract duality

The observations on the interior polynomial contained in Proposition 6.2 and Theorem 6.3 illustrate more than just the (already proved) order-independence of the coefficients. The values observed also remain the same if we exchange the roles of hyperedges and vertices, that is, they are invariant under abstract duality. We conjecture that this is true for the rest of the polynomial as well. I.e., the interior polynomial is in fact an invariant of bipartite graphs.

Conjecture 7.1. *Let $G = (V_0, V_1, E)$ be a connected bipartite graph which induces the hypergraphs \mathcal{G}_0 and \mathcal{G}_1 as in (2). Then $I_{\mathcal{G}_0} = I_{\mathcal{G}_1}$.*

So far, supporting evidence includes Postnikov’s Theorem 3.11, which says that $I_{\mathcal{G}_0}(1) = I_{\mathcal{G}_1}(1)$, and Theorem 6.3. We present some more below.

Conjecture 7.1 was certainly true in Example 5.6, where it was also illustrated that abstract duality will usually not preserve the exterior polynomial. (Indeed, the linear terms can already be different.) Let us work out another family of examples.

Example 7.2. We consider the complete bipartite graph (V_0, V_1, E) , where $|V_0| = n$ and $|V_1| = m$. For both induced hypergraphs \mathcal{G}_0 and \mathcal{G}_1 the condition (3b) in Theorem 3.4 becomes vacuous and we see that their hypertree polytopes are the simplices

$$Q_{\mathcal{G}_0} = (m - 1)\Delta_{V_0} \quad \text{and} \quad Q_{\mathcal{G}_1} = (n - 1)\Delta_{V_1},$$

respectively, both of which contain $\binom{n+m-2}{n-1}$ hypertrees. (The number of lattice points in a standard d -dimensional simplex of sidelength k is $\binom{k+d}{d}$.)

Fix now some order on V_0 . If $\mathbf{f}: V_0 \rightarrow \mathbf{N}$ is a hypertree in \mathcal{G}_0 and $e \in V_0$, then e is internally inactive with respect to \mathbf{f} if and only if $\mathbf{f}(e) > 0$ and e is not the smallest in the order. Hence to enumerate hypertrees in \mathcal{G}_0 with internal inactivity k , we need to

- (a) choose the set of the k internally inactive hyperedges and
- (b) partition the sum of hypertree entries, $m - 1$, between them and the smallest hyperedge so that the latter may receive 0 but the others get a positive amount.

This is an easy exercise whose solution is $\binom{n-1}{k} \binom{m-1}{k}$. As this is a symmetric expression in n and m , Conjecture 7.1 holds for complete bipartite graphs.

We leave it to the reader to verify that the number of hypertrees in \mathcal{G}_0 with external inactivity k is $\binom{m+k-2}{k}$ for all $0 \leq k \leq n - 1$. For this of course it does matter if we interchange n and m .

Let us temporarily call a bipartite graph *good* if it satisfies [Conjecture 7.1](#). Bipartite graphs with an isomorphism that interchanges their color classes are good. We have seen a number of operations that, when applied to good graphs, result in other good graphs. These include adding on a new valence one point ([Lemma 6.6](#)), joining two bipartite graphs at a vertex ([Corollaries 6.8 and 6.9](#)) or at an edge ([Theorem 6.7](#)).

We will introduce one more powerful way of generating new good bipartite graphs from old. It requires the following ‘upside down’ version of [Proposition 6.14](#).

Proposition 7.3. *Let $\mathcal{H} = (V, E)$ be a hypergraph with a vertex $v \in V$ belonging to exactly two hyperedges e_1 and e_2 . Define the deletion of v as $\mathcal{H}' = \overline{\mathcal{H}} \setminus \{v\}$ and the contraction of v as $\mathcal{H}'' = \overline{\mathcal{H}} / v$. If $\text{Bip } \mathcal{H} - \{v\}$ is connected, we have*

$$I_{\mathcal{H}}(\xi) = I_{\mathcal{H}'}(\xi) + \xi I_{\mathcal{H}''}(\xi). \tag{14}$$

If we also suppose that v is not the only common element of e_1 and e_2 , we obtain the relation $X_{\mathcal{H}}(\eta) = X_{\mathcal{H}'}(\eta) + \eta X_{\mathcal{H}''}(\eta)$ as well.

Notice that (14) matches the first equation in (13) exactly. There is however a discrepancy between the formulas concerning exterior polynomials, not to mention the extra assumption. By the way, the assumption that v not be the only common element of e_1 and e_2 cannot be dropped. If $\text{Bip } \mathcal{H}$ is a $2k$ -cycle with $k \geq 2$, then (for any choice of v) we have $X_{\mathcal{H}}(\eta) = 1 + (k - 1)\eta$, $X_{\mathcal{H}'}(\eta) = 1$, and $X_{\mathcal{H}''}(\eta) = 1 + (k - 2)\eta$. The stated equation on exterior polynomials only holds if $k = 2$, which is exactly when the extra assumption is satisfied.

In the proof of [Proposition 7.3](#) we will use the following claim. It is true for integer bases of arbitrary integer extended polymatroids, but for convenience we will only state it for hypertrees.

Lemma 7.4. *Let \mathbf{x} be a hypertree in the hypergraph \mathcal{H} . Assume that the distinct hyperedges a, b (“givers”) and c, d (“receivers”) are such that a simultaneous transfer of valence is possible at \mathbf{x} , i.e., that $\mathbf{y} = \mathbf{x} - \mathbf{i}_{\{a\}} - \mathbf{i}_{\{b\}} + \mathbf{i}_{\{c\}} + \mathbf{i}_{\{d\}}$ is also a hypertree. Then \mathbf{x} is such that either a can transfer valence to c and b to d , or a can transfer valence to d and b to c .*

Proof. If \mathbf{x} is such that any giver can transfer valence to any receiver, then we are done. Assume that one of the transfers, say from a to d , is not possible. It suffices to show that in this case a can transfer valence to c and b can transfer valence to d .

This follows from [Lemma 4.6](#) and some elementary geometry applied to the $abcd$ cross-section (through \mathbf{x}) of $Q_{\mathcal{H}}$. Namely, if the plane π through $\mathbf{x}, \mathbf{y}, \mathbf{x} - \mathbf{i}_{\{a\}} + \mathbf{i}_{\{c\}}$, and $\mathbf{x} - \mathbf{i}_{\{b\}} + \mathbf{i}_{\{d\}}$ was not a supporting plane of the cross-section, then because of the directions of the other possible supporting planes and the presence of \mathbf{x} and \mathbf{y} , the point $\mathbf{x} - \mathbf{i}_{\{a\}} + \mathbf{i}_{\{d\}}$ would also be a hypertree. Thus π is a supporting plane and so it forms a rectangular intersection with $Q_{\mathcal{H}}$, which proves that $\mathbf{x} - \mathbf{i}_{\{a\}} + \mathbf{i}_{\{c\}}$ and $\mathbf{x} - \mathbf{i}_{\{b\}} + \mathbf{i}_{\{d\}}$ are indeed hypertrees. \square

The typical situation in which this lemma will be applied is when we know that $\mathbf{x}, \mathbf{x} - \mathbf{i}_{\{a\}} + \mathbf{i}_{\{c\}}$, and $\mathbf{x} - \mathbf{i}_{\{a\}} - \mathbf{i}_{\{b\}} + \mathbf{i}_{\{c\}} + \mathbf{i}_{\{d\}}$ are hypertrees and we conclude that \mathbf{x} is such that b can transfer valence either to d or to c ; also that d can receive a transfer from a or from b .

Yet another way of phrasing the statement would be that for certain regular octahedra in \mathbf{Z}^E , if $Q_{\mathcal{H}}$ contains one pair of opposite vertices then it has to contain another pair. Note that the polyhedron used in [Example 5.9](#) violates this principle.

Proof of Proposition 7.3. Just like in the proof of [Proposition 6.14](#), first we are going to split the set of hypertrees in \mathcal{H} in two so that the two sides are in one-to-one correspondences with

hypertrees in \mathcal{H}' and \mathcal{H}'' , respectively, and then study activities. The geometry however will be different this time and therefore the argument more complex.

If \mathbf{f} is a hypertree in \mathcal{H}' , define $\mathbf{f}'(e_1) = \mathbf{f}(e_1) + 1$ and $\mathbf{f}'(x) = \mathbf{f}(x)$ for other hyperedges. This is a hypertree in \mathcal{H} because a realization of \mathbf{f} in $\text{Bip } \mathcal{H}' = \text{Bip } \mathcal{H} - \{v\}$ together with the edge of $\text{Bip } \mathcal{H}$ between e_1 and v gives a realization of \mathbf{f}' . The association $\mathbf{f} \mapsto \mathbf{f}'$ is an isometric embedding, in particular injective.

Note that \mathcal{H}'' contains a distinguished hyperedge e where e_1 and e_2 merge. If \mathbf{g} is a hypertree in \mathcal{H}'' then \mathcal{H} has hypertrees which agree with \mathbf{g} away from e_1 and e_2 because the same edges that constitute a realization of \mathbf{g} in $\text{Bip } \mathcal{H}''$ form a two-component forest in $\text{Bip } \mathcal{H}$ which can be turned into a spanning tree by adding to it both edges adjacent to v . The set of these,

$$L_{\mathbf{g}} = \left\{ \mathbf{h} \in Q_{\mathcal{H}} \cap \mathbf{Z}^E \mid \mathbf{h}|_{E \setminus \{e_1, e_2\}} = \mathbf{g}|_{E \setminus \{e_1, e_2\}} \right\},$$

is analogous to the line segments that we used in the proof of [Theorem 5.4](#). Let us associate to \mathbf{g} the hypertree $\mathbf{g}'' \in L_{\mathbf{g}}$ with the highest value at e_2 .

The sets $L_{\mathbf{g}}$, which we will also refer to as the *fibers*, lie on parallel lines. Any hypertree \mathbf{h} in \mathcal{H} is part of a set $L_{\mathbf{g}}$. To see this, use [Lemma 3.3](#) to construct a realization of \mathbf{h} containing both edges adjacent to v and then contract those edges to a point to obtain a spanning tree in $\text{Bip } \mathcal{H}''$ that induces the right \mathbf{g} .

Furthermore, if $\mathbf{g}_1 \neq \mathbf{g}_2$, then the lines containing $L_{\mathbf{g}_1}$ and $L_{\mathbf{g}_2}$ are disjoint (so that the hypertree \mathbf{g} of the previous paragraph is uniquely determined by \mathbf{h}). This is because $E \setminus \{e_1, e_2\}$ represents all but one of the hyperedges in \mathcal{H}'' and by (3c), the last value (at e) is uniquely determined by the rest. In particular, the correspondence $\mathbf{g} \mapsto \mathbf{g}''$ is also injective.

There do not exist \mathbf{f} and \mathbf{g} so that $\mathbf{f}' = \mathbf{g}''$ because

- (a) For any \mathbf{f} , the hypertree \mathbf{f}' is such that a transfer of valence is possible from e_1 to e_2 : just repeat the argument that \mathbf{f}' is a hypertree but with the roles of e_1 and e_2 reversed.
- (b) At a hypertree of the form \mathbf{g}'' , the same transfer is never possible.

On the other hand, elements of $L_{\mathbf{g}} \setminus \{\mathbf{g}''\}$, i.e., all hypertrees \mathbf{h} in \mathcal{H} that are such that a transfer of valence is possible from e_1 to e_2 , are of the form \mathbf{f}' for some \mathbf{f} . This is because each such \mathbf{h} has a realization that does not contain the edge γ of $\text{Bip } \mathcal{H}$ connecting e_2 and v . To show this, take any realization Ξ of \mathbf{h} . If it does not contain γ , we are done. Otherwise, by [Lemma 3.3](#), we may assume that Ξ contains both edges of $\text{Bip } \mathcal{H}$ adjacent to v . By removing those two edges, we separate Ξ to two smaller trees Ξ_1 and Ξ_2 . Let $E_1 \ni e_1$ be the set of hyperedges contained by Ξ_1 and let $E_2 \ni e_2$ be contained by Ξ_2 so that E is the disjoint union of E_1 and E_2 .

There have to be elements of E_2 that are adjacent in $\text{Bip } \mathcal{H}$ to vertices of Ξ_1 . This is because otherwise $\Xi \cap (\text{Bip } \mathcal{H}|_{E_2}) = \Xi_2 \cup \{\gamma\}$ would be a spanning tree in $\text{Bip } \mathcal{H}|_{E_2}$, which by [Lemma 4.11](#) would mean that E_2 is tight at \mathbf{h} and that would rule out the possibility of the assumed transfer of valence from e_1 to e_2 . Whenever we locate this kind of element a in E_2 , we ‘switch it over’ to E_1 by our usual method: add to Ξ an edge α connecting a to a vertex of Ξ_1 ; this creates a unique cycle C in $\Xi \cup \{\alpha\}$ which necessarily contains γ ; remove the edge of C adjacent to a which is not α .

By iterating this procedure, we move through realizations of \mathbf{h} with smaller and smaller associated sets E_2 . Therefore after finitely many steps, the hyperedge a that we operate on will be e_2 itself. The edge that we remove in that step is necessarily γ . This finishes the argument that a hypertree in \mathcal{H} is always either of the form \mathbf{f}' or of the form \mathbf{g}'' but never both.

In order to study activities, let us now order E so that e_1 and e_2 are the smallest hyperedges. (Later we will use $e_1 < e_2$ in the internal case and $e_2 < e_1$ in the external case, but for now we leave this unspecified.) We will use this order in \mathcal{H} to define \bar{t} and $\bar{\varepsilon}$, as well as in \mathcal{H}' to define \bar{t}' and $\bar{\varepsilon}'$ values. In \mathcal{H}'' , let e be the smallest hyperedge and otherwise use the same order in the definitions of \bar{t}'' and $\bar{\varepsilon}''$. We start with two easy claims.

- (i) If a hyperedge a is inactive (in either sense) with respect to \mathbf{f} in \mathcal{H}' , then a is inactive with respect to \mathbf{f}' in \mathcal{H} . This is because if the hypertrees \mathbf{f}_1 and \mathbf{f}_2 in \mathcal{H}' are related by a single transfer of valence then \mathbf{f}'_1 and \mathbf{f}'_2 are related by the same transfer.
- (ii) Regarding \mathcal{H}'' , an implication of the opposite kind is easy: if the hyperedge $a \in E \setminus \{e_1, e_2\}$ is inactive with respect to some element of $L_{\mathbf{g}}$ (in \mathcal{H}), then a is inactive with respect to \mathbf{g} (in \mathcal{H}'') as well. Just take the hypertree that results from the transfer of valence which ‘makes’ a inactive and notice that it is part of a fiber $L_{\mathbf{g}_0}$; here \mathbf{g}_0 differs from \mathbf{g} by the single transfer of valence which establishes the claim.

The rest of the proof is concerned with the degree to which the converses of (i) and (ii) above are true. We start with (ii).

Lemma 7.5. *If the hyperedge $a \in E \setminus \{e_1, e_2\}$ is inactive, in either sense, with respect to the hypertree \mathbf{g} in \mathcal{H}'' , then a is inactive with respect to any element $\mathbf{h} \in L_{\mathbf{g}}$ (e.g. $\mathbf{h} = \mathbf{g}'$) in \mathcal{H} .*

To show this, we separate two cases. Let $\mathbf{h} \in L_{\mathbf{g}}$.

- I. If \mathbf{g} is such that a transfer of valence is possible from a to e (resp. from e to a) resulting in the hypertree \mathbf{g}_0 in \mathcal{H}'' , then we claim that \mathbf{h} is such that a transfer of valence is possible from a to e_1 or e_2 (resp. from e_1 or e_2 to a). This follows by fixing a 2-dimensional plane containing $L_{\mathbf{g}}$ and $L_{\mathbf{g}_0}$ (so that only the e_1, e_2 , and a -coordinates can vary) and constructing the type of elementary argument that we used to prove [Lemmas 5.2](#) and [5.5](#).
- II. If \mathbf{g} is such that no transfer of valence is possible from a to e (resp. from e to a), then by the argument in (ii) it follows that \mathbf{h} is such that there is no transfer of valence from a neither to e_1 nor to e_2 (resp. neither from e_1 nor from e_2 to a). As a is inactive, there is another hyperedge $b \in E \setminus \{e_1, e_2\}$, smaller than a , so that \mathbf{g} is such that valence can be transferred from a to b (resp. from b to a), resulting in the hypertree \mathbf{g}_0 . From [Lemma 5.2](#) and (ii) we obtain that \mathbf{h} is such that valence cannot be transferred from b to e_1 or e_2 (resp. from e_1 or e_2 to b) either. Thus [Lemma 4.6](#), applied to $p = e_1, q = e_2, r = a, s = b$ (resp. $p = a, q = b, r = e_1, s = e_2$) implies that $L_{\mathbf{g}}$ and $L_{\mathbf{g}_0}$ are located in a rectangular cross-section of $Q_{\mathcal{H}}$ so that every hypertree along $L_{\mathbf{g}}$ is such that a transfer of valence is possible from a to b (resp. from b to a). This completes the proof of the lemma.

Let us now assume that e_1 and e_2 have a second common vertex $u \neq v$. This is equivalent to saying that for any hypertree \mathbf{g} in \mathcal{H}'' , the fiber $L_{\mathbf{g}}$ has at least two elements,² which in turn makes it meaningful for us to let \mathbf{g}''' denote the element of $L_{\mathbf{g}}$ adjacent to \mathbf{g}' . (Note that \mathbf{g}''' is of the form \mathbf{f}' for some \mathbf{f} .) Indeed, if u exists, then any hypertree in \mathcal{H} is such that a transfer of valence is possible between e_1 and e_2 : using [Lemma 3.3](#), force a realization to contain the edges e_1u and e_2u , then replace one of them with the unique element of $\{e_1v, e_2v\}$ that is originally not in the realization to achieve the transfer. Conversely, if $e_1 \cap e_2 = \{v\}$, then there are (greedy) hypertrees in \mathcal{H} which take the maximal values $|e_1| - 1$ and $|e_2| - 1$ at e_1 and e_2 , respectively, and hence form singleton fibers.

² As an alternative to the short argument below, the equivalence can also be deduced from Postnikov’s [Proposition 3.12](#).

In addition to our previous assumptions (including the existence of u), we now set $e_2 < e_1$ and prove the following items.

- $\bar{\varepsilon}(\mathbf{g}'') = \bar{\varepsilon}''(\mathbf{g}) + 1$ for all hypertrees \mathbf{g} in \mathcal{H}'' . It is clear that e_1 is externally inactive with respect to \mathbf{g}'' whereas of course e is active with respect to \mathbf{g} . The rest follows from Lemma 7.5 and (ii).
- $\bar{\varepsilon}(\mathbf{f}') = \bar{\varepsilon}'(\mathbf{f})$ for all hypertrees \mathbf{f} in \mathcal{H}' . A converse of (i) is indeed true in this case. If \mathbf{f}' is such that e_2 can transfer valence to e_1 , then the same is obviously true for \mathbf{f} . If \mathbf{f}' is such that e_1 or e_2 can transfer valence to some other hyperedge a then, even if the result of the transfer happens to be of the form \mathbf{g}'' , it is because the transfer came from e_1 and it is easy to see that a transfer of valence from e_2 to a turns \mathbf{f}' into \mathbf{g}''' . Finally, if \mathbf{f}' and a are such that no transfer of valence is possible from e_1 or e_2 to a but the hyperedge b can transfer valence to a , resulting in the hypertree \mathbf{h} , then the usual combination of Lemmas 5.2 and 4.6 shows that the fibers of \mathbf{f}' and \mathbf{h} lie in a rectangular cross-section and therefore \mathbf{h} is not of the form \mathbf{g}'' either.

The last two equations together settle the claim in the Proposition concerning exterior polynomials.

Finally, we turn to internal activities and establish (14). For the remainder of the proof, we will use the order $e_1 < e_2$ among the two smallest hyperedges. Under the simplifying assumption that all fibers contain at least two hypertrees, the result follows relatively easily just like above in the external case (with the roles of e_1 and e_2 suitably exchanged). Without that assumption, both components of the argument break down: some hypertrees \mathbf{g}'' may be such that e_2 cannot transfer valence to e_1 , i.e., \mathbf{g}'' does not pick up an extra internally inactive hyperedge beyond those of \mathbf{g} ; on the other hand there may also be hypertrees \mathbf{f}' which do pick up extra internally inactive hyperedges beyond those of \mathbf{f} because some transfer of valence can turn them into one of the \mathbf{g}'' without a convenient \mathbf{g}''' nearby. Our task is to show that the two problems cancel each other out.

Let us call a hypertree in \mathcal{H} *lonely* if it is the unique element of its fiber, i.e., if it is such that no transfer of valence is possible between e_1 and e_2 . We claim that if \mathbf{h} is a lonely hypertree, then there is a hyperedge $a \in E \setminus \{e_1, e_2\}$ so that \mathbf{h} is such that both e_1 and e_2 can transfer valence to a . After noting that any realization Ξ of \mathbf{h} has to contain both edges of $\text{Bip } \mathcal{H}$ adjacent to v , we proceed as earlier (on page 853) to define the subtrees and hyperedge sets $\Xi_i \supset E_i \ni e_i, i = 1, 2$. Now there has to be a hyperedge a in E_1 or E_2 that is connected by an edge α of $\text{Bip } \mathcal{H}$ to a vertex in Ξ_2 or Ξ_1 , respectively, for otherwise removing v would disconnect the graph. Adding α to Ξ creates a cycle through v so that by removing either edge of $\text{Bip } \mathcal{H}$ adjacent to v , we realize the two desired transfers of valence to a .

If \mathbf{h} is a lonely hypertree then select the smallest hyperedge $a_{\mathbf{h}}$ with the property above and let the hypertree $t(\mathbf{h})$ be obtained from \mathbf{h} by a single transfer of valence from e_2 to $a_{\mathbf{h}}$. From the construction it is obvious that there is a hypertree \mathbf{f} in \mathcal{H}' so that $\mathbf{f}' = t(\mathbf{h})$. It is also clear from Lemma 5.2 that $t(\mathbf{h})$ is part of a two-element fiber; in particular it is such that e_2 cannot transfer valence to e_1 .

For a hypertree \mathbf{f} in \mathcal{H}' , let us introduce

$$T_{\mathbf{f}} = \{ \mathbf{h} \in Q_{\mathcal{H}} \cap \mathbf{Z}^E \mid \mathbf{h} \text{ is a lonely hypertree with } t(\mathbf{h}) = \mathbf{f}' \}.$$

Clearly, if $\mathbf{h} \neq \mathbf{j}$ belong to the same set $T_{\mathbf{f}}$, then $a_{\mathbf{h}} \neq a_{\mathbf{j}}$ because \mathbf{f} and $a_{\mathbf{h}}$ determine \mathbf{h} via $t(\mathbf{h}) = \mathbf{f}'$. The proof of (14) will be immediate from the following items.

- 1) $\bar{t}(\mathbf{g}'') = \bar{t}''(\mathbf{g}) + 1$ for all hypertrees \mathbf{g} in \mathcal{H}'' so that \mathbf{g}'' is not lonely. This follows easily from (ii), Lemma 7.5, and the observation that e_2 is internally inactive with respect to \mathbf{g}'' in this case.
- 2) $\bar{t}(\mathbf{g}'') = \bar{t}''(\mathbf{g})$ whenever \mathbf{g}'' is lonely. In these cases e_2 is internally active; otherwise see above.
- 3) $\bar{t}(\mathbf{f}') = \bar{t}'(\mathbf{f}) + |T_{\mathbf{f}}|$ for all hypertrees \mathbf{f} in \mathcal{H}' .
- 4) For any hypertree \mathbf{f} in \mathcal{H}' , the values of \bar{t} among the elements of $T_{\mathbf{f}}$ are $\bar{t}'(\mathbf{f}), \bar{t}'(\mathbf{f}) + 1, \dots, \bar{t}'(\mathbf{f}) + |T_{\mathbf{f}}| - 1$, each occurring exactly once.

We will prove (3) and (4) momentarily but let us first take care of (14). For each hypertree \mathbf{g} in \mathcal{H}'' , we need a hypertree in \mathcal{H} with \bar{t} one higher than $\bar{t}''(\mathbf{g})$. If it is not lonely then \mathbf{g}'' can play this role by (1); otherwise (2) says that \mathbf{g}'' is not good but it belongs to a unique set $T_{\mathbf{f}}$ which, by (4), either contains an element with the desired \bar{t} or if it does not then, according to (3), \mathbf{f}' will do the job. Also, for each hypertree \mathbf{f} in \mathcal{H}' we need a hypertree in \mathcal{H} with the same \bar{t} . If $T_{\mathbf{f}} = \emptyset$ then, by (3), \mathbf{f}' will do; otherwise (4) says that the element of $T_{\mathbf{f}}$ with the lowest \bar{t} will suffice. As we used each hypertree in \mathcal{H} exactly once, (14) follows.

To show (3), let us now fix a hypertree \mathbf{f} in \mathcal{H}' . For each $\mathbf{h} \in T_{\mathbf{f}}$, the hyperedge $a_{\mathbf{h}}$ is internally inactive with respect to \mathbf{f}' because a transfer of valence from $a_{\mathbf{h}}$ to e_2 turns \mathbf{f}' into \mathbf{h} . We claim that the same hyperedges $a_{\mathbf{h}}$ are internally active with respect to \mathbf{f} , implying via (i) that $\bar{t}(\mathbf{f}') \geq \bar{t}'(\mathbf{f}) + |T_{\mathbf{f}}|$. For this it suffices to prove that any downward transfer of valence from $a_{\mathbf{h}}$ will turn \mathbf{f}' into a hypertree of the form \mathbf{g}'' for some \mathbf{g} . The hypertree \mathbf{f}' is such that valence cannot be transferred from $a_{\mathbf{h}}$ to e_1 because \mathbf{h} was such that valence could not be transferred from e_2 to e_1 . A transfer of valence from $a_{\mathbf{h}}$ to e_2 turns \mathbf{f}' into \mathbf{h} , which is lonely and therefore of the form \mathbf{g}'' . If \mathbf{f}' is such that a transfer of valence is possible from $a_{\mathbf{h}}$ to some hyperedge $e_2 < b < a_{\mathbf{h}}$ (resulting in the hypertree \mathbf{j}) then, by the way $a_{\mathbf{h}}$ was chosen in the definition of $t(\mathbf{h})$, the hypertree \mathbf{h} is such that valence cannot be transferred from e_1 to b . In other words, \mathbf{j} is such that valence cannot be transferred from e_1 to e_2 which, by the characterization (b), establishes our claim.

To see why $\bar{t}(\mathbf{f}') \leq \bar{t}'(\mathbf{f}) + |T_{\mathbf{f}}|$ holds too, we let $x \in E$ be internally inactive with respect to \mathbf{f}' and show that it is either internally inactive with respect to \mathbf{f} as well or it is one of the $a_{\mathbf{h}}$ for some $\mathbf{h} \in T_{\mathbf{f}}$. By (a), \mathbf{f}' is such that valence can be transferred from e_1 to e_2 resulting in some hypertree \mathbf{k} . If \mathbf{f}' is such that a downward transfer of valence from x turns it into \mathbf{f}'_0 for some \mathbf{f}_0 then x is internally inactive with respect to \mathbf{f} and we are done. By Lemma 5.2, this kind of situation arises if \mathbf{f}' is such that a transfer of valence is possible from x to e_1 . Assume now that this is not possible but \mathbf{f}' is such that valence can be transferred from x to a hyperedge $e_2 < y < x$ resulting in the hypertree \mathbf{j} . We may assume that \mathbf{j} is such that valence cannot be transferred from e_1 to e_2 . If we apply Lemma 7.4 with $\mathbf{x} = \mathbf{j}$ and e_1, y as givers and e_2, x as receivers, we obtain that \mathbf{j} is such that y can transfer valence to e_2 . In other words, \mathbf{f}' is such that a transfer of valence is possible from x to e_2 (which would have been the last case to be examined anyway). Let the result of that transfer be the hypertree \mathbf{h} . We may assume that \mathbf{h} is not of the form \mathbf{f}'_0 for any \mathbf{f}_0 , i.e., that \mathbf{h} is such that e_1 cannot transfer valence to e_2 .

Since we have assumed that \mathbf{f}' was such that x cannot transfer valence to e_1 , we see that \mathbf{h} is such that e_2 cannot transfer valence to e_1 . Hence \mathbf{h} is lonely. Because \mathbf{h} is such that both e_1 and e_2 can transfer valence to x , we have $a_{\mathbf{h}} \leq x$. If $a_{\mathbf{h}} = x$ and therefore $t(\mathbf{h}) = \mathbf{f}'$, we are done; if $a_{\mathbf{h}} \neq x$ then \mathbf{f}' is such that x can transfer valence to $a_{\mathbf{h}}$ (the result being $t(\mathbf{h})$, which is of the form \mathbf{f}'_0 for some \mathbf{f}_0), showing that x is internally inactive with respect to \mathbf{f} . This finishes the proof of (3).

What remains is to examine the set $T_{\mathbf{f}} \cup \{\mathbf{f}'\}$ of lattice points. Recall that $\mathbf{f}' = t(\mathbf{h})$ and the hyperedge $a_{\mathbf{h}}$ determine the hypertree \mathbf{h} . Let us label the elements of $T_{\mathbf{f}}$ so that $a_{\mathbf{h}_1} > a_{\mathbf{h}_2} > \dots > a_{\mathbf{h}_{|T_{\mathbf{f}}|}}$ holds. Other than the ones belonging to e_2 and the $a_{\mathbf{h}_i}$, $i = 1, 2, \dots, |T_{\mathbf{f}}|$, all components

of the elements of $T_f \cup \{f'\}$ are identical. The remaining components are almost the same too as each h_i is derived from f' by a transfer of valence from a_{h_i} to e_2 . From this description it is clear that $T_f \cup \{f'\}$ is the set of vertices of an inverted $|T_f|$ -dimensional unit simplex. We need to show two things about these points.

First, if a hyperedge x is not one of the a_h for $h \in T_f$ then it is internally active either with respect to all elements of $T_f \cup \{f'\}$ or with respect to none of those. In other words for any $h \in T_f$, we have that x is internally inactive with respect to h if and only if it is internally inactive with respect to f' . Let first x be internally inactive with respect to h . If h is such that a transfer of valence is possible from x to one of the e_i , then an application of Lemma 5.2 at h shows that it is such that x can transfer valence to a_h (when $i = 1$, we use the fact that $k = f' - i_{\{e_1\}} + i_{\{e_2\}}$ is a hypertree). In other words, f' is such that x can transfer valence to e_2 . If h is such that x can transfer valence to a hyperedge $e_2 < y < x$, then by applying Lemma 7.4 to f' with a_h and x as givers and e_2 and y as receivers, we see that x is internally inactive there.

Let now x be internally inactive with respect to f' . If f' is such that a transfer of valence is possible from x to e_1 , then by Lemma 5.2, the same is possible to e_2 . If f' is such that x can transfer valence to some $e_2 < y < x$, then apply Lemma 7.4 at h with e_2 and x as givers and a_h and y as receivers to show that h is such that either x can transfer valence to y (in which case we are done) or x can transfer valence to e_2 . So what is left is to study the case when $j = f' - i_{\{x\}} + i_{\{e_2\}} = h - i_{\{x\}} + i_{\{a_h\}}$ is a hypertree and to show that x is internally inactive with respect to h .

The conclusion is immediate if $a_h < x$ so we assume $a_h > x$. If j is not lonely then it is easy to apply Lemma 7.4 (coupled with the fact that h is lonely) to show that h is such that x can transfer valence to e_1 or to e_2 . If j is lonely, then it has a corresponding hyperedge a_j . Because j is such that both e_1 and e_2 can transfer valence to x (the results being f' and k , respectively), we have $a_j \leq x$; but in fact $a_j < x$ for otherwise we would have $j \in T_f$ but we assumed x was not of the form a_h for any $h \in T_f$. Now apply Lemma 7.4 to h with x and e_i as givers and a_h and a_j as receivers, where i can be 1 or 2. In at least one of the two cases, e_i cannot transfer valence to a_j because of the definition of a_h (recall that $a_j < x < a_h$). So x can transfer valence to a_j and thus it is internally inactive with respect to h .

Second, we claim that with respect to h_i , the hyperedges $e_2, a_{h_{|T_f|}}, \dots, a_{h_i}$ are internally active and $a_{h_{i-1}}, \dots, a_{h_1}$ are internally inactive. This can again be separated into two parts. One is that a_h is internally active with respect to h for all $h \in T_f$. The other is that if $h \neq j$ are in T_f with $a_h > a_j$, then a_h is internally inactive with respect to j , whereas a_j is internally active with respect to h . As to the first claim, if h is such that a_h can transfer valence to e_1 or e_2 , then the existence of f' and k in \mathcal{QH} and Lemma 5.2 lead to a contradiction with the fact that h is lonely. If the transfer from a_h is to some hyperedge $e_2 < x < a_h$, then again Lemma 5.2 leads to the conclusion that h is such that both e_1 and e_2 can transfer valence to x , which contradicts the definition of a_h .

Let now $h \neq j$ be elements of T_f with $a_h > a_j$. As j is such that a_h can transfer valence to a_j (the result being h), a_h is indeed internally inactive with respect to j . Assume now that a_j is internally inactive with respect to h . If h is such that a_j can transfer valence to e_1 or e_2 , then it is easy to deduce from Lemma 7.4 (applied at $x = f'$ or k with givers a_h, a_j and receivers e_1, e_2) the contradiction that j or h cannot be lonely. If h is such that a_j can transfer valence to some hyperedge $e_2 < x < a_j$, resulting in the hypertree p , then we arrive at a contradiction as follows. We may assume that p is lonely for otherwise we can apply Lemma 7.4 with $x = h$ (the givers are a_j and one of e_1 or e_2 , the receivers are x and the other of e_1 and e_2 , and we use the

loneliness of \mathbf{h}) to show that \mathbf{h} is such that a_j can transfer valence to e_1 or e_2 . Note also that \mathbf{p} can be obtained from \mathbf{j} by a transfer of valence from a_h to x .

Again from Lemma 7.4, this time applied at $\mathbf{x} = \mathbf{f}'$ with givers a_h, a_j and receivers e_2, x , we deduce that \mathbf{f}' is such that either a_h or a_j can transfer valence to x . A similar conclusion can be drawn at the hypertree \mathbf{k} (now the receivers are e_1 and x). If the same option held true at both \mathbf{f}' and \mathbf{k} , then we would get a contradiction with the definition of a_h or a_j . Let us therefore assume that \mathbf{f}' is such that a_h cannot transfer valence to x and that \mathbf{k} is such that a_j cannot transfer valence to x ; the other option, with the roles of a_h and a_j interchanged, can be handled in the exact same way. In other words, \mathbf{h} is such that e_1 can transfer valence to x but e_2 cannot. Likewise, \mathbf{j} is such that e_2 can transfer valence to x but e_1 cannot. This means that there exists a set U_1 of hyperedges that is tight at \mathbf{h} , contains x , and does not contain e_2 . By passing to a smaller set if necessary, we may assume that U_1 is such that $\text{Bip } \mathcal{H}|_{U_1}$ is connected (cf. the first part of Lemma 4.11). Because of our assumed transfers, e_1 and a_j are also elements of U_1 . Similarly, we have a set $U_2 \subset E$ with $\text{Bip } \mathcal{H}|_{U_2}$ connected that is tight at \mathbf{j} , contains x, e_2 , and a_h but does not contain e_1 .

It is easy to see that both U_1 and U_2 remain tight at \mathbf{p} . Let Γ be an arbitrary realization of \mathbf{p} . By Lemma 4.11, Γ intersects $\text{Bip } \mathcal{H}|_{U_1}$ and $\text{Bip } \mathcal{H}|_{U_2}$ in trees. In particular, Γ contains a path from x to e_1 which misses e_2 as well as a path from x to e_2 which misses e_1 . Since \mathbf{p} is lonely, Γ also contains the two-edge path from e_1 to e_2 through v . The union of the three paths contradicts the assumption that Γ is a tree.

This completes the proof of (4) and hence that of the Proposition. \square

Theorem 7.6. *Let G be a connected bipartite graph with a vertex u that is only joined by edges to the vertices v and w . If the graphs $G - \{u\}$ (where we erase u and the edges uv and uw from G) and G/vuw (where we contract the path vuw to a point) both satisfy Conjecture 7.1, then so does G .*

Proof. Immediate from a comparison of Propositions 6.14 and 7.3. \square

Theorem 7.6 and the list of operations before Proposition 7.3 reduce Conjecture 7.1 to connected bipartite graphs that cannot be disconnected by the removal of a vertex or two adjacent vertices, and in which each vertex is at least trivalent. One such graph is discussed below.

Example 7.7. Consider the bipartite graph shown in Fig. 8. Both of its induced hypergraphs have the interior polynomial

$$1 + 10\xi + 48\xi^2 + 146\xi^3 + 302\xi^4 + 410\xi^5 + 277\xi^6 + 49\xi^7 + \xi^8.$$

This is a non-trivial example of Conjecture 7.1 because the structure is not symmetric: even though both color classes have nine elements with matching valences, there is no isomorphism that interchanges them. This is demonstrated by the slight difference between the corresponding exterior polynomials. If the full dots play the role of hyperedges, we obtain

$$1 + 8\eta + 36\eta^2 + 110\eta^3 + 235\eta^4 + 344\eta^5 + 318\eta^6 + 162\eta^7 + 30\eta^8,$$

whereas if the hollow dots represent hyperedges, we get

$$1 + 8\eta + 36\eta^2 + 110\eta^3 + 235\eta^4 + 348\eta^5 + 326\eta^6 + 159\eta^7 + 21\eta^8.$$

All three polynomials are outputs of a computer code by Péter Juhász.




Fig. 8. The graph of Example 7.7.

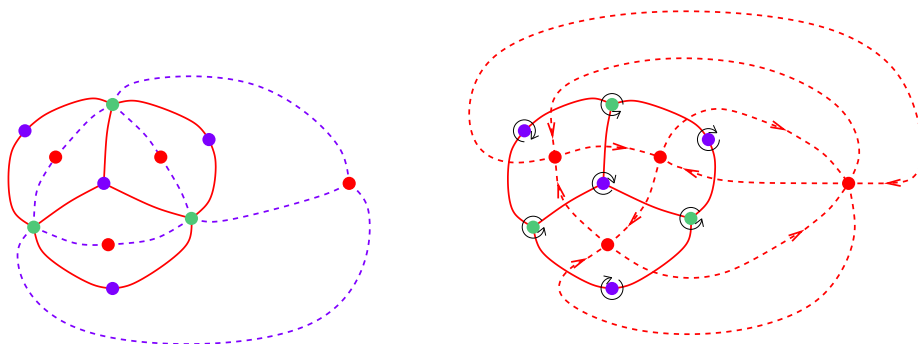


Fig. 9. Left: the planar dual of a plane hypergraph \mathcal{H} of three hyperedges. Right: the planar dual of the corresponding plane bipartite graph $\text{Bip } \mathcal{H}$.

8. Planar duality

We call a hypergraph $\mathcal{H} = (V, E)$ *planar* if the corresponding bipartite graph $\text{Bip } \mathcal{H}$ is planar, i.e., if it admits an embedding into the 2-sphere S^2 . In the rest of the paper we will almost always assume that such an embedding is fixed, but for extra clarity, we will talk about *plane* (hyper)graphs to mean a graph together with a particular embedding. A *region* of a plane (hyper)graph is a connected component of the complement of the image of the embedding. The set of regions will customarily be denoted with R . If the graph is non-empty and connected, each region is homeomorphic to a disk.

For the rest of the paper, it will be convenient to allow multiple edges in (plane) bipartite graphs. Such objects induce the same two hypergraphs (and therefore the same two hypertree polytopes and the same interior and exterior polynomials) as the bipartite graph which results from reducing each set of multiple edges to a single edge. This is all fairly natural given that a spanning tree may contain at most one element from a set of multiple edges.

Definition 8.1. Let $\mathcal{H} = (V, E)$ be a plane hypergraph so that its associated bipartite graph $\text{Bip } \mathcal{H}$ is connected. We define the *planar dual hypergraph* \mathcal{H}^* of \mathcal{H} by keeping the set E of hyperedges and letting the new set of vertices be the set R of regions of \mathcal{H} . We let a region belong to a hyperedge if it is incident with the point representing it on the plane.

The planar dual hypergraph \mathcal{H}^* is also planar because $\text{Bip } \mathcal{H}^*$ can be represented by a planar diagram where we place a new vertex in each region $r \in R$ and connect it to all points around the boundary ∂r that represent elements of E . See the left panel of Fig. 9 for an example. In fact

we will identify R and the set of points just introduced and think of the planar dual hypergraph as the plane hypergraph with the embedding that we have just described.

It is easy to see that our notion generalizes the usual duality of plane graphs. Indeed, a plane graph $G = (V, E)$ is also a hypergraph. Passing to $\text{Bip } G$ means placing a new vertex to the midpoint of each edge. Now in G^* , these same midpoints are connected to the points which represent the two regions on either side of the original edge. Finally, if we forget the midpoint, these two connections merge to form the usual dual edge.

Proposition 8.2. *For a plane hypergraph \mathcal{H} , the planar dual of the planar dual of \mathcal{H} is naturally isomorphic to \mathcal{H} .*

Proof. By an examination of the left panel of Fig. 9. We explained after Definition 8.1 how $\text{Bip } \mathcal{H}^*$ is embedded in S^2 . Each point representing a vertex $v \in V$ is in some region of $\text{Bip } \mathcal{H}^*$. It is easy to see that this correspondence is onto and one-to-one. \square

The hypertree polytopes $Q_{\mathcal{H}}$ and $Q_{\mathcal{H}^*}$ of a planar dual pair are subsets of the same Euclidean space \mathbf{R}^E . Next, we address the effect of planar duality on the hypertree polytope, as well as the interior and exterior polynomials.

Theorem 8.3. *Let the hypergraphs $\mathcal{H} = (V, E)$ and $\mathcal{H}^* = (R, E)$ be planar duals. Then $Q_{\mathcal{H}}$ and $Q_{\mathcal{H}^*}$ are reflections of each other in a point. Furthermore, we have $I_{\mathcal{H}^*} = X_{\mathcal{H}}$ and $X_{\mathcal{H}^*} = I_{\mathcal{H}}$.*

In Fig. 9, left panel we see that the hypergraph \mathcal{G}_0 of our earlier examples is planar and in fact planar self-dual. This explains why (10) showed its interior and exterior polynomials to coincide.

The following argument contains simplifications made by A. Bene.

Proof of Theorem 8.3. Proposition 8.2 implies that for the claim on the polytopes, it suffices to show that if $\mathbf{f}: E \rightarrow \mathbf{N}$ is a hypertree in \mathcal{H} then $\mathbf{f}^*(e) = |e| - 1 - \mathbf{f}(e)$ defines a hypertree in \mathcal{H}^* . (So the center in which $Q_{\mathcal{H}}$ and $Q_{\mathcal{H}^*}$ are symmetric is the point whose e -coordinate is $\frac{1}{2}(|e| - 1)$ for all $e \in E$.)

To this end, choose a spanning tree $\Gamma \subset \text{Bip } \mathcal{H}$ that realizes \mathbf{f} . Form its planar dual Γ^d in the classical sense, i.e., put a vertex in each element of R and connect these by edges that bisect each edge of $\text{Bip } \mathcal{H}$ that is not an edge in Γ . This is well known to be a spanning tree in the dual graph, in particular it is cycle-free, connected, and contains all points of R . We are going to use a continuous deformation followed by some discrete steps to turn Γ^d into a spanning tree $\Gamma^* \subset \text{Bip } \mathcal{H}^*$ that realizes \mathbf{f}^* .

Each edge γ^d of Γ^d bisects an edge γ of $\text{Bip } \mathcal{H}$ and γ has exactly one end e in E . Let us push the midpoint of γ^d along γ so that it gets very close to e . If $r \in R$ and $e \in E$ are adjacent with two edges γ_1, γ_2 of $\text{Bip } \mathcal{H}$, then we merge the halves of γ_1^d and γ_2^d adjacent to r , as shown in the right side of Fig. 10. The result of the modifications so far is still a plane tree which we will continue to denote with Γ^d .

Now let $e \in E$ be arbitrary. Write the edges of $\text{Bip } \mathcal{H}$ emanating from e , in a cyclic order determined by the embedding, as $\gamma_1, \dots, \gamma_{|e|}$. Let the edges among these that do not belong to Γ form k maximal subsequences of consecutive elements. Let the lengths of these subsequences be s_1, \dots, s_k . In our construction so far, each of these k groups resulted in a point e_i near e with $s_i + 1$ edges connecting it to elements of R . See Fig. 10, which shows $k = 3$ with $s_1 = 2, s_2 = 1,$ and $s_3 = 4$. We wish to push these points e_1, \dots, e_k all the way to e so that the edges adjacent to them become edges of $\text{Bip } \mathcal{H}^*$.




Fig. 10. Modifying a tree Γ^d in $(\text{Bip } \mathcal{H})^*$ toward a tree Γ^* in $\text{Bip } \mathcal{H}^*$. Solid edges represent $\text{Bip } \mathcal{H}$; the thick ones are in Γ and the thin ones are not.

The simplest case is that of $k = 1$ when we can just identify e_1 with e . Note that after this, the number of edges of Γ^d adjacent to e is $s_1 + 1 = |e| - (\mathbf{f}(e) + 1) + 1 = \mathbf{f}^*(e) + 1$. When $k \geq 2$, we still start with pushing e_1 to e . When we do the same to e_2 , however, a unique cycle is created in Γ^d and we remove one of its two edges adjacent to e in order to have a tree again. Then we push e_3 to e and remove another edge, and so forth until the last point e_k . In the tree that results at the end, the valence of e is $(s_1 + 1) + s_2 + \dots + s_k = |e| - (\mathbf{f}(e) + 1) + 1 = \mathbf{f}^*(e) + 1$. After applying this procedure to all elements of $e \in E$ with $k \geq 1$, we arrive at a tree that is adjacent to all such hyperedges as well as to all elements of R .

We intentionally left out the case $k = 0$ which of course corresponds to $\mathbf{f}^*(e) = 0$. To finish the construction of the spanning tree $\Gamma^* \subset \text{Bip } \mathcal{H}^*$ realizing \mathbf{f}^* , we add one more edge for each such hyperedge so that it is connected to an arbitrarily chosen adjacent element of R .

The claim on the polynomials is now easily obtained by the following observation. The correspondence $\mathbf{f} \leftrightarrow \mathbf{f}^*$ is a bijection between hypertrees in \mathcal{H} and \mathcal{H}^* . The hypertree \mathbf{f} can be transformed into \mathbf{g} by a single transfer of valence if and only if \mathbf{f}^* can be transformed into \mathbf{g}^* by the opposite transfer. Therefore with respect to any order of the hyperedges, we have $\bar{i}(\mathbf{f}^*) = \bar{e}(\mathbf{f})$ and $\bar{e}(\mathbf{f}^*) = \bar{i}(\mathbf{f})$. \square

Definition 8.4. If the hypertree \mathbf{f} in the plane hypergraph \mathcal{H} and the hypertree \mathbf{f}^* in \mathcal{H}^* are related as above (namely, $\mathbf{f}(e) + \mathbf{f}^*(e) = |e| - 1$ for all hyperedges e), then we will call them *planar dual hypertrees*.

9. Trinities

9.1. Basic observations

Starting from a plane hypergraph and iterating the constructions of planar and abstract duality, a total of six plane hypergraphs can be built. The corresponding bipartite graphs form a triple; see Fig. 11 for an example. The picture thus obtained has a rich combinatorics and admits several equivalent definitions (cf. Remark 9.5). To highlight the perfect symmetry between the constituent parts, we chose to build our discussion around the following notion.

Definition 9.1. A *trinity* is a triangulation of the sphere S^2 together with a three-coloring of the 0-simplices. (I.e., 0-simplices joined by a 1-simplex have different colors.) According to dimension, we will refer to the simplices as *points*, *edges*, and *triangles*.

Most claims of this subsection can be repeated for three-colored triangulations of other closed, orientable surfaces. The material in the rest of the section, however, is quite specific to the sphere. We hope to return to this point in our forthcoming paper with Bene.

We will use the names red, emerald, and violet for the colors in the trinity and denote the respective sets of points with R , E , and V . Let us color each edge in the triangulation with the color that does not occur among its ends. Then E and V together with the red edges form a bipartite graph that we will call the *red graph* and denote with G_R . Each region of the red graph contains a unique red point. Likewise, the *emerald graph* G_E has red and violet points, emerald edges, and regions marked with emerald points. Finally, the *violet graph* contains R and E as vertices, violet edges, and a violet point in each of its regions.

Definition 9.2. We will refer to the plane bipartite graphs G_R , G_E , and G_V above as the *constituent bipartite graphs* of the trinity. The hypergraphs induced by the constituent bipartite graphs are said to be *contained* in the trinity.

A trinity contains six plane hypergraphs. It can be uniquely reconstructed from each of the six as follows.

Proposition 9.3. *Given a trinity with point set $R \cup E \cup V$ as above, let $\mathcal{H} = (V, E)$ be the hypergraph with emerald hyperedges and violet vertices so that $\text{Bip } \mathcal{H} = G_R$. Then we also have $(R, E) = \overline{\mathcal{H}^*}$, $(E, R) = \overline{\mathcal{H}^*}$, $(V, R) = \left(\overline{\mathcal{H}^*}\right)^* = \overline{(\mathcal{H})^*}$, $(R, V) = \left(\overline{\mathcal{H}}\right)^*$, and $(E, V) = \overline{\mathcal{H}}$.*

Proof. Obvious from an examination of Fig. 11. □

The triangles of a trinity can also be colored but not with the original three colors. Notice that each triangle is adjacent with exactly one edge and one point of each color. Compared to the orientation of the sphere, the cyclic order of the colors around each triangle may be positive or negative. If two triangles share an edge, these orientations are opposite. Hence the triangles have a black and white checkerboard coloring according to orientation, cf. Fig. 11. In particular, the dual graph of a trinity is a plane bipartite cubic (i.e., uniform trivalent) graph. (We remark that the graph used in Example 7.7 belongs to this class.) It turns out that the converse is also true.

Proposition 9.4. *Planar duals of plane bipartite cubic graphs are three-colorable.*

Proof. Let θ be a plane bipartite cubic graph and let us call its color classes black and white. Direct each edge in θ from its black to its white endpoint. Now if we assign the modulo 3 coefficient 1 to the thus directed edges, the result is a cycle C and hence a homology class in $H_1(S^2, \mathbf{Z}_3) = 0$. Therefore the modulo 3 intersection number of any closed loop with C is zero.

Choose a base region of θ and assign the remainder class $0 \in \mathbf{Z}_3$ to it. Given any other region, connect it to the base with a path in S^2 and assign to the region the modulo 3 intersection number of the path and C . By the above, this is well defined and it is clearly a three-coloring with the colors 0, 1, and 2. □

Remark 9.5. According to Proposition 9.4 and the paragraph above it, the notion of a trinity is equivalent to that of a plane bipartite cubic graph.

Finally, notice that the sets of red edges, emerald edges, violet edges, white triangles, and black triangles all have the same cardinality n . In particular, adjacency defines natural bijections




Fig. 11. A trinity of plane bipartite graphs.

between white triangles and edges of each color. Now if we apply Euler’s formula to the trinity, we get $|R| + |E| + |V| - 3n + 2n = 2$, that is that in any trinity,

$$\text{the total number of points exceeds that of the white triangles by 2.} \tag{15}$$

9.2. The Tree Trinity Theorem

Let T be a trinity. Let us take the planar dual G_R^* of the red graph G_R in the classical sense, i.e., the vertex set of G_R^* is R and its edges are in a one-to-one correspondence with the red edges of the trinity. This graph has a natural orientation (more precisely, two natural orientations that are opposites of each other), defined as follows. Put a positive spin (as in a small spinning top) to the emerald points and a negative spin to the violet points. If $e \in E$ and $v \in V$ are connected by a (red) edge, then the two spins induce the same orientation on the dual edge. See the right panel of Fig. 9 for an illustration. Notice that at each red point r , the edges of G_R^* that are adjacent to it are oriented toward and away from r in an alternating fashion. In particular, each vertex of G_R^* has the same number of incoming and outgoing edges.

Similar properties hold for the graphs G_E^* and G_V^* . It is in fact this triple of directed graphs that Tutte called a trinity. In [7], he proved a general property of directed graphs which will be stated after the next definition.

Definition 9.6. Let D be a finite directed graph (possibly with multiple edges) and fix a vertex $r \in D$, called the *root*. A spanning tree in D is called a (*spanning*) *arborescence* with respect to

r (or *rooted at r*) if the unique path in the tree from r to any other vertex is oriented toward that vertex.

A directed graph is called *balanced* if each of its points has the same in-degree as its out-degree.

Theorem 9.7 (Tutte). *Let D be a finite, balanced directed graph. The number of its spanning arborescences does not depend on the choice of root. The same holds for trees that are oriented toward the root and the counts of the two kinds of tree coincide.*

The resulting invariant of balanced directed graphs is denoted with $\rho(D)$ and called the *arborescence number*. The following will be a useful characterization.

Lemma 9.8. *Spanning arborescences of a finite directed graph with root r are exactly those cycle-free subgraphs which also have the property that for any vertex other than r , exactly one edge of the subgraph is directed toward that vertex.*

Proof. Let $D = (V, E)$ be a finite directed graph. If it is disconnected then both properties above describe the empty set. If D is connected and A is a spanning arborescence with respect to r , then of course it is cycle-free and it has at least one edge directed into any vertex other than r , namely the last edge on the path from r to that vertex. But because A contains exactly $|V| - 1$ edges, these already exhaust all of them.

Conversely, if a subgraph is cycle-free than it may have at most $|V| - 1$ edges. If it has one incoming edge for all points other than r , then it has to have exactly $|V| - 1$ of them (so it is a spanning tree) and none can be directed into r . For any vertex $v \neq r$, the first edge of the unique path from r to v has to be directed toward v . The same holds for the second edge lest they share terminal points with the first. Etc. \square

Let us now return to trinitities and recall a beautiful result by Tutte.

Theorem 9.9 (Tree Trinity Theorem). *The arborescence numbers of the three directed graphs associated to a trinity are the same: $\rho(G_R^*) = \rho(G_E^*) = \rho(G_V^*)$.*

We will call this common value the *arborescence number of the trinity* and denote it with $\rho(T)$. We quickly sketch a proof, based on [9], as elements of it will be relevant later.

Proof. Recall that the triangles of a trinity have a black-and-white color scheme. We may concretize our choice of orientation in the three directed graphs by requiring that the head of each edge be in a white triangle while the tail end of each is in a black triangle. Compare Fig. 9 (right panel) and 12. Then by Lemma 9.8, a spanning arborescence in one of our directed graphs, say G_R^* , is an assignment of an adjacent white triangle to each non-root red point (the unique incoming edge will reach the point through the chosen triangle). We adapt this point of view for the rest of the proof.

Let us now distinguish a white triangle t_0 (called the outer triangle) and choose its adjacent points $r_0 \in R$, $e_0 \in E$, and $v_0 \in V$ as the roots in the three directed graphs. Then t_0 can never be selected into an arborescence of any color. Notice that by (15), the rest of the triangles and the rest of the points form equinumerous sets.

Fix a spanning arborescence A in G_R^* . We are going to associate to it two other spanning arborescences, in G_E^* and G_V^* respectively, so that

$$\begin{aligned} &\text{the union of the three mappings is a bijection from the set of non-root points to} \\ &\text{the set of non-outer white triangles.} \end{aligned} \tag{16}$$




Fig. 12. The trinity of Fig. 11 and the directed graph G_R^* .

In other words, no white triangle will get assigned to two different points. We also claim that, given A , there is a unique way to do this.

The dual A^* of A is a spanning tree in G_R . This tree A^* consists of those edges whose duals are not in A , in other words of those red edges whose adjacent white triangles have not been assigned to a red point. Note that the edge connecting e_0 and v_0 is part of A^* .

As A^* is a tree, it has at least two leaves, i.e., valence one points. If those are e_0 and v_0 , then A^* consists of a single edge. Hence there cannot be any other emerald or violet points and we have no assigning to do. Otherwise, as each leaf is adjacent to a unique white triangle that has not yet been assigned, it makes sense to pick a leaf in A^* that is neither e_0 nor v_0 , make the obvious assignment, and delete the leaf (with its single adjacent edge) from A^* . We continue this procedure until A^* is reduced to the single edge of t_0 between e_0 and v_0 . By that time, as A^* is spanning, all other emerald and violet points have been associated to an adjacent white triangle so that no white triangle was used twice.

We have to check (cf. Lemma 9.8) that the subgraphs we constructed in G_E^* and G_V^* are cycle-free. This is shown using proof by contradiction, arguing that a violet or emerald cycle would prevent some red point from being connected to the red root r_0 in A . See [9] for details.

As the color red played no special role above, any triple of spanning arborescences satisfying (16) can be uniquely reconstructed from any of its three members. In particular, we see that the choice of t_0 specifies a bijection between any two of the three sets of arborescences. \square

A mapping satisfying (16) will be called a *Tutte matching* (with respect to some fixed outer white triangle t_0). Tutte matchings are exactly the non-zero expansion terms of a determinant that first appeared in the work of Berman.

Theorem 9.10 (Berman). *The common arborescence number of the three directed graphs of a trinity is the absolute value of the determinant of the adjacency matrix between non-outer white triangles and non-root points.*

This is a square matrix by the observation (15). In his proof, Berman [1] checks that any Tutte matching is the union of three spanning arborescences (which boils down to cycle-freeness and another argument by contradiction) and that all non-zero expansion terms have the same sign.

Example 9.11. The adjacency matrix associated with the trinity (and triangle t_0) of Fig. 11 is

$$M = \begin{matrix} & t_1 & t_2 & t_3 & t_4 & t_5 & t_6 & t_7 & t_8 \\ \begin{matrix} r_1 \\ r_2 \\ r_3 \\ e_1 \\ e_2 \\ v_1 \\ v_2 \\ v_3 \end{matrix} & \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0. \end{pmatrix} \end{matrix}$$

Its determinant is 7, which equals the number of spanning arborescences in each of the three dual directed graphs, such as the one shown in Fig. 12.

10. Hypertrees in a trinity

10.1. Hypertrees and arborescences

The main result of this last section is the following.

Theorem 10.1. *Let $G = (V_0, V_1, E)$ be a plane bipartite graph. The arborescence number of its planar dual G^* agrees with the number of hypertrees in the hypergraphs $\mathcal{G}_0 = (V_1, V_0)$ and $\mathcal{G}_1 = (V_0, V_1)$ induced by G . I.e.,*

$$\rho(G^*) = |Q_{\mathcal{G}_0} \cap \mathbf{Z}^{V_0}| = |Q_{\mathcal{G}_1} \cap \mathbf{Z}^{V_1}|.$$

Proof. Let us denote the vertex set of G^* with R and fix a root $r_0 \in R$. Any spanning arborescence A rooted at r_0 has a planar dual spanning tree A^* in G and that in turn induces a hypertree $\mathbf{f}_A: V_0 \rightarrow \mathbf{N}$. We will show that the mapping $A \mapsto \mathbf{f}_A$ is one-to-one and onto. (A similar statement holds of course for hypertrees in \mathcal{G}_1 instead of \mathcal{G}_0 . Thus what we find as a byproduct is a set of spanning trees in G that simultaneously realize all hypertrees in \mathcal{G}_0 and \mathcal{G}_1 .)

For any $v \in V_0$, the edges of G adjacent to it give rise to a directed cycle C_v in G^* and the edge set of G^* is the disjoint union of these cycles. Note that each cycle C_v travels around the corresponding point v following the same orientation.

First we show that if two arborescences A and B contain the same number of edges from each C_v (which is equivalent to $\mathbf{f}_A = \mathbf{f}_B$), then $A = B$. Assume this is not so. Then at least one cycle C_{v_1} contains an edge α_1 of A that is not an edge in B . Let the head of α_1 be r_1 . Obviously $r_1 \neq r_0$ because A has an edge pointing to it. The other arborescence B also has to have an edge β_1 pointing to r_1 but that has to be part of a different cycle C_{v_2} . The edge β_1 cannot belong to A because then A would have two edges pointing to r_1 . Then, because A and B intersect C_{v_2} in the same number of edges, C_{v_2} also contains an edge α_2 which belongs to A but does not belong to B .

Iterating our argument, we find cycles $C_{v_1}, C_{v_2}, C_{v_3}, \dots$ in G^* , along with edges $\alpha_i \in C_{v_i} \cap (A \setminus B)$ and $\beta_i \in C_{v_{i+1}} \cap (B \setminus A)$, until the first repetition occurs. After discarding some cycles at the beginning and relabeling the rest, we may assume that $C_{v_{k+1}} = C_{v_1}$ (and $\alpha_{k+1} = \alpha_1$) was the first such coincidence. The situation is depicted in Fig. 13. The interiors of the cycles C_{v_i} are disjoint from G^* , hence the non-root points r_1, \dots, r_k separate G^* into the smaller directed graphs D_1 and D_2 , see Fig. 13, which have only those points in common. Both




Fig. 13. A cycle of cycles in the directed graph G^* .

D_1 and D_2 have to contain vertices other than r_1, \dots, r_k because otherwise A would contain the cycle $\alpha_1, \dots, \alpha_k$ or B would contain the cycle β_k, \dots, β_1 . Only one of D_1 or D_2 , however, can contain the root r_0 . Now the contradiction is apparent from the fact that no directed path in A can pass from D_1 to D_2 and no directed path in B can pass from D_2 to D_1 .

To prove that any hypertree $\mathbf{f}: V_0 \rightarrow \mathbf{N}$ can be obtained in the form $\mathbf{f} = \mathbf{f}_A$, we employ the following strategy. Start from an arbitrary spanning tree Γ in G that realizes \mathbf{f} and take its dual tree $\Gamma^* \subset G^*$. We are going to change Γ step-by-step through other realizations of the same hypertree \mathbf{f} so that the dual moves closer and closer to being an arborescence. It is sufficient to keep track of the changes made to the dual tree Γ^* . The condition that we preserve the hypertree translates to requiring that for each $v \in V_0$, the number of edges in $C_v \cap \Gamma^*$ stays invariant throughout the process.

It is of course crucial to describe carefully what we mean by getting closer to an arborescence. In a rooted tree, edges γ have a well defined distance $d(\gamma) \in \mathbf{N}$ to the root (those adjacent to the root have $d = 0$ etc.). If the tree is directed, then there is also a clear sense for each edge to be directed toward the root or away from the root. The first kind of edge will be called *bad* and the second kind *good*. The tree is an arborescence if and only if all of its edges are good. Let us associate the following quantities to a finite, directed, rooted tree (T, r) .

- Let $n(T, r)$ denote the smallest value of d among bad edges. (So that within a radius of $n(T, r)$ from the root, the tree is an arborescence.) If there are no bad edges, let $n(T, r)$ be the maximal distance from r to any point in T .
- For $1 \leq m \leq n(T, r)$, let $\lambda_{T,r}(m)$ be the number of edges γ with $d(\gamma) = m - 1$. These values are positive. For $m > n(T, r)$, we define $\lambda_{T,r}(m) = 0$.

Then, for a pair of rooted trees we write $(T_1, r_1) < (T_2, r_2)$ if either

1. the sequence $\lambda_{T_1,r_1}(1), \lambda_{T_1,r_1}(2), \lambda_{T_1,r_1}(3), \dots$ is smaller than the sequence $\lambda_{T_2,r_2}(1), \lambda_{T_2,r_2}(2), \lambda_{T_2,r_2}(3), \dots$ in lexicographic order, or
2. the two sequences coincide (implying $n(T_1, r_1) = n(T_2, r_2) = n$) but the number of bad edges with $d = n$ is higher in T_1 than in T_2 .

It is easy to see that $<$ is a partial order of rooted trees.

This may sound complicated but the actual idea of the proof is very simple. If Γ^* is not a spanning arborescence already, then it contains a bad edge γ with $d(\gamma) = n(\Gamma^*, r_0)$. If we remove γ from Γ^* , then the tree falls apart into a root component and a non-root component. Now γ is part of a cycle C_v which contains points from both components. Therefore C_v also has

an edge γ' that goes from a point of the root component to a point of the non-root component. Let $\Gamma^{*'}$ denote the tree obtained from Γ^* by replacing γ with γ' .

Note that γ' is a good edge of the new tree. The difficulty is that in what used to be the non-root component, distances to the root and relative orientations may have changed. What we do know though is that each of those edges is farther away from the root than γ' . Let us separate two cases according to whether the non-root component is reattached ‘close to’ or ‘far from’ the root.

- (i) If $d(\gamma') < n(\Gamma^*, r_0)$ (where $d(\gamma')$ is of course measured in $\Gamma^{*'}$), then $(\Gamma^*, r_0) < (\Gamma^{*'}, r_0)$ by (1). Indeed, with the addition of γ' , $\lambda(d(\gamma') + 1)$ went up by 1 whereas the λ values before it stayed the same.
- (ii) If $d(\gamma') \geq n(\Gamma^*, r_0)$, then $(\Gamma^*, r_0) < (\Gamma^{*'}, r_0)$ either by (1) (if γ was the unique bad edge of Γ^* with $d(\gamma) = n(\Gamma^*, r_0)$, in which case $\lambda(n(\Gamma^*, r_0) + 1)$ moves from 0 to a positive value) or by (2).

We conclude that if the dual of a realization of \mathbf{f} is not a spanning arborescence, then it is possible to change it to another tree, dual to another realization, that is larger in the sense of $<$. Because there are altogether finitely many spanning trees in G^* , it is now obvious that a spanning arborescence will be reached after finitely many improvements. Thus we find that our map $A \mapsto \mathbf{f}_A$ is onto. \square

Theorem 10.1 of course verifies Postnikov’s **Theorem 3.11** in the case of a planar bipartite graph G . Postnikov [4, Lemma 12.6] describes a sufficient condition on a set of spanning trees of G in order for them to define a triangulation of the root polytope and therefore to simultaneously realize all hypertrees of the two induced hypergraphs. It is interesting to note that our set of spanning trees satisfies his condition, as follows.

Proposition 10.2. *Let G be a plane bipartite graph with (classical) dual G^* as before. Let A and B denote two spanning arborescences of G^* with respect to the same root. Then there is no cycle $\alpha_1, \beta_1, \alpha_2, \beta_2, \dots, \alpha_k, \beta_k$ in G composed of edges $\alpha_1, \alpha_2, \dots, \alpha_k$ of A^* and $\beta_1, \beta_2, \dots, \beta_k$ of B^* .*

Proof. It is easy to see that edges of A could only cross such a cycle in one direction and edges of B could only cross it in the opposite direction. That leads to the usual contradiction of some points becoming inaccessible from the root for one of the two arborescences. \square

Tutte’s **Theorem 9.9** says that the three directed graphs in a trinity have the same number of spanning arborescences, just like the two graphs in a planar dual pair have the same number of spanning trees. We may now add that the six hypergraphs contained in the trinity have the same number of hypertrees. Namely, the following is immediate from **Theorem 10.1**.

Corollary 10.3. *The arborescence number ρ associated to a trinity is also the number of hypertrees in any of the six hypergraphs contained in the trinity.*

Let us make one more observation. We saw how each Tutte matching contains three spanning arborescences (one for each of G_R^* , G_E^* , and G_V^*) and how the (classical) planar duals of these realize hypertrees in each of the six hypergraphs. This way, each hypertree turns up in relation to exactly one Tutte matching. In fact we also have the following.

Proposition 10.4. *In any trinity, the six hypertrees induced by a Tutte matching form three pairs of planar duals (in the sense of **Definition 8.4**).*

Proof. Pick an arbitrary, say emerald point from the trinity and denote it with e . It is matched to one adjacent white triangle (if the point is the root, substitute the outer white triangle here). Out of its other $|e| - 1$ adjacent white triangles, suppose that m are matched to red points and $|e| - 1 - m$ to violet points. Then there are $|e| - m$ red edges and $m + 1$ violet edges adjacent to e that belong to the spanning trees of G_R and G_V , respectively, induced by the Tutte matching. This causes the e -coordinates of the corresponding hypertrees to be $|e| - m - 1$ and m , respectively, the sum of which is indeed $|e| - 1$. \square

10.2. *Determinant formulas*

We are going to extend Berman’s [Theorem 9.10](#) to write hypertree polytopes associated with plane hypergraphs in determinant form. Let M denote the adjacency matrix with rows indexed by the non-root points and columns indexed by the non-outer white triangles. If a triangle t_i is adjacent to the emerald point e_j and the non-root violet point v_k , then at the intersection of row v_k and column t_i , change the entry 1 to e_j . (After it becomes a matrix entry, we will treat e_j as an indeterminate associated with the original point.)

Call this matrix the *enhanced adjacency matrix* and denote it with $M_{e \rightarrow v}$. By varying the colors, altogether six such matrices can be associated to a trinity. Notice that even though M contains no row indexed by a root, indeterminates belonging to roots do appear in enhanced adjacency matrices.

It would be rather pointless to write the symbol e_j in the row indexed by e_j because that would simply multiply the determinant by e_j . The twist that we introduced, on the other hand, turns out to be quite useful.

Theorem 10.5. *Let the plane bipartite graphs G_R , G_E , and G_V be the constituents of a trinity. Fix an outer white triangle t_0 which is adjacent to the roots $r_0 \in R$, $e_0 \in E$, and $v_0 \in V$ and form the enhanced adjacency matrix $M_{e \rightarrow v}$ as above. Then for the hypergraph $\mathcal{H} = (V, E)$ and its hypertree polytope, we have*

$$Q_{\mathcal{H}} \cap \mathbf{Z}^E = \det M_{e \rightarrow v}$$

in the following sense. The determinant on the right hand side is a sum of monomials in the indeterminates $e \in E$. Either each monomial has coefficient $+1$ or each has -1 . If we write the exponents in the monomials as vectors, the set we obtain is exactly the left hand side.

As a special case, we may use this theorem to write the vertices of the spanning tree polytope of a plane graph as a determinant, too. To avoid abuse of notation, we may re-state the result as

$$\pm \det M_{e \rightarrow v} = \sum_{\mathbf{f} \in Q_{\mathcal{H}} \cap \mathbf{Z}^E} \prod_{e \in E} e^{\mathbf{f}(e)}.$$

However it is unlikely that the first form will lead to confusion.

Proof. The claim on the uniform signs of course follows from Berman’s [Theorem 9.10](#) on the adjacency matrix M . There and in the proof of [Theorem 9.9](#) we saw that the non-zero terms in the expansion of $\det M$, i.e., the Tutte matchings (of non-root points in the trinity to adjacent non-outer white triangles), are also triples of spanning arborescences in the directed graphs G_R^* , G_E^* , and G_V^* . Furthermore, all spanning arborescences from each of the three directed graphs occur as part of exactly one triple.

Then in the proof of [Theorem 10.1](#) we found that the planar duals A^* of the spanning arborescences A of G_R^* realize each hypertree in $Q_{\mathcal{H}}$ exactly once. Recall that the value of

the hypertree f_A at an emerald point e_j is the number of adjacent (to e_j) edges of A^* minus one. A red edge is in A^* if and only if its adjacent white triangle is not assigned to its red point, i.e., if the triangle is the outer one or if it is assigned to an emerald or violet point.

If $e_j \neq e_0$ then e_j is not adjacent to t_0 and out of all white triangles adjacent to e_j , exactly one is assigned to an emerald point (namely, e_j). Thus $f_A(e_j)$ is the number of white triangles that are adjacent to e_j and which are assigned to their violet points. The same holds for e_0 where now t_0 plays the role of the correction term -1 .

Finally, we just have to check the effect of enhancing M into $M_{e \rightarrow v}$ on the expansion term that corresponds to A . In that monomial, by definition, the indeterminate e_j will appear once for each time that a (non-outer) white triangle t_i , which is adjacent to the point e_j , gets assigned to its adjacent violet point v_k . This completes the proof. \square

Example 10.6. The hypertree polytopes of Example 3.9 can be recomputed as the determinants of the following two matrices $M_{e \rightarrow v}$ and $M_{v \rightarrow e}$ (warning: the labels of the points have changed, cf. Fig. 11):

$$\begin{matrix}
 & t_1 & t_2 & t_3 & t_4 & t_5 & t_6 & t_7 & t_8 \\
 r_1 & \left(\begin{array}{cccccccc}
 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\
 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 \\
 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\
 0 & e_0 & 0 & e_1 & e_2 & 0 & 0 & 0 \\
 0 & 0 & e_0 & 0 & 0 & 0 & 0 & e_1 \\
 0 & 0 & 0 & 0 & 0 & e_2 & e_1 & 0
 \end{array} \right) ; \\
 e_1 & & & & & & & & \\
 e_2 & & & & & & & & \\
 v_1 & & & & & & & & \\
 v_2 & & & & & & & & \\
 v_3 & & & & & & & & \\
 \\
 & t_1 & t_2 & t_3 & t_4 & t_5 & t_6 & t_7 & t_8 \\
 r_1 & \left(\begin{array}{cccccccc}
 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\
 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & v_1 & 0 & 0 & v_3 & v_2 \\
 v_0 & 0 & 0 & 0 & v_1 & v_3 & 0 & 0 \\
 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 \\
 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\
 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0
 \end{array} \right) . \\
 r_2 & & & & & & & & \\
 r_3 & & & & & & & & \\
 e_1 & & & & & & & & \\
 e_2 & & & & & & & & \\
 v_1 & & & & & & & & \\
 v_2 & & & & & & & & \\
 v_3 & & & & & & & &
 \end{matrix}$$

The first is equal to $e_0^2 e_1 + e_0 e_1^2 + e_0^2 e_2 + e_0 e_1 e_2 + e_1^2 e_2 + e_0 e_2^2 + e_1 e_2^2$, and the second one is $v_0 v_1 + v_1^2 + v_0 v_2 + v_1 v_2 + v_0 v_3 + v_1 v_3 + v_2 v_3$. Observe that the sequences of exponents $((2, 1, 0)$ etc. and $(1, 1, 0, 0)$ etc., respectively) are exactly the hypertrees that we saw earlier.

Of course, as we have already done in the example above, we may use Theorem 10.5 to compute any of the six hypertree polytopes associated to a trinity. This leads to another proof of Theorem 8.3 in terms of manipulating determinants. It is quick but it is probably much less instructive than the proof given in Section 8.

Indeed, starting from $M_{e \rightarrow v}$, pull out e_j from any column that belongs to a white triangle adjacent to e_j . Do this for all emerald points e_j . Then for the non-root emerald points, pull out e_j^{-1} from the row indexed by e_j . The result is a monomial times the enhanced adjacency matrix $M_{e \rightarrow r}$ but with e_j^{-1} written everywhere where e_j should appear. Therefore the determinant of the matrix is $-Q_{\mathcal{H}^*} \cap \mathbf{Z}^E$.

If we examine the monomial factor, we see that the exponent of e_j in it is exactly $|e_j| - 1$. Here $|e_j|$ means the size of e_j as a hyperedge, which is the same in \mathcal{H} as in \mathcal{H}^* . It can also be described as the number of violet and red points connected to e_j by red, respectively violet edges of the trinity, as well as the number of white triangles adjacent to e_j . The claim is obvious for the non-root emerald points and the exponent of e_0 is $|e_0| - 1$ because e_0 is adjacent to the outer white triangle t_0 which had no column in the matrix.

10.3. *Summary and two final notes*

Let us summarize our findings on hypertree polytopes and interior and exterior polynomials associated to the six hypergraphs contained in a trinity. The six polytopes form three centrally symmetric pairs by [Theorem 8.3](#), i.e., there are only three ‘shapes’ associated to the trinity. These polytopes have different dimensions but by [Corollary 10.3](#) each contains the same number of integer lattice points, namely ρ , where ρ is the arborescence number of the trinity.

A priori, there are twelve interior and exterior polynomials associated to the six hypergraphs. But again by [Theorem 8.3](#), that number is reduced to six, as indicated in the table below. These polynomials all have non-negative integer coefficients with their sum equal to ρ . Furthermore, if [Conjecture 7.1](#) holds, then $I = I'$, $X = X'$, and $Y = Y'$, so there are in fact only three different polynomials.

Hypergraph	(V, E)	(R, E)	(E, R)	(V, R)	(R, V)	(E, V)
Interior p.	I	X	X'	Y	Y'	I'
Exterior p.	X	I	Y	X'	I'	Y'
Polytope	$Q_{(V,E)} \cong -Q_{(R,E)}$	$Q_{(E,R)} \cong -Q_{(V,R)}$	$Q_{(R,V)} \cong -Q_{(E,V)}$			

Note the cyclic nature of the table: alternately, each hypergraph is either the planar dual or the abstract dual of the one to its left, and the pattern wraps up (so that (V, E) naturally follows after (E, V)). Cf. [Proposition 9.3](#).

Our first ‘final note’ is that $Q_{(V,E)} \cap \mathbf{Z}^E$, $Q_{(E,R)} \cap \mathbf{Z}^R$, and $Q_{(R,V)} \cap \mathbf{Z}^V$ (that is, the three different arrangements of hypertrees that can be derived from the trinity) are all projections of a single set Q of lattice points. This can be defined by superimposing $M_{e \rightarrow v}$, $M_{r \rightarrow e}$, and $M_{v \rightarrow r}$ to obtain the matrix $M_{e \rightarrow v \rightarrow r}$ and taking its determinant. In our running example, the result is

$$M_{e \rightarrow v \rightarrow r} = \begin{matrix} & \begin{matrix} t_1 & t_2 & t_3 & t_4 & t_5 & t_6 & t_7 & t_8 \end{matrix} \\ \begin{matrix} r_1 \\ r_2 \\ r_3 \\ e_1 \\ e_2 \\ v_1 \\ v_2 \\ v_3 \end{matrix} & \begin{pmatrix} 0 & 0 & 0 & 0 & v_1 & 0 & v_3 & 0 \\ v_0 & v_1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & v_2 & v_1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & r_3 & 0 & 0 & r_1 & r_0 \\ r_2 & 0 & 0 & 0 & r_1 & r_0 & 0 & 0 \\ 0 & e_0 & 0 & e_1 & e_2 & 0 & 0 & 0 \\ 0 & 0 & e_0 & 0 & 0 & 0 & 0 & e_1 \\ 0 & 0 & 0 & 0 & 0 & e_2 & e_1 & 0 \end{pmatrix} \end{matrix},$$

with determinant $e_0^2 e_1 r_0^2 v_0 v_1^2 + e_0 e_1^2 r_0 r_3 v_0 v_1 v_2 + e_1^2 e_2 r_1 r_2 v_1^2 v_2 + e_0^2 e_2 r_0 r_1 v_0 v_1 v_3 + e_0 e_2^2 r_0 r_2 v_1^2 v_3 + e_0 e_1 e_2 r_1 r_3 v_0 v_2 v_3 + e_1 e_2^2 r_2 r_3 v_1 v_2 v_3$. It is obvious that $Q_{(V,E)} \cap \mathbf{Z}^E = \det M_{e \rightarrow v}$, $Q_{(E,R)} \cap \mathbf{Z}^R = \det M_{r \rightarrow e}$, and $Q_{(R,V)} \cap \mathbf{Z}^V = \det M_{v \rightarrow r}$ are all obtained from $\det M_{e \rightarrow v \rightarrow r}$ by substituting 1 for the unneeded variables. In particular, all four determinants contain the same number of lattice




Fig. 14. Collapsing moves used in the definition of Jaeger's polynomial.

points, which is the arborescence number of the trinity. It is not yet clear, however, whether $Q = \det M_{e \rightarrow v \rightarrow r}$ is such that $(\text{Conv } Q) \cap (\mathbf{Z}^E \oplus \mathbf{Z}^V \oplus \mathbf{Z}^R) = Q$.

Finally, there exists a curious relationship between our results and an invariant introduced by Jaeger [3]. We noted in Section 9.1 that the planar dual of a trinity is a plane bipartite cubic graph. Jaeger associated a one-variable polynomial $S(u)$ (with non-negative integer coefficients) to such objects. His definition does not make use of the three-coloring of Proposition 9.4 and indeed, his polynomial is different from any of the six (three) that we introduced. (For example, the trinity of Fig. 11 has $S(u) = 6u^3 + u^5$.) The sum of the coefficients is, however, the same.

Proposition 10.7. *Let T be a trinity. Then $S(T^*, 1) = \rho(T)$, the arborescence number of T .*

Proof. As the definition of S is based on a recursion, it is natural to prove our claim by induction. Jaeger sets his initial condition $S(u) = 1$ at a ‘free loop,’ but an equivalent notion results if we start by setting $S(u) = u$ for the theta graph (which has two vertices and three edges connecting them). The theta graph is dual to the trinity formed by one black and one white triangle with a common boundary. Each of the three directed graphs associated to the latter has a single vertex and one loop edge and, therefore, only one arborescence. So our claim holds for the theta graph.

The invariant S obeys two kinds of recursive ‘collapsing’ rules. Fig. 14 shows the associated pictures, drawn not for the plane bipartite cubic graphs but for their dual trinities (with an arbitrarily chosen local coloring). The part of the trinity outside of the emerald contour does not change under the collapsing operations. (The red and violet points in the pictures may have other neighbors that are not shown.) It is easy to show that if a trinity contains four or more triangles than at least one of the moves can always be carried out.

The rules themselves are $S(T^*, u) = uS(T_0^*, u)$ for the collapsing move on the left and $S(T^*, u) = S(T_1^*, u) + S(T_2^*, u)$ for the one on the right. In the first case, $\rho(T) = \rho(T_0)$ by an application of Corollary 10.3 and either Lemma 6.6 (if we choose to work with the red or the violet graph) or the observation that a double edge has been removed from the emerald graph.

In the second case, separate the points of T by color into the usual sets R , E , and V . Notice that the set V_1 of violet points in T_1 is obtained from V by identifying v' and v'' to a single point v . The set of emerald points in both T_1 and T_2 is $E \setminus \{e\}$. Furthermore, the hypergraph $(V_1, E \setminus \{e\})$ contained in T_1 and the hypergraph $(V, E \setminus \{e\})$ contained in T_2 are the contraction and deletion, respectively, of the hyperedge e from the hypergraph (V, E) contained in T . So we get the desired conclusion from Corollary 10.3 and Proposition 6.14. \square

Acknowledgments

Part of the research reported here was completed while I worked at the University of Tokyo. My understanding of issues pertaining to the current paper has greatly benefited from discussions

with my collaborators Alex Bene, András Juhász, and Jake Rasmussen. I also want to express my gratitude to László Fehér, László Lovász, Hitoshi Murakami, Richárd Rimányi, Oliver Riordan, Gábor Tardos, and Vera Vértesi for stimulating discussions. I thank Péter Juhász for checking many examples by computer. Last but not least, I am grateful to the referee for a number of helpful suggestions for improvement in the exposition.

References

- [1] K.A. Berman, A proof of Tutte's trinity theorem and a new determinant formula, *SIAM J. Algebraic Discrete Methods* 1 (1) (1980) 64–69.
- [2] J. Ellis-Monaghan, C. Merino, Graph polynomials and their applications I: the Tutte polynomial, in: Matthias Dehmer (Ed.), *Structural Analysis of Complex Networks*, Birkhauser, 2011, pp. 219–256.
- [3] F. Jaeger, A new invariant of plane bipartite cubic graphs, *Discrete Math.* 101 (1992) 149–164.
- [4] A. Postnikov, Permutohedra, Associahedra, and Beyond, *Int. Math. Res. Not.* (6) (2009) 1026–1106.
- [5] R. Schneider, Convex bodies: the Brunn–Minkowski theory, in: *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, 1993.
- [6] A. Schrijver, Combinatorial optimization. Polyhedra and efficiency, in: *Algorithms and Combinatorics* 24, Springer-Verlag, Berlin, 2003.
- [7] W.T. Tutte, The dissection of equilateral triangles into equilateral triangles, *Proc. Cambridge Philos. Soc.* 44 (1948) 463–482.
- [8] W.T. Tutte, A contribution to the theory of chromatic polynomials, *Canad. J. Math.* 6 (1954) 80–91.
- [9] W.T. Tutte, Duality and trinity. Infinite and finite sets, in: (Colloq. Keszthely, 1973; dedicated to P. Erdős on his 60th birthday), vol. III, in: *Colloq. Math. Soc. János Bolyai*, vol. 10, North-Holland, Amsterdam, 1975, pp. 1459–1472.