Submitted exclusively to the London Mathematical Society
doi:10.1112/0000/000000

Root polytopes, Tutte polynomials, and a duality theorem for bipartite graphs

Tamás Kálmán and Alexander Postnikov

Abstract

Let G be a connected bipartite graph with colour classes E and V and root polytope Q. Regarding the hypergraph $\mathcal{H} = (V, E)$ induced by G, we prove that the interior polynomial of $\mathcal{H}$ is equivalent to the Ehrhart polynomial of Q, which in turn is equivalent to the h-vector of any triangulation of Q. It follows that the interior polynomials of $\mathcal{H}$ and its transpose $\overline{\mathcal{H}} = (E, V)$ agree.

When G is a complete bipartite graph, our result recovers a well known hypergeometric identity due to Saalschütz. It also implies that certain extremal coefficients in the Homfly polynomial of a special alternating link can be read off of an associated Floer homology group.

1. *Introduction*

The interior polynomial I [**5**] is an invariant of hypergraphs (and of integer polymatroids) that generalises the specialisation $T(x, 1)$ of the Tutte polynomial of ordinary graphs (matroids). It first arose as a byproduct of the first author's study of polynomial invariants of knots. In that context, it was natural to conjecture that $I_{\mathcal{H}} = I_{\overline{\mathcal{H}}}$, where $\overline{\mathcal{H}}$ is the 'abstract dual' (also known as transpose) of the hypergraph $\mathcal{H}$ resulting from interchanging the roles of its vertices and hyperedges. That is, I is an invariant of the bipartite graph G that captures the common structure of $\mathcal{H}$ and $\overline{\mathcal{H}}$.

In this paper we verify that conjecture. As an intermediary between the two polynomials, we insert the Ehrhart polynomial ε_G of the root polytope Q_G, which was introduced and studied by the second author [**9**]. (For the definition of Q_G see Section 3.) We also show that ε_G has an equivalent description as the h-vector of an arbitrary triangulation of Q_G. (In particular, all triangulations of Q_G have the same h-vector h_G.) Explicitly, we will prove the following.

Theorem 1.1. *If E and V are the colour classes of the connected bipartite graph G, then the interior polynomial I of either hypergraph induced by G is related to h_G via the formula*

$$I(x) = x^{|E|+|V|-1} h_G(x^{-1}). \tag{1.1}$$

We chose the somewhat unusual labels for the colour classes to ease the transition to the hypergraph point of view. In the proof, we argue that if the sequence $a_0, a_1, \ldots, a_{|E|+|V|-2}$ of rational numbers (which will later turn out to consist of non-negative integers, with at least $\max\{\,|E|, |V|\,\} - 1$ zeros at the end) is such that for any non-negative integer s, we have

$$\varepsilon_G(s) := \left|(s \cdot Q_G) \cap (\mathbf{Z}^E \oplus \mathbf{Z}^V)\right| = \sum_{k=0}^{|E|+|V|-2} a_k \binom{s + |E| + |V| - 2 - k}{|E| + |V| - 2}, \tag{1.2}$$

2000 *Mathematics Subject Classification* 05C31 (primary), 05C65 (secondary).

The first author was supported by consecutive Japan Society for the Promotion of Science (JSPS) Grant-in-Aids for Young Scientists B (nos. 21740041 and 25800037). The second author was supported in part by NSF Grants DMS-1100147 and DMS-1362336.

then both sides[†] of (1.1) agree with $\sum_{k=0}^{|E|+|V|-2} a_k x^k$. This boils down to a delicate analysis of interior faces in a triangulation of Q_G in which we will rely on previous work by the second author [**9**]. We note that a key result of [**9**] is the claim $I_{\mathcal{H}}(1) = I_{\overline{\mathcal{H}}}(1)$, which is proven by equating both sides to $h_G(1)$. That is, in [**9**] it is shown that $\mathcal{H}$ and $\overline{\mathcal{H}}$ have the same number of so-called hypertrees (see Definition 2.2) and that that number is the number of simplices in each triangulation of Q_G.

If G is a complete bipartite graph on $(m+1)+(n+1)$ vertices then $Q_G = \Delta_m \times \Delta_n$ is the product of an m- and an n-dimensional unit simplex. If we consider (1.2) in this special case and for a_k we substitute the coefficients of the interior polynomial from a separate computation [**5**, Example 7.2], we obtain the well known identity

$$\binom{s+m}{m}\binom{s+n}{n} = \sum_{k=0}^{\min\{m,n\}} \binom{m}{k}\binom{n}{k}\binom{s+m+n-k}{m+n} \tag{1.3}$$

due to Saalschütz. See Example 5.3. The other form in which (1.3) often appears,

$$\binom{q}{m}\binom{q}{n} = \sum_{k=0}^{\min\{m,n\}} \binom{m}{k}\binom{n}{k}\binom{q+k}{m+n}, \tag{1.4}$$

is related to (1.3) by an application of Ehrhart reciprocity to Q_G.

Both our main theorem and its proof are entirely combinatorial. Yet the result serves as a crucial step in a knot theoretical program advanced by Juhász, Murakami, Rasmussen, and the first author. Namely, in [**6**] it is shown that certain extremal coefficients in the Homfly polynomial of a special alternating link agree with coefficients in the h-vector h_G, where G is the Seifert graph of the link. By the main result of this paper, we see that those same Homfly coefficients are also the coefficients of the interior polynomial of either one of the corresponding two hypergraphs. By definition, interior polynomials are derived from the structure of the set of hypertrees of the hypergraph (see [**5**] and Section 2). Finally, in [**4**] it is proven that the hypertrees in question appear as the supporting Spin^c structures of some Floer homology groups that are naturally associated to the special alternating link. Hence it is possible to read certain Homfly coefficients directly out of Floer homology. To the best of our knowledge such a result is the first of its kind. In order to keep this paper concise, we omit a full explanation and mention only one topological corollary of Theorem 1.1.

Corollary 1.2. Regarding a positive special alternating link diagram with Seifert graph G, the coefficient of $z^{b_1(G)}$ in its Homfly polynomial $P(v,z)$ is $v^{b_1(G)} I(v^2)$, where I is the common interior polynomial of the two hypergraphs induced by G. (Here the first Betti number $b_1(G)$ is the highest exponent of z in P [**8**].)

Observations by Murakami and the first author [**6**] concerning generalised parking functions [**10**] for the dual of a plane bipartite graph G yield that their natural enumerator also coincides with the interior polynomial of G. See Corollary 5.9.

Some of our results provide new information on the Tutte polynomial $T(x,y)$ in the classical (graph and matroid) context. As a special case of Theorem 2.10 we see that $T(1,y)$ and $T(x,1)$ are equivalent to lattice point counts in the Minkowski sum of a simplex or inverted simplex of variable sidelength with the spanning tree polytope (for graphs) or base polytope

[†]The dimension of Q_G is $|E|+|V|-2$. The binomial coefficient on the right hand side of (1.2) is the number of lattice points in a standard simplex of that dimension and of sidelength $s-k$. When considered for all k, these binomial coefficients form a basis in the sense of Lemma 3.8.

(for matroids). Our duality theorem (Corollary 5.4) can be used to express $T(x,1)$ in the case of a graph as a sum written in terms of activities associated to vertices instead of edges, cf. Corollary 5.7.

This paper updates its predecessor [**5**] by replacing the two longest proofs. Namely, Corollary 5.4 implies that the two deletion-contraction formulas established there [**5**, Proposition 6.14 and Theorem 7.3] are equivalent, so a separate proof of the latter is no longer necessary. We also give a new proof of the well-definedness of the interior (and exterior) polynomial, shorter and more elegant than the original [**5**, Theorem 5.4].

The paper is organised as follows. In Section 2 we recall the construction of the interior polynomial and re-prove that it is well-defined. In Section 3 we summarise some facts about the root polytope and establish the equivalence of its Ehrhart polynomial with the common h-vector of its triangulations. Section 4 contains a few lemmas needed for the proof of Theorem 1.1, which is given in Section 5. At the end of the paper we discuss several corollaries.

Acknowledgements. We are grateful to Péter Frenkel for recognising Saalschütz's identity, and to Suho Oh and Dylan Thurston for stimulating conversations.

2. *The interior polynomial*

In this section we set notation, recall some necessary material from [**5**], and give a new, shorter proof of one of that paper's key results.

2.1. *Hypertrees*

A *hypergraph* is a pair $\mathcal{H} = (V, E)$, where V is a finite set (*vertices*) and E is a finite multiset of non-empty subsets of V (*hyperedges*). The sets V and E are the colour classes of a bipartite graph $\operatorname{Bip}\mathcal{H}$, where we connect $v \in V$ to $e \in E$ with an edge if $v \in e$. We call $\operatorname{Bip}\mathcal{H}$ the *bipartite graph associated to the hypergraph* $\mathcal{H}$. We will often refer to E as the emerald colour class of $\operatorname{Bip}(V, E)$ and to V as the violet colour class.

The construction of $\operatorname{Bip}\mathcal{H}$ is reversible if we specify one of the colour classes, A and B say, in the bipartite graph G. Let the resulting hypergraphs be

$$\mathcal{H}_0 = (A, B) \quad \text{and} \quad \mathcal{H}_1 = (B, A). \tag{2.1}$$

Definition 2.1. The bipartite graph G above is said to *induce* the hypergraphs $\mathcal{H}_0$ and $\mathcal{H}_1$. Two hypergraphs are *abstract duals* if they are obtained in the form (2.1). I.e., the abstract dual (also known as the *transpose*) $\overline{\mathcal{H}} = (E, V)$ of a hypergraph $\mathcal{H} = (V, E)$ is defined by interchanging the roles of its vertices and hyperedges.

A certain generalisation of the notion of a spanning tree to hypergraphs plays a central role in this paper. It appears in both authors' previous work. Although it was first introduced as the 'left (or right) degree vector' [**9**], we will use the following terminology instead.

Definition 2.2. Let $\mathcal{H} = (V, E)$ be a hypergraph so that its associated bipartite graph $\operatorname{Bip}\mathcal{H}$ is connected. (In this case we call $\mathcal{H}$ itself *connected*.) A *hypertree* in $\mathcal{H}$ is a function (vector) $\mathbf{f}\colon E \to \mathbf{N} = \{0, 1, \dots\}$ so that a spanning tree of $\operatorname{Bip}\mathcal{H}$ can be found which has valence $\mathbf{f}(e) + 1$ at each $e \in E$. Such a spanning tree is said to *realise* or to *induce* $\mathbf{f}$. We denote the set of all hypertrees in $\mathcal{H}$ with $B_{\mathcal{H}}$.

It is easy to show that all hypertrees $\mathbf{f}$ in $\mathcal{H}$ satisfy $\sum_{e\in E}\mathbf{f}(e) = |V| - 1$. The set $B_{\mathcal{H}}$ is such that $(\operatorname{Conv} B_{\mathcal{H}}) \cap \mathbf{Z}^E = B_{\mathcal{H}}$, cf. Lemma 2.7, where Conv denotes the usual convex hull. We will call $\operatorname{Conv} B_{\mathcal{H}}$ the *hypertree polytope* of $\mathcal{H}$.

The definition of the interior polynomial uses hypertrees and the following concept, which is a natural extension of internal activity in graphs (and matroids). A small difference is that if we specialise our version to graphs, external edges of a spanning tree become internally active; another is that we count inactive hyperedges instead of active ones. But since the number of external edges is the same for all spanning trees (namely, the first Betti number of the graph, also known as its nullity), all this just mirrors and shifts the distribution of the 'classical' statistic.

DEFINITION 2.3. Let (V, E) be a connected hypergraph. A hyperedge $e \in E$ is *internally active* with respect to the hypertree $\mathbf{f}$ if one cannot decrease $\mathbf{f}(e)$ by 1 and increase $\mathbf{f}$ at a hyperedge smaller than e by 1 so that another hypertree results. Let $\iota(\mathbf{f})$ denote the number of internally active hyperedges with respect to $\mathbf{f}$.

If the hypertree $\mathbf{f} \in B_{\mathcal{H}}$ and the hyperedges $e, e' \in E$ are such that changing the value $\mathbf{f}(e)$ to $\mathbf{f}(e) - 1$ and the value $\mathbf{f}(e')$ to $\mathbf{f}(e') + 1$ results in another hypertree $\mathbf{f}'$, then we say that $\mathbf{f}$ and $\mathbf{f}'$ are related by a *transfer of valence* from e to e'.

We call a hyperedge *internally inactive* with respect to a hypertree if it is not internally active and denote the number of such hyperedges (for a given $\mathbf{f}$) by $\bar{\iota}(\mathbf{f}) = |E| - \iota(\mathbf{f})$. This value will be called the *internal inactivity* of $\mathbf{f}$.

DEFINITION 2.4. Let $\mathcal{H} = (V, E)$ be a connected hypergraph. For some fixed linear order on E we consider the generating function of internal inactivity, $I_{\mathcal{H}}(\xi) = \sum_{\mathbf{f}\in B_{\mathcal{H}}} \xi^{\bar{\iota}(\mathbf{f})}$, and call it the *interior polynomial* of $\mathcal{H}$.

If $\mathcal{H}$ is a graph with Tutte polynomial $T(x, y)$, then its interior polynomial is $\xi^{|V|-1}T(1/\xi, 1)$. For any $\mathcal{H}$, the polynomial $I_{\mathcal{H}}$ is independent of the order that is used to define $\bar{\iota}$ [**5**, Theorem 5.4]. We give a new proof of this fact in the next subsection. There is a similar notion of external activity and a corresponding exterior polynomial of hypergraphs [**5**]. However it does not play much of a role in this paper other than in the planar case, cf. [**5**, Theorem 8.3] and Corollary 5.8.

EXAMPLE 2.5. Let the complete bipartite graph $K_{2,3}$ induce the hypergraph $\mathcal{H}$ with three hyperedges and $\overline{\mathcal{H}}$ with two hyperedges. Both have three hypertrees, namely $(1, 0, 0), (0, 1, 0), (0, 0, 1)$ in $\mathcal{H}$ and $(2, 0), (1, 1), (0, 2)$ in $\overline{\mathcal{H}}$. In both cases, up to isomorphism there is only one way to order hyperedges and the resulting interior polynomial is $1 + 2\xi^2$.

EXAMPLE 2.6. Regarding the bipartite graph pictured in Figure 1, if we treat the four emerald points a, b, c, d as hyperedges and write hypertrees $\mathbf{f}$ in the form of a vector $(\mathbf{f}(a), \mathbf{f}(b), \mathbf{f}(c), \mathbf{f}(d))$, then $(1, 0, 1, 2)$ becomes a hypertree. This is because the spanning tree pictured on the right realises it.

For any order of the hyperedges, the smallest one is always internally active with respect to any hypertree. If a hypertree assigns 0 to a hyperedge, then that hyperedge is automatically internally active with respect to the hypertree. So if we use the order $a < b < c < d$ in our example, then with respect to $(1, 0, 1, 2)$ the hyperedges a and b are internally active. On the other hand, c and d are internally inactive. This is demonstrated by the upper right and lower

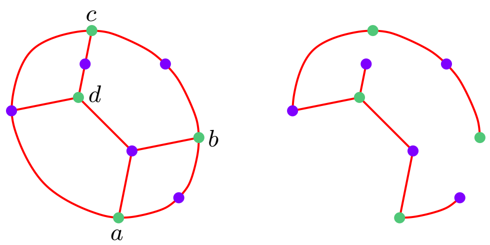

FIGURE 1. *A plane bipartite graph with one of its spanning trees.*

left spanning trees (and their induced hypertrees) of Figure 5. Indeed, one shows a transfer of valence from c to b and the other a transfer of valence from d to a. Hence $\bar{\iota}(1,0,1,2)=2$ so that the hypertree contributes 1 to the coefficient of ξ^2 in I.

The other 15 hypertrees may be found by solving the system of inequalities given in Lemma 2.7 below, whose necessity is particularly easy to see, or by any number of ad hoc methods†. (We do not, however, recommend enumerating all spanning trees as there are 217 of those.) After checking their internal inactivities, we find the interior polynomial to be $I(\xi)=1+4\xi+7\xi^2+4\xi^3$. The same polynomial results if the violet points play the role of hyperedges.

At times it will be convenient to rely on submodular function techniques, that is to say, on Lemma 2.8 below. This is made possible by the following observations.

For a bipartite graph G with colour classes E and V and for a subset $E'\subset E$, we let $G\big|_{E'}$ denote the graph formed by E', all edges of G adjacent to elements of E', and their endpoints in V. We let $c(E')$ denote the number of connected components of $G\big|_{E'}$ and we also let $\bigcup E'=V\cap\big(G\big|_{E'}\big)$ (this notation is natural if we view (V,E) as a hypergraph). Finally we let $\mu(\varnothing)=0$ and otherwise

$$\mu(E')=|\bigcup E'|-c(E').$$

Then μ is a non-decreasing (i.e., $E''\subset E'$ implies $\mu(E'')\le\mu(E')$) submodular function on the power set of E. The latter means that for all $A,B\subset E$ we have

$$\mu(A)+\mu(B)\ge\mu(A\cup B)+\mu(A\cap B).$$

This fact is probably 'part of the folklore.' For a proof, see [**5**, Proposition 4.7].

LEMMA 2.7 [**5**, Theorem 3.4]. *If G is connected, then $\mu(E)=|V|-1$ and the hypertrees* $\mathbf{f}$ *in (V,E) are exactly the integer solutions of the system of inequalities*

$$\mathbf{f}(e)\ge 0 \text{ for all } e\in E;\quad \sum_{e\in E}\mathbf{f}(e)=|V|-1;\quad \sum_{e\in E'}\mathbf{f}(e)\le\mu(E') \text{ for all } E'\subset E.$$

We note that the non-negativity of $\mathbf{f}$ follows from the other constraints. We say that the set $E'\subset E$ is *tight* at $\mathbf{f}$ if $\sum_{e\in E'}\mathbf{f}(e)=\mu(E')$ holds. If E' is tight at $\mathbf{f}$ and it so happens that for another hypertree $\mathbf{g}$, we have $\sum_{e\in E'}\mathbf{f}(e)=\sum_{e\in E'}\mathbf{g}(e)$ (for example, if $\mathbf{f}$ and $\mathbf{g}$ differ by a transfer of valence between elements of E'), then E' is also tight at $\mathbf{g}$. The next lemma follows immediately from [**12**, Theorem 44.2].

†In a planar case such as this, [**5**, Theorem 10.5] even provides an explicit determinant formula for the set of hypertrees.

LEMMA 2.8. *If the sets $A, B \subset E$ are both tight at the hypertree $\mathbf{f}$, then so are $A \cup B$ and $A \cap B$.*

Let now Γ be a spanning tree in G which induces the hypertree $\mathbf{f}\colon E \to \mathbf{N}$. For any subset $E' \subset E$, the connected components of $\Gamma\big|_{E'}$ induce a partition of E'.

LEMMA 2.9. *Let $E' \subset E$ be a tight set at the hypertree $\mathbf{f}$ and $\Gamma \subset G$ a realisation of $\mathbf{f}$. Then each part in the partition of E' induced by Γ is itself tight at $\mathbf{f}$.*

Proof. An equivalent definition of μ is that a cycle-free subgraph of $G\big|_{E'}$ is a spanning forest of $G\big|_{E'}$ if and only if it has $|E'| + \mu(E')$ edges. Since now E' is tight, we have $|E'| + \mu(E') = |E'| + \sum_{e\in E'} \mathbf{f}(e) = \sum_{e\in E'}(\mathbf{f}(e)+1)$, which is the number of edges in $\Gamma\big|_{E'}$ (note that $\mathbf{f}(e)+1$ is the degree of e in Γ and each edge of $\Gamma\big|_{E'}$ has a unique endpoint in E'). Thus $\Gamma\big|_{E'}$ is a spanning forest in $G\big|_{E'}$; in particular, the connected components of $\Gamma\big|_{E'}$ and those of $G\big|_{E'}$ induce the same partition of E' and for each part E'' of that partition, $\Gamma\big|_{E''}$ is a spanning tree in $G\big|_{E''}$. By the same logic as above the latter implies $\sum_{e\in E''}(\mathbf{f}(e)+1) = |E''| + \mu(E'')$, i.e., $\sum_{e\in E''} \mathbf{f}(e) = \mu(E'')$, as claimed. □

2.2. *Order-independence*

The fact that the interior polynomial is independent of the order imposed on the hyperedges was proved in [**5**] using a straightforward but rather long argument. The reasoning that establishes the main theorem of this paper can also be viewed as a rather circuitous proof of order-independence. Here we give a third argument which is shorter than the other two. It also serves to emphasise that interior (and exterior) polynomials correspond to Ehrhart-type lattice point counts not just in the sense of (1.2) but in a more direct way as well. (One might expect this observation to be the basis of a short proof of Theorem 1.1, but so far no such connection has been found.)

We fix the connected hypergraph $\mathcal{H} = (V, E)$ whose set of hypertrees is $B_{\mathcal{H}}$. Let the *inverted standard simplex* ∇_E be the convex hull of the set $\{\, -\mathbf{i}_{\{e\}} \in \mathbf{R}^E \mid e \in E \,\}$. Here and in the rest of the paper $\mathbf{i}_S$ denotes the indicator function of the set $S \subset E$. For a positive integer k we are going to count lattice points in the Minkowski sum

$$\operatorname{Conv} B_{\mathcal{H}} + k\nabla_E.$$

First of all, each of these is of the form $\mathbf{f} + \mathbf{v}$, where $\mathbf{f}$ is a hypertree and $\mathbf{v}$ is an integer vector in $k\nabla_E$, i.e., a vector with non-positive integer entries whose sum is $-k$. This follows from [**12**, Corollary 46.2c] since $\operatorname{Conv} B_{\mathcal{H}}$ is the base polytope of an integer polymatroid (by Lemma 2.7 and the submodularity of μ), and ∇_E is a translate of another such polytope by the integer vector $-\mathbf{i}_E$: indeed ∇_E corresponds to the $(|E|-1)$-uniform matroid on the set E.

After fixing an order on E, we partition our lattice points $\mathbf{f} + \mathbf{v}$ according to the lowest possible $\mathbf{f} \in B_{\mathcal{H}}$ in the following version of the lexicographic order:

$$\mathbf{f}_1 \text{ is smaller than } \mathbf{f}_2 \text{ if they are different and if for the smallest element } e \text{ of } E \text{ with } \mathbf{f}_1(e) \neq \mathbf{f}_2(e), \text{ we have } \mathbf{f}_1(e) > \mathbf{f}_2(e) \quad (2.2)$$

(this is rather natural because the sum of the entries in $\mathbf{f}$ is fixed). For a hypertree $\mathbf{f}$, let $P_{\mathbf{f}}$ be the corresponding set of the partition.

We claim that $P_{\mathbf{f}}$ is the set of lattice points in a(n inverted) simplex of sidelength k and dimension $\iota(\mathbf{f}) - 1 = |E| - 1 - \bar{\iota}(\mathbf{f})$. More precisely, if we let

$$\nabla_{\mathbf{f}} = \operatorname{Conv}\{\, -\mathbf{i}_{\{e\}} \mid e \in E \text{ is internally active with respect to } \mathbf{f} \,\},$$

then we have

$$P_{\mathbf{f}} = \mathbf{f} + (k\nabla_{\mathbf{f}} \cap \mathbf{Z}^E). \tag{2.3}$$

As this is the key idea of the proof we included Figure 2 to illustrate it. The solid black dots of the figure represent the hypertrees of some hypergraph $\mathcal{H} = (V, E)$ of three hyperedges e_0, e_1, e_2. They are included in the regular triangle with vertices $(|V| - 1)\mathbf{i}_{\{e_i\}}$, $i = 0, 1, 2$, although we changed those labels to e_i for simplicity (for the example we let $|V| = 17$ but that is not important). Instead of Conv $B_{\mathcal{H}} + k\nabla_E$, we constructed its integer translate Conv $B_{\mathcal{H}} + k(\nabla_E + \mathbf{i}_{\{e_0\}})$ (using $k = 5$) which is in the same (hyper)plane as $B_{\mathcal{H}}$. Figure 2 also shows the partition that corresponds to the order $e_0 < e_1 < e_2$. It consists of the lattice points in

- a single triangle shaded grey, one of whose vertices is the unique hypertree with $\bar{\iota} = 0$;
- one line segment for each hypertree with $\bar{\iota} = 1$, which are the hypertrees represented by the medium-sized black dots;
- a singleton for each hypertree with $\bar{\iota} = 2$.

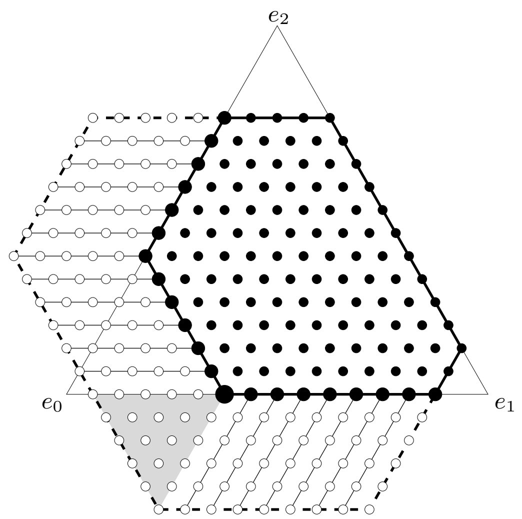

FIGURE 2. *Minkowski sum of a hypertree polytope and an inverted simplex.*

To prove (2.3), we establish the two-way inclusion as follows.

Proof that $\mathbf{f} + (k\nabla_{\mathbf{f}} \cap \mathbf{Z}^E) \subset P_{\mathbf{f}}$. Assume that $\mathbf{u} \in \mathbf{f} + (k\nabla_{\mathbf{f}} \cap \mathbf{Z}^E)$ is also an element of $\mathbf{g} + (k\nabla_E \cap \mathbf{Z}^E)$ for some hypertree $\mathbf{g}$. We claim that $\mathbf{f} \leq \mathbf{g}$ lexicographically. Suppose that $\mathbf{f} \neq \mathbf{g}$ and let e denote the smallest hyperedge where they differ. Let A be the set of hyperedges that are larger than e and internally active with respect to $\mathbf{f}$. Since $\mathbf{f}$ is such that no element of A can transfer valence to e, for all $a \in A$ there exists a set T_a of hyperedges that is tight at $\mathbf{f}$, contains e, and does not contain a. Put $T = \cap_{a \in A} T_a$ if $A \neq \varnothing$ and $T = E$ otherwise. Then the set T is tight at $\mathbf{f}$ (by Lemma 2.8) and contains e. Hence if we had $\mathbf{g}(e) > \mathbf{f}(e)$ then T would need to have an element e' so that $\mathbf{g}(e') < \mathbf{f}(e')$. Because of the way we chose e, this e' would obviously satisfy $e < e'$ and such elements of T are internally inactive with respect to $\mathbf{f}$. Now from $\mathbf{u} \in \mathbf{f} + (k\nabla_{\mathbf{f}} \cap \mathbf{Z}^E)$ it follows that $\mathbf{u}(e') = \mathbf{f}(e')$; on the other hand $\mathbf{u} \in \mathbf{g} + (k\nabla_E \cap \mathbf{Z}^E)$ implies $\mathbf{g}(e') \geq \mathbf{u}(e') = \mathbf{f}(e')$, which is a contradiction. Therefore $\mathbf{g}(e) < \mathbf{f}(e)$ and thus $\mathbf{f} < \mathbf{g}$ as claimed.

Proof that $P_{\mathbf{f}} \subset \mathbf{f} + (k\nabla_{\mathbf{f}} \cap \mathbf{Z}^E)$. Assume the contrary, namely that a vector $\mathbf{u} \in P_{\mathbf{f}}$ exists so that $\mathbf{u}(e) < \mathbf{f}(e)$ for a hyperedge e that is internally inactive with respect to $\mathbf{f}$. Let $\mathbf{g}$ be a hypertree that results from $\mathbf{f}$ by a single transfer of valence from e to a smaller hyperedge. Then $\mathbf{g} < \mathbf{f}$ lexicographically and $\mathbf{u} \in \mathbf{g} + (k\nabla_E \cap \mathbf{Z}^E)$ since each component of $\mathbf{u}$ is bounded above by the corresponding component of $\mathbf{g}$. This contradicts the definition of $P_{\mathbf{f}}$ and thus concludes the proof of (2.3).

From (2.3) it follows that if the interior polynomial, using the given order, of $\mathcal{H}$ is $I_{\mathcal{H}}(\xi) = a_0 + a_1\xi + a_2\xi^2 + \cdots$, then for all k, the polytope $(\operatorname{Conv} B_{\mathcal{H}}) + k\nabla_E$ contains

$$a_0\binom{k+|E|-1}{|E|-1} + a_1\binom{k+|E|-2}{|E|-2} + a_2\binom{k+|E|-3}{|E|-3} + \cdots$$

lattice points. Since the binomial coefficients $\{\binom{k+n}{n} \mid n = 0,1,2,\dots\}$ constitute a basis of the polynomial ring $\mathbf{Q}[k]$ (this is obvious from the fact that the degree of $\binom{k+n}{n}$ is n) and $(\operatorname{Conv} B_{\mathcal{H}}) + k\nabla_E$ does not depend on the order, the order-independence of $a_0, a_1, a_2, \dots$ follows.

There is a similar proof for the order-independence of the exterior polynomial [**5**] using the standard simplex $\Delta_E = -\nabla_E$ instead of ∇_E. Both arguments extend to the case of integer polymatroids without difficulty to yield the following result.

THEOREM 2.10. *Let S be a finite ground set and $\mu\colon \mathcal{P}(S) \to \mathbf{Z}$ a submodular and non-decreasing function on its power set. Let*

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

*be the polymatroid associated to μ and $B_\mu = \{\, \mathbf{x} \in P_\mu \mid \mathbf{x}\cdot\mathbf{i}_S = \mu(S)\,\}$ its base polytope. Then the interior and exterior polynomials I_μ and X_μ [**5**], in the standard basis $1, \mathrm{id}, \mathrm{id}^2, \dots, \mathrm{id}^{|S|-1}$, have the same coefficients as the polynomials (in k)*

$$\left|(B_\mu + k\nabla_S) \cap \mathbf{Z}^S\right| \quad \text{and} \quad \left|(B_\mu + k\Delta_S) \cap \mathbf{Z}^S\right|,$$

respectively, in the basis $\binom{k+|S|-1}{|S|-1}, \dots, \binom{k+2}{2}, k+1, 1$.

Theorem 2.10 offers new interpretations of the specialisations $T(x,1)$ and $T(1,y)$ of the Tutte polynomial of a matroid (graph) in terms of lattice point counts in Minkowski sums of the base polytope (spanning tree polytope) and a simplex. As the next step, it is natural to try to count lattice points in the base polytope plus k times the inverted simplex plus l times the standard simplex for independent variables k and l. Recently, Cameron and Fink [**2**] used this idea to find a new interpretation, and indeed a generalisation to polymatroids, of the two-variable polynomial $T(x,y)$.

3. *The root polytope*

In this section we recall the general notions of f-vector and h-vector. We discuss the root polytope, introduced in [**3**] and [**9**]. Some of the results also appeared in [**9**] and are included for the convenience of the reader, but some others, notably Proposition 3.6 and Theorem 3.10, are new.

3.1. *Triangulations*

A *triangulation* of a polytope Q is a collection of maximal simplices, each spanned by the vertices of Q, so that every two intersect in a common face and their union is Q. Here we consider the empty set as a face. To any d-dimensional simplicial complex we may associate the *f-vector*

$$f(y) = y^{d+1} + f_0\, y^d + f_1\, y^{d-1} + \cdots + f_{d-2}\, y^2 + f_{d-1}\, y + f_d, \tag{3.1}$$

where f_k, for $k \geq 0$, is the number of its k-dimensional simplices. The *h-vector* of the same complex is defined as $h(x) = f(x-1)$. The latter is well suited for *shellable complexes*, i.e., complexes with a shelling order. For a pure simplicial complex (that is, one in which all maximal simplices have the same dimension), a *shelling order* $\sigma_1 < \sigma_2 < \cdots < \sigma_{f_d}$ lists the maximal simplices so that each σ_i, $i \geq 1$, intersects the set $\sigma_1 \cup \cdots \cup \sigma_{i-1}$ in a union of c_i codimension one faces. We always have $c_1 = 0$ and assume that $c_i \geq 1$ for $i \geq 2$. When a shelling order exists, it is not hard to show [**13**] that

$$h(x) = f(x-1) = \sum_{i=1}^{f_d} x^{d+1-c_i}. \tag{3.2}$$

The root polytope may be defined for a large class of graphs [**9**, Section 13]. In this paper however we only need the following special case.

Definition 3.1. Let G be a bipartite graph with colour classes E and V. In the space $\mathbf{R}^E \oplus \mathbf{R}^V$ we consider the vectors $\mathbf{i}_{\{e\}} + \mathbf{i}_{\{v\}}$, where $e \in E$, $v \in V$, and ev is an edge in G. We let Q_G be their convex hull and call Q_G the *root polytope* of G.

The sum of the coordinates for each vertex of Q_G is 2. Furthermore, the sum of the E-coordinates equals that of the V-coordinates. When G is connected there are no more affine relations and the dimension of Q_G is $|E| + |V| - 2$.

In [**9**] the vertices of Q_G are given in the form $\mathbf{i}_{\{e\}} - \mathbf{i}_{\{v\}}$. (The two versions are isometric via multiplying each V-coordinate by -1.) Indeed, the root polytope is inspired by the standard proof that the graphical matroid of G is representable.

Example 3.2. The root polytope of the complete bipartite graph with colour classes E and V is the direct product of the unit simplices Δ_E and Δ_V. In Figure 3 we picture it in the case of $|E| = 3$ and $|V| = 2$, mainly to set notation for Example 3.13. Note that we replaced the symbols $\mathbf{i}_{\{x\}}$ with $\mathbf{x}$ for better readability.

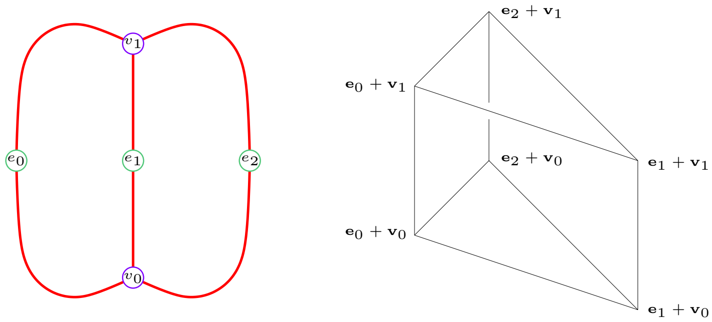

Figure 3. *The complete bipartite graph $K_{2,3}$ and its root polytope.*

The vertices of Q_G obviously correspond to edges of G. It is not hard to verify [**9**, Lemma 12.5] that a set of vertices is affinely independent if and only if the corresponding set of edges is cycle-free. In other words, simplices in Q_G (spanned by vertices) correspond to forests in G. In particular, maximal simplices in Q_G are in a one-to-one correspondence with spanning trees in G. Next we establish an elementary property of these simplices.

LEMMA 3.3. *If the vertices of the simplex σ are the vertices $\mathbf{v}_1, \ldots, \mathbf{v}_m$ of Q_G, then for each positive integer s, the set of integer points in $s \cdot \sigma$ agrees with the set*

$$A = \{ \lambda_1 \mathbf{v}_1 + \cdots + \lambda_m \mathbf{v}_m \mid \lambda_1, \ldots, \lambda_m \in \mathbf{N}, \lambda_1 + \cdots + \lambda_m = s \}.$$

Proof. Each point of $s \cdot \sigma$ can be uniquely written as $\sum_{i=1}^m \lambda_i \mathbf{v}_i$, where the λ_i are non-negative reals summing to s. Since the $\mathbf{v}_i$ are integer vectors, the relation $A \subset (s \cdot \sigma) \cap (\mathbf{Z}^E \oplus \mathbf{Z}^V)$ is clear. Conversely, let $\mathbf{p}$ be an integer point in $s \cdot \sigma$. We need to check that the unique solution (in $\lambda_1, \ldots, \lambda_m$) of $\sum_{i=1}^m \lambda_i \mathbf{v}_i = \mathbf{p}$ is integer.

In terms of the forest Σ that corresponds to σ, the unknowns λ_i are associated to the edges of Σ and the coordinates of $\mathbf{p}$ are associated to the vertices. The condition is that the value at each vertex be the sum of the values on the adjacent edges. Now this system of equations is in 'triangular form': Σ has at least one degree 1 vertex and the corresponding equation forces the unknown on the adjacent edge to take an integer value. Then we can remove this vertex-edge pair and look for another degree 1 vertex. This way a simple inductive argument can be constructed which shows that if the system of equations has a solution (which we assumed) then that solution needs to be integer. □

As to the relative position of two maximal simplices in Q_G, we recall the following basic observation.

LEMMA 3.4 [**9**, Lemma 12.6]. *Let Γ_1 and Γ_2 be spanning trees in G. The following two statements are equivalent.*

(i) *The simplices in Q_G that correspond to the Γ_i intersect in a (possibly empty) common face.*
(ii) *There does not exist a cycle $\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_{2k}$ of edges in G, where $k \geq 2$, so that all odd-index edges are from Γ_1 and all even-index edges are from Γ_2.*

If two maximal simplices (spanning trees) satisfy the equivalent conditions of Lemma 3.4 then we call them *compatible*. A triangulation of Q_G is then a collection of pairwise compatible maximal simplices whose union is Q_G. Since all maximal simplices have the same volume [**9**, Lemma 12.5], each triangulation of Q_G consists of the same number of them. Theorem 3.10 below shows that the triangulations have much more in common.

3.2. *Facets*

Our next goal is to describe the root polytope by a system of linear inequalities. In other words, we shall discuss the facets (codimension one strata of the boundary) of Q_G. These turn out to be in a one-to-one correspondence with certain cuts in G, as follows. A *cut* in a graph is a set of edges that is obtained by splitting the set of vertices into the disjoint union of two subsets, and then taking the set of edges between the two subsets. A non-empty cut is *minimal* if it does not contain any other non-empty cuts. For example, all *star-cuts* (when the vertex set is split into a non-isolated singleton and the rest) in a two-connected graph are minimal.

In a bipartite graph G with colour classes E and V we may speak of *directed cuts*, that is cuts which arise from a splitting $V \cup E = S \cup T$ so that each edge in the cut is adjacent to S at an element of V (and then to T at an element of E). In other words, G has no edges between $E \cap S$ and $V \cap T$. For instance, star-cuts in a bipartite graph are directed.

Lemma 3.5. *Let the splitting $V \cup E = S \cup T$ induce the non-empty directed cut C in the bipartite graph G. Then the root polytope Q_G has a supporting hyperplane which does not contain any of the vertices corresponding to elements of C but contains all other vertices.*

Proof. Assume that all the edges of C are adjacent to S at elements of V. The linear functional $\lambda_C \colon \mathbf{R}^E \oplus \mathbf{R}^V \to \mathbf{R}$ which is defined to take the value 1 at $\mathbf{i}_{\{x\}}$ if $x \in (S \cap V) \cup (T \cap E)$ and the value -1 if $x \in (S \cap E) \cup (T \cap V)$ is such that its value at vertices of Q_G corresponding to elements of C is 2, while at all other vertices λ_C vanishes. Thus the hyperplane $\Pi_C = \ker \lambda_C$ has the required properties. □

Proposition 3.6. *Let G be a connected bipartite graph. For each minimal directed cut C of G, the supporting hyperplane Π_C of Lemma 3.5 intersects the root polytope Q_G in a facet. Furthermore, each facet of Q_G arises this way.*

Proof. Let the minimal directed cut C belong to the splitting $V \cup E = S \cup T$. The subgraphs of G induced by S and T, respectively, are connected for otherwise C would not be minimal. By picking spanning trees in each, we obtain a two-component forest in G whose edge set is disjoint from C. The vertices of Q_G that correspond to edges of the forest are affinely independent, i.e., they form an $(|E| + |V| - 3)$-dimensional simplex. The fact that this simplex lies in Π_C proves the first claim.

Conversely, every facet of Q_G contains some $|E| + |V| - 2$ affinely independent vertices which span a codimension 1 (rel. Q_G) simplex σ. The cycle-free subgraph of G to which these correspond is a two-component forest F. The vertex sets of the two components define a minimal cut C in G. The cut C is directed because if it was not then F would have two extensions to spanning trees of G that satisfy the compatibility condition of Lemma 3.4 — but that would mean that the corresponding maximal simplices lie on opposite sides of σ, which is clearly impossible. Finally, as Π_C contains σ, it has to intersect Q_G in the given facet. □

Example 3.7. The root polytope of Figure 3 has five facets which correspond (in the manner described in Lemma 3.5) to the star-cuts at each of the five vertices. In this case the graph has no other minimal directed cuts.

3.3. *Ehrhart polynomials*

For $s \in \mathbf{N}$ let us define $\varepsilon_G(s) = |(s \cdot Q_G) \cap (\mathbf{Z}^E \oplus \mathbf{Z}^V)|$. This is well known to be a polynomial in s (which can then be extended to all $s \in \mathbf{C}$), called the *Ehrhart polynomial* of Q_G. Let us write ε_G as in (1.2),

$$\varepsilon_G(s) = \sum_{k=0}^{|E|+|V|-2} a_k C_k(s), \tag{3.3}$$

using the binomial coefficients $C_k(s) = \binom{s+|E|+|V|-2-k}{|E|+|V|-2}$ and rational numbers a_k. There is a unique such expression because the degree of ε_G is the dimension of Q_G and the following observation.

Lemma 3.8. *For each $d \in \mathbf{N}$, the expressions $C_0(s), \dots, C_d(s)$ form a basis over $\mathbf{Q}$ in the space $P_d[s]$ of polynomials with rational coefficients and of degree no greater than d.*

Why we prefer this basis will be made clear by Theorem 3.10. Then in Remark 3.12 we explain how one naturally arrives at it.

Proof. The degree of each $C_k(s)$ is exactly $|d$. As their number equals the dimension, it suffices to show that the $C_k(s)$ span $P_d[s]$. Since an element of $P_d[s]$ is determined by its values at $0, 1, 2, \ldots, d$, it suffices to arbitrarily fix rational numbers $r_0, r_1, \ldots, r_d$ and to find a linear combination of the $C_k(s)$ that takes the value r_j at j for each $j = 0, 1, \ldots, d$.

The roots of $C_k(s)$ are the consecutive integers $k-d, \ldots, k-1$. In particular, $C_0(s)$ is the only polynomial among the $C_k(s)$ that does not vanish at 0. This determines the coefficient of $C_0(s)$ in the desired linear combination. Since only $C_0(s)$ and $C_1(s)$ take non-zero values at 1, the coefficient of $C_1(s)$ also gets determined. By a trivial induction proof, the same is true for all coefficients and the resulting linear combination obviously satisfies our requirement. □

Remark 3.9. The set A of Lemma 3.3 has cardinality $\binom{s+m-1}{m-1}$. Thus the simplices σ that appear in Lemma 3.3 have Ehrhart polynomials $\varepsilon_\sigma(s) = \binom{s+\dim\sigma}{\dim\sigma}$. It follows that the number of lattice points in s times the *relative interior* of σ is $\binom{s-1}{\dim\sigma}$. This is the same observation that the standard proof of Ehrhart reciprocity is based on and in any case it follows from that principle since $\binom{s-1}{\dim\sigma} = (-1)^{\dim\sigma}\binom{-s+\dim\sigma}{\dim\sigma}$.

Theorem 3.10. *If ε_G is the Ehrhart polynomial of Q_G (as described in (3.3)), then the h-vector h of any triangulation of Q_G satisfies*

$$x^{|E|+|V|-1}h(x^{-1}) = a_0 + a_1\,x + a_2\,x^2 + \cdots + a_{|E|+|V|-2}\,x^{|E|+|V|-2}. \tag{3.4}$$

In particular, all triangulations of Q_G share the same h-vector.

We will denote the common h-vector described in Theorem 3.10 by h_G.

It is obvious from (3.1) that the left hand side of (3.4) is a polynomial. Since the Euler characteristic of a convex polytope is 1, we have

$$h(0) = f(-1) = (-1)^{d+1} + \sum_{l=0}^{d}(-1)^{d-l}f_l = (-1)^{d+1}\left(1 - \sum_{l=0}^{d}(-1)^l f_l\right) = 0,$$

implying that both sides of (3.4) are in $P_{|E|+|V|-2}$. Their degree is in fact much lower than $|E|+|V|-2$, but we will only see that after proving Theorem 1.1: indeed by [**5**, Proposition 6.1] the degree of the interior polynomial is at most $\min\{\,|E|, |V|\,\}-1$. This does not contradict the requirement that the degree of the Ehrhart polynomial ε_G be $\dim Q_G = |E|+|V|-2$ because each $C_k(s)$ in (3.3) has that degree.

Proof. Let us put $|E|+|V|-2 = d$. Fix a triangulation $\mathcal{T}$ of Q_G with h-vector h and corresponding f-vector $f(y) = h(y+1)$, cf. (3.1). That means that

$$\begin{aligned} h(x) = f(x-1) &= (x-1)^{d+1} + \sum_{l=0}^{d} f_l(x-1)^{d-l} \\ &= \sum_{i=1}^{d+1}\left((-1)^{d+1-i}\binom{d+1}{i} + \sum_{l=0}^{d-i} f_l(-1)^{d-l-i}\binom{d-l}{i}\right)x^i, \end{aligned}$$

where we used the fact that $h(0) = 0$. Thus the left hand side of (3.4) (putting $i = d + 1 - k$ above) is

$$\sum_{k=0}^{d} \left((-1)^k \binom{d+1}{d+1-k} + \sum_{l=0}^{k-1} f_l (-1)^{k-l-1} \binom{d-l}{d+1-k} \right) x^k. \tag{3.5}$$

In order to address the right hand side of (3.4), first we express $\binom{s-1}{l}$ (for $l = 0, 1, \ldots, d$) in the basis $C_k(s) = \binom{s+d-k}{d}$ ($k = 0, 1, \ldots, d$) of Lemma 3.8. (Why we need this will be made clear by (3.6) below.) To do so we will rely on the identity $\binom{n}{m} = \sum_{i=0}^{p} (-1)^i \binom{p}{i} \binom{n+p-i}{m+p}$. (This holds for any integers n, m and non-negative integer p. It is very easy to prove by induction on p based on just the basic relation in Pascal's triangle.) Applying it directly to $\binom{s-1}{l}$ with $p = d - l$ gives

$$\begin{aligned} \binom{s-1}{l} &= \sum_{i=0}^{d-l} (-1)^i \binom{d-l}{i} \binom{s-1+d-l-i}{l+d-l} \\ &= \sum_{i=0}^{d-l-1} (-1)^i \binom{d-l}{i} \binom{s-1+d-l-i}{d} + (-1)^{d-l} \binom{s-1}{d}. \end{aligned}$$

Here the last term is not one of our basis elements. To express it in the basis we apply our identity again, this time to $\binom{s-1}{-1} = 0$ with $p = d + 1$:

$$\begin{aligned} \binom{s-1}{l} = &\sum_{i=0}^{d-l-1} (-1)^i \binom{d-l}{i} \binom{s+d-(l+i+1)}{d} \\ &+ (-1)^{d-l}(-1)^{d+1} \left(-\sum_{j=0}^{d} (-1)^j \binom{d+1}{j} \binom{s+d-j}{d} \right) \\ = &\sum_{k=0}^{d} \left((-1)^{k-l-1} \binom{d-l}{k-l-1} + (-1)^{l+k} \binom{d+1}{k} \right) \binom{s+d-k}{d}. \end{aligned}$$

If $k \le l$ then of course $\binom{d-l}{k-l-1} = 0$.

Now as the relative interiors of the simplices of $\mathcal{T}$ partition Q_G, by Remark 3.9 we obtain

$$\begin{aligned} \varepsilon_G(s) &= \sum_{\text{simplices } \sigma \text{ in } \mathcal{T}} \binom{s-1}{\dim \sigma} = \sum_{l=0}^{d} f_l \binom{s-1}{l} \\ &= \sum_{k=0}^{d} \sum_{l=0}^{d} \left((-1)^{k-l-1} \binom{d-l}{k-l-1} + (-1)^{l+k} \binom{d+1}{k} \right) \cdot f_l \cdot \binom{s+d-k}{d}, \end{aligned} \tag{3.6}$$

that is that in (3.3) (and the right hand side of (3.4)) we have

$$\begin{aligned} a_k &= \sum_{l=0}^{d} \left((-1)^{k-l-1} \binom{d-l}{k-l-1} + (-1)^{l+k} \binom{d+1}{k} \right) \cdot f_l \\ &= (-1)^k \binom{d+1}{k} \sum_{l=0}^{d} (-1)^l f_l + \sum_{l=0}^{k-1} (-1)^{k-l-1} \binom{d-l}{k-l-1} f_l \\ &= (-1)^k \binom{d+1}{k} + \sum_{l=0}^{k-1} (-1)^{k-l-1} \binom{d-l}{k-l-1} f_l. \end{aligned} \tag{3.7}$$

Comparing the last formula to (3.5) completes the proof. □

Remark 3.11. From (3.7) we see that $a_0 = 1$ and $a_1 = -(d+1) + f_0$. Here f_0 is the number of edges in G and $d + 1 = |E| + |V| - 1$ is the number of edges in a spanning tree of G, i.e., we have $a_1 = b_1(G)$. These values are consistent with [**5**, Proposition 6.2 and Theorem 6.3] and our expectation (cf. Theorem 1.1) that they be the two lowest-degree coefficients in the common interior polynomial of the hypergraphs (V, E) and (E, V). As to a_2, see Proposition 5.5.

Remark 3.12. The proof of Theorem 3.10 becomes far more transparent if we assume $\mathcal{T}$ to be shellable†. This also makes the appearance of $C_k(s)$ more natural. Indeed if there is a shelling order with the quantities c_i as in (3.2), then the left hand side of (3.4) is $\sum x^{c_i}$ where the summation is over all maximal simplices. We may also count lattice points in $s \cdot Q_G$ using the shelling order, going over the maximal simplices one-by-one. If we do so then the i'th simplex contributes lattice points as described in Lemma 3.3, except that the points along c_i of its facets have already been counted. Geometrically that means that what is left to count are the lattice points in a simplex of full dimension $|E| + |V| - 2$ but with reduced sidelength $s - c_i$. (By sidelength we mean the number of lattice points along a side of the simplex minus one. In the absence of shellability one is forced to work with lower-dimensional simplices too, and that causes complications.) The number of those points is exactly $C_{c_i}(s)$. Therefore a_k, that is the number of times $C_k(s)$ appears in $\varepsilon_G(s)$, is equal to the number of maximal simplices with $c_i = k$, and that in turn is exactly the coefficient of x^k on the left hand side.

3.4. *Cross-sections*

A remarkable property of triangulations of the root polytope is that they 'pair up' hypertrees in the hypergraph $\mathcal{H} = (V, E)$ and its abstract dual (transpose) $\overline{\mathcal{H}} = (E, V)$. To set up the correspondence, we recall the following facts from [**9**]. First, Q_G naturally projects onto the standard unit simplices $\Delta_E \subset \mathbf{R}^E$ and $\Delta_V \subset \mathbf{R}^V$. The preimage of the barycentre $\mathbf{i}_V/|V| \in \Delta_V$ under the second projection $\mathrm{pr}_2 \colon Q_G \to \Delta_V$ is the $(|E|-1)$-dimensional Minkowski sum

$$S_E = \frac{1}{|V|}\left(\sum_{v\in V} \Delta_v\right) + \frac{1}{|V|} \cdot \mathbf{i}_V, \tag{3.8}$$

where $\Delta_v \subset \Delta_E \subset \mathbf{R}^E$ is the convex hull of the set $\{\, \mathbf{i}_{\{e\}} \mid ev \text{ is an edge in } G\,\}$. In particular, for a maximal simplex γ in Q_G that corresponds to the spanning tree Γ in G, we have

$$\gamma \cap S_E = \frac{1}{|V|}\left(\sum_{v\in V} \Delta_v^\Gamma\right) + \frac{1}{|V|} \cdot \mathbf{i}_V, \tag{3.9}$$

where Δ_v^Γ is the face of Δ_v spanned by unit vectors corresponding to neighbours of v in Γ. Moreover, a quick dimension count shows that the Minkowski sum in the latter formula is a direct sum, i.e., its elements have unique representations as sums of one vector from each summand. Below we will sometimes replace S_E with its homothetic image‡ $P_E = \sum_{v\in V} \Delta_v$ so that our formulas look simpler, however we will continue to think of P_E as a cross-section of Q_G. We will refer to the set described in (3.9), as well as to the homothetic set

$$M_\Gamma = \sum_{v\in V} \Delta_v^\Gamma \subset P_E, \tag{3.10}$$

as the *Minkowski cell* associated to the spanning tree Γ.

†It seems unlikely for all triangulations of Q_G to be shellable, but at present we do not know any counterexamples.

‡Both S_E and P_E belong, as rather special cases, to the class of *generalised permutohedra* [**9**].

Next, we recall that M_Γ contains a unique translate of the unit simplex Δ_E and it is $\mathbf{f} + \Delta_E$ [**9**] where $\mathbf{f} \in \mathbf{R}^E$ is the hypertree in $\mathcal{H}$ induced by Γ. (The interior of M_Γ is disjoint from other integer translates of Δ_E.) Let us denote the barycentre of this translated simplex with

$$\mathbf{f}^+ = \mathbf{f} + \frac{1}{|E|} \cdot \mathbf{i}_E. \tag{3.11}$$

We will call both $\mathbf{f}^+$ and (depending on context) the corresponding point

$$\frac{1}{|V|} \cdot \mathbf{f}^+ + \frac{1}{|V|} \cdot \mathbf{i}_V \in \gamma \cap S_E \subset S_E \subset Q_G, \tag{3.12}$$

the *emerald marker* of the simplex γ. Note that the set of emerald markers is a simple dilation of the set $B_\mathcal{H}$ of hypertrees. If markers are meant in the sense of (3.11), then in fact it is a translation of $B_\mathcal{H}$, so that for instance vectors connecting hypertrees are the same as vectors connecting the corresponding markers. Of course each maximal simplex in Q_G also contains a unique violet marker, i.e., an element of the set $\frac{1}{|E|}\left(B_{\overline{\mathcal{H}}} + \frac{1}{|V|}\mathbf{i}_V\right) + \frac{1}{|E|} \cdot \mathbf{i}_E$.

Now if we fix a triangulation of Q_G, then each marker will lie in the interior of a unique simplex; as this simplex has unique markers of each colour (which are essentially the two hypertrees that are induced by the spanning tree corresponding to the simplex), we see that $B_\mathcal{H}$ and $B_{\overline{\mathcal{H}}}$ are equinumerous. In Section 5, we exploit this picture further to improve the conclusion to $I_\mathcal{H} = I_{\overline{\mathcal{H}}}$.

Example 3.13. We examine the root polytope of the complete bipartite graph $G = K_{2,3}$. Let the emerald colour class be $E = \{ e_0, e_1, e_2 \}$ along with violet colour class $V = \{ v_0, v_1 \}$. In this case Q_G is the product of the 2-simplex Δ_E and the 1-simplex Δ_V. In Example 3.2 we labelled its vertices (using the symbol $\mathbf{x}$ in place of $\mathbf{i}_{\{x\}}$). In the upper left panel of Figure 4 we indicated the cross-sections S_E and S_V, as well as all markers for both colours (three each). We also chose a triangulation of the root polytope and showed the Minkowski cells that occur in the cross-sections. These are just three isometric segments in the one-dimensional cross-section S_V, whereas in S_E we get two triangles and a rhombus. By examining the spanning trees corresponding to the three maximal simplices (shown on the right), the interested reader may check the validity of Lemma 3.4 and the formula (3.10).

The f- and h-vectors of our triangulation are $f(y) = y^4 + 6y^3 + 12y^2 + 10y + 3$ and $h(x) = f(x-1) = x^4 + 2x^3$, respectively. The coefficients of the latter are explained by the fact that the triangulation is shellable so that the second and third maximal simplices are both attached along a single facet, cf. (3.2). Comparing the h-vector to the interior polynomial of Example 2.5, we see that Theorem 1.1 holds in this case.

We will need one more piece of information about the spanning tree $\Gamma \subset G$, corresponding maximal simplex $\gamma \subset Q_G$, and induced hypertree $\mathbf{f} \in B_\mathcal{H}$. In (3.11) we gave the coordinates of the emerald marker $\mathbf{f}^+$ (as it appears in the enlarged cross-section P_E) of γ in the natural basis of $\mathbf{R}^E$. However it will also be useful for us to describe the marker in terms of the direct sum decomposition (3.10) of the Minkowski cell M_Γ. It is convenient to identify the vertices of Δ_v^Γ with the edges of Γ adjacent to v. Then the barycentric coordinates of a point in Δ_v^Γ can be thought of as non-negative real numbers (weights) written on these edges, subject to the condition that their sum is 1. (Since we have $|V|$ such conditions on the $|E| + |V| - 1$ weights, this correctly identifies the dimension of M_Γ as $|E| - 1$.) By taking the sum of the weights assigned to edges adjacent to a vertex $e \in E$, we recover the standard e-coordinate in $\mathbf{R}^E$.

Lemma 3.14. *Let Γ be a spanning tree of G that induces the hypertree $\mathbf{f}$. The weights (on the edges of Γ) that produce the emerald marker $\mathbf{f}^+ \in M_\Gamma$ are given as follows. For each*

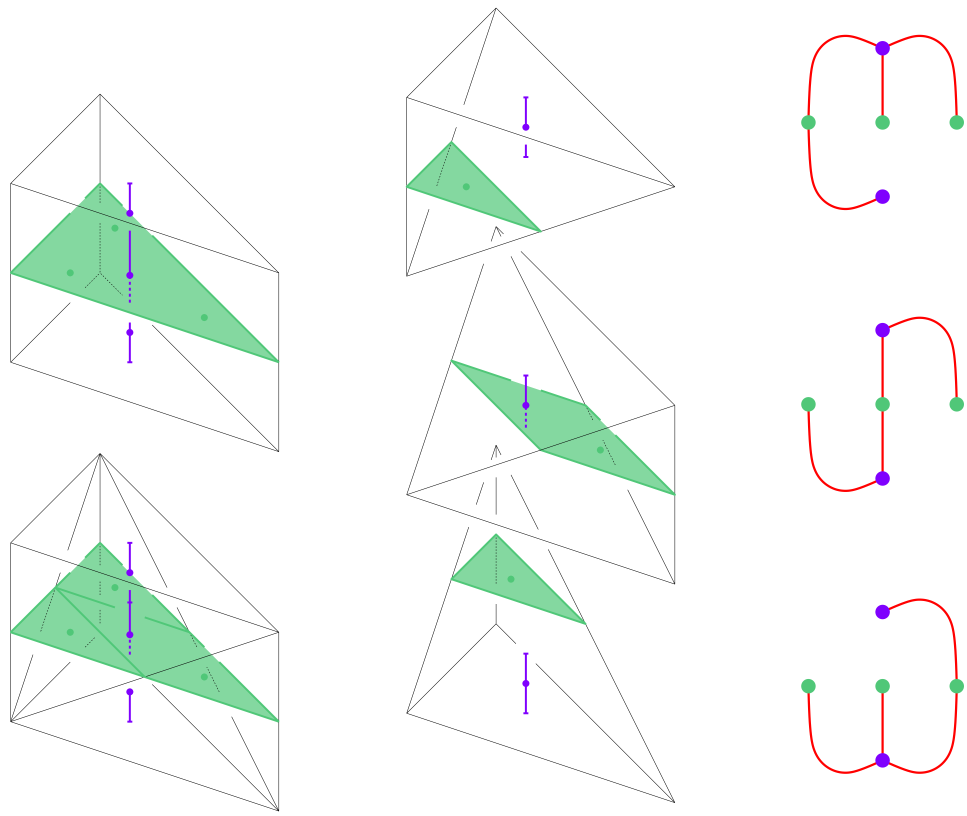

FIGURE 4. *The root polytope of $K_{2,3}$ with cross-sections and markers (upper left), and a triangulation (lower left and middle) with the corresponding spanning trees (right). Cf. Figure 3.*

$v \in V$, the elements of E are partitioned according to the edge of G adjacent to v through which they can be reached from v by a path in Γ. Let the weight of the edge be the size of the corresponding set divided by $|E|$.

Proof. Our assignment obviously satisfies all $|V|$ constraints and furthermore, if the vertex $e \in E$ has degree $d = \mathbf{f}(e) + 1$ in Γ, then the sum of the weights on the edges adjacent to e is

$$\frac{1}{|E|}\left((|E|-1)(d-1)+d\right) = d-1+\frac{1}{|E|} = \mathbf{f}(e)+\frac{1}{|E|} = \mathbf{f}^{+}(e),$$

as claimed. The formula is valid because when we compute the d weights, e appears in all d relevant counts of emerald vertices, whereas all other elements of E appear $d-1$ times. □

4. *Preparatory results*

In this section we prove a few more technical results needed for our main theorem. Regarding a polytope, we say that a simplex spanned by some of its vertices is *interior* if it is not part of the boundary of the polytope, i.e., if the (relative) interior points of the simplex are interior points of the polytope. E.g., all maximal simplices are interior but 0-dimensional simplices (vertices) are never interior.

Let Γ be a tree in our usual bipartite graph G of colour classes E and V, and let $C \subset \Gamma$ be a connected subgraph. If ε is an edge in $\Gamma \setminus C$, then there is a clear sense in which one endpoint of ε is closer to C than the other. Depending on the colour class of this endpoint, we say that the emerald or the violet endpoint of ε *faces* C.

LEMMA 4.1. *Let G be a connected bipartite graph with root polytope Q_G and fix an arbitrary triangulation $\mathcal{T}$ of Q_G. Let Σ be a cycle-free subgraph of G so that the corresponding simplex $\sigma \subset Q_G$ is an interior face of $\mathcal{T}$, and let C be a connected component of Σ. Then there is a unique maximal simplex in $\mathcal{T}$ containing σ so that in the corresponding spanning tree $\Gamma_C \supset \Sigma$, all edges of $\Gamma_C \setminus \Sigma$ have their violet endpoint face C.*

Because of the resemblance between Γ_C and an arborescence in a directed graph, we may refer to this statement as the 'arborescence lemma.'

Proof. We start with uniqueness. (This part of the proof works for all faces σ of $\mathcal{T}$, including those along the boundary of Q_G.) Let us assume that Γ_1 and Γ_2 are different spanning trees in G so that both satisfy the conditions stipulated for Γ_C. We will show that they violate the compatibility condition of Lemma 3.4 and hence the corresponding simplices cannot both be part of $\mathcal{T}$. Let us pick an edge $\varepsilon \in \Gamma_1 \setminus \Gamma_2$ and let e be the emerald endpoint of ε. Let us also fix an arbitrary vertex x of C and consider the unique paths $p_1 \subset \Gamma_1$ and $p_2 \subset \Gamma_2$ from e to x. The first edge along p_1 is clearly ε. Now let us find the first vertex y (after e) along p_1 that is also a vertex along p_2 (since x is such a vertex, y is well-defined). If we follow p_1 from e to y and then follow p_2 from y back to e, we obtain a cycle as described in condition ((ii)) of Lemma 3.4. Indeed, each time we step from an emerald to a violet vertex, that edge is either in Σ, hence in Γ_1, or if not then it has to be along the first half of our loop (from e to y) which is again part of Γ_1. Similarly, steps taken from violet to emerald vertices are either in Σ, hence in Γ_2, or along the second half of the loop which is part of Γ_2 as well.

We will establish the existence of Γ_C by an iterative argument that is very similar to the second half of the proof of Theorem 10.1 in [**5**]. A spanning tree Γ containing Σ (or rather, the set of edges $\Gamma \setminus \Sigma$) can be viewed as a rooted tree in which the connected components of Σ play the role of vertices and C is the root. Let us call an edge of $\Gamma \setminus \Sigma$ *bad* if its emerald endpoint faces C and let us call other edges of $\Gamma \setminus \Sigma$ *good*.

In order to be more precise, for any spanning tree Γ containing Σ, let us consider the tree Γ^{red} which results from contracting each connected component of Σ to a point (vertex). The edges of Γ^{red} inherit a good/bad classification from Γ. The vertex corresponding to C will be treated as the root of Γ^{red}. Given a rooted tree, we may speak of the distance $d(\varepsilon) \in \mathbf{N}$ of each edge ε to the root (those adjacent to the root have $d = 0$ etc.). Let us now associate the following quantities to Γ.

- Let $n(\Gamma)$ denote the smallest value of d among bad edges of Γ^{red}. If there are no bad edges, let $n(\Gamma)$ be one more than the maximal value of d among the edges of Γ^{red}.
- For $0 \le m < n(\Gamma)$, let $\lambda_\Gamma(m)$ be the number of edges ε of Γ^{red} with $d(\varepsilon) = m$. These values are positive. For $m \ge n(\Gamma)$, we let $\lambda_\Gamma(m) = 0$.

Then for a pair of spanning trees Γ_1, Γ_2 extending Σ, we write $\Gamma_1 \prec \Gamma_2$ if either

(i) the sequence $\lambda_{\Gamma_1}(0), \lambda_{\Gamma_1}(1), \lambda_{\Gamma_1}(2), \ldots$ is smaller in lexicographic order than the sequence $\lambda_{\Gamma_2}(0), \lambda_{\Gamma_2}(1), \lambda_{\Gamma_2}(2), \ldots$, or

(ii) the two sequences coincide (implying $n(\Gamma_1) = n(\Gamma_2) = n$) but the number of bad edges with $d = n$ is higher in Γ_1^{red} than in Γ_2^{red}.

The relation $\prec$ is a so called strict weak order on the set of spanning trees of G containing Σ, in particular it is (obviously) transitive and asymmetric. To finish the proof of the Lemma, it suffices to show that if a tree $\Gamma \supset \Sigma$ corresponds to a maximal simplex in $\mathcal{T}$ and does not satisfy the conditions for Γ_C (i.e., Γ has bad edges), then there is another tree $\widetilde{\Gamma}$ containing

Σ, with $\Gamma \prec \widetilde{\Gamma}$, which also appears as a maximal simplex in $\mathcal{T}$. Let us pick a bad edge ε of Γ from among those that are closest to the root in Γ^{red}. As Σ represents an interior face of $\mathcal{T}$, so does the larger set $\Gamma \setminus \{\varepsilon\}$. This means that this codimension one simplex bounds maximal simplices on both sides. I.e., there exists an edge δ of G so that $\widetilde{\Gamma} = (\Gamma \setminus \{\varepsilon\}) \cup \{\delta\}$ corresponds to a maximal simplex of $\mathcal{T}$. An application of Lemma 3.4 to Γ and $\widetilde{\Gamma}$ shows that δ is a good edge for $\widetilde{\Gamma}$.

It is easy to check that $\Gamma \prec \widetilde{\Gamma}$ is indeed true. If δ is closer to the root in $\widetilde{\Gamma}^{\mathrm{red}}$ than ε was in Γ^{red}, or if ε was the only bad edge of Γ^{red} at distance $d(\varepsilon)$ from the root, then the relation holds by part ((i)) of its definition. Otherwise it holds by part ((ii)). This completes the proof. □

The following 'path lemma' will soon be useful as well. The term 'transfer of valence' was introduced right after Definition 2.3.

Lemma 4.2. *Let Γ be a spanning tree of the connected bipartite graph G that induces the hypertree $\mathbf{f}$ in $\mathcal{H} = (V, E)$ (here as usual, E and V are the colour classes of G). Let a, b, c be distinct elements of E so that the unique path from c to a in Γ passes through b. Suppose that $\mathbf{f}$ is such that c can transfer valence to a. Then $\mathbf{f}$ is also such that b can transfer valence to a.*

Proof. Assume that b cannot transfer valence to a. Then there exists a set T of hyperedges that is tight at $\mathbf{f}$, contains a, but does not contain b. (Cf. Lemma 2.7 and the discussion right after it.) We have to have $c \in T$ for otherwise no transfer would be possible from c to a. For the same reason, in view of Lemma 2.9, a and c have to be in the same connected component of the forest $\Gamma\big|_T$. Hence a and c can be connected by a path in $\Gamma\big|_T$, but since $b \notin T$, that is different from the path, in Γ, that passes through b. This contradicts the cycle-freeness of Γ. □

Let $\mathcal{T}$ be a fixed triangulation of Q_G and let us also fix an order of the colour class E (which we also think of as the set of hyperedges in the hypergraph (V, E)). Let γ be a maximal simplex of $\mathcal{T}$ and Γ the corresponding spanning tree of G. Let Γ induce the hypertree $\mathbf{f}$ and let $e \in E$ be an internally inactive hyperedge with respect to $\mathbf{f}$. Let e' be the smallest hyperedge so that $\mathbf{f}$ is such that e can transfer valence to e', and let the hypertree $\mathbf{g}$ be the result of that transfer.

We will refer to the hyperedge e' above as the *favourite* of e (with respect to the chosen order). If e is indeed internally inactive then we have $e' < e$; for an internally active hyperedge, let us define its favourite to be itself.

Let us now connect the two emerald markers $\mathbf{f}^+$ and $\mathbf{g}^+$ (cf. (3.11)) by a straight line segment, which we call the *feeler* for the pair $(\mathbf{f}, e)$. (So technically, the feeler is defined to be a subset of the Minkowski sum P_E discussed in Section 3, but we may also think of it as a subset of the cross-section $S_E \subset Q_G$ since P_E and S_E are related by a dilation.) Let γ' be the maximal simplex that the feeler enters right after it leaves γ. (In Lemma 4.3 we will see that the feeler exits γ through a relative interior point of a facet so that it enters the interior of the next maximal simplex. It is possible for $\mathbf{g}^+$ to be the emerald marker of γ', but it will typically not be the case.) If Γ' is the spanning tree that corresponds to γ', then the symmetric difference of Γ and Γ' consists of exactly two edges of G. In the next Lemma we identify one of those two edges (under slightly more general hypotheses).

Lemma 4.3. *Let the hypertrees $\mathbf{f}, \mathbf{g} \in B_{\mathcal{H}}$ differ by a single transfer of valence from e to e'. Let us connect the emerald markers $\mathbf{f}^+$ and $\mathbf{g}^+$ by a line segment l in P_E. If $\mathbf{f}^+$ is an interior point of the Minkowski cell M_Γ (for a spanning tree Γ inducing $\mathbf{f}$, cf.* (3.10)*) then l leaves M_Γ*

through a codimension 1 *stratum of its boundary which corresponds to removing from* Γ *the first edge along the unique path from* e *to* e'*.*

Proof. In Lemma 3.14 we described the starting point $\mathbf{f}^+$ of l in terms of the direct sum decomposition (3.10) of M_Γ. We will extend that description to the initial segment of l. Let the path p in Γ from e to e' be $\varepsilon_1, \delta_1, \varepsilon_2, \delta_2, \dots, \varepsilon_k, \delta_k$ and assume that at $\mathbf{f}^+$ these edges have the weights $u_1, v_1, u_2, v_2, \dots, u_k, v_k$, respectively, as given in Lemma 3.14. (Note that emerald and violet vertices alternate along p and since e and e' are both emerald, there has to be an even number of edges.) Then we claim that l is parametrised by the weights

$$u_1 - t, v_1 + t, u_2 - t, v_2 + t, \dots, u_k - t, v_k + t \tag{4.1}$$

along p, while all other weights are constant. Indeed, this ensures that sums of weights on edges adjacent to a violet vertex remain fixed at 1; the sums of weights on edges adjacent to an emerald vertex also remain constant except in the cases of e and e', which is consistent with the direction of l.

The segment l reaches the boundary of γ (which is the maximal simplex corresponding to Γ) when it reaches the boundary of its (dilated) cross-section M_Γ, i.e., when one of the weights becomes zero. By (4.1), this will occur on one of the ε_i. In fact, from Lemma 3.14 it easily follows that $u_1 < u_2 < \cdots < u_k$, from which we see that it is the edge ε_1 (and only that) that gets removed from Γ when we pass from γ to the adjacent maximal simplex. □

Note that Lemma 4.3 does not provide any information on the unique edge in $\Gamma' \setminus \Gamma$. That depends on the nature of the triangulation $\mathcal{T}$. If the feeler intersects only the maximal simplices marked by $\mathbf{f}^+$ and $\mathbf{g}^+$, then this edge will be adjacent to e'; otherwise we have basically no control over it.

Remark 4.4. Let us substitute $t = u_1$ in (4.1) and use the resulting values as new weights on the edges $\varepsilon_1, \delta_1, \varepsilon_2, \delta_2, \dots, \varepsilon_k, \delta_k$, while we keep the weights provided by $\mathbf{f}^+$ on other edges of Γ. These are the barycentric coordinates (in the simplex γ that corresponds to Γ) of the point where the feeler leaves the Minkowski cell, i.e., the simplex. But they also correspond to the following operation. Let C be the component of $\Gamma \setminus \{\varepsilon_1\}$ that does not contain e'. Now remove ε_1 from Γ and then identify e and e' to obtain a new tree. (One can imagine this as transporting C along the path p from e to e'.) Let the merged vertex have multiplicity 2. If we apply the edge-weight formula of Lemma 3.14 to this tree, it is easy to see that we obtain exactly the weights we have just described. This observation may sound ad hoc but it will be useful in the proof of Theorem 5.1 below.

The last result in this section is an extension of [**5**, Lemma 7.4]. It allows us to 'generate' (i.e., guarantees the existence of) hypercubes of hypertrees in certain situations. We state it in the hypergraph context even though it holds for integer polymatroids as well.

Lemma 4.5. *Let* $\{ a_1, \dots, a_m \}$ *and* $\{ x_1, \dots, x_m \}$ *be disjoint collections of hyperedges in the connected hypergraph* $\mathcal{H} = (V, E)$*. Let* $\mathbf{f} \colon E \to \mathbf{N}$ *be a hypertree which is such that* x_j *can transfer valence to* a_j *for all* j *but* x_j *cannot transfer valence to* a_i *for any* $i < j$*. Then* $\mathbf{f}$ *is such that any subset of the transfers mentioned above is simultaneously possible, that is, for any subset* $J \subset \{ 1, \dots, m \}$*, the vector* $\mathbf{f} + \sum_{j \in J} (\mathbf{i}_{\{a_j\}} - \mathbf{i}_{\{x_j\}})$ *is another hypertree.*

Proof. Without loss of generality we may assume that $J = \{1, \ldots, m\}$ and $m \geq 2$. If the statement fails to be true then there exists a smallest index k so that $\mathbf{u} = \mathbf{f} + \sum_{j=1}^{k-1}(\mathbf{i}_{\{a_j\}} - \mathbf{i}_{\{x_j\}})$ is a hypertree but $\mathbf{f} + \sum_{j=1}^{k}(\mathbf{i}_{\{a_j\}} - \mathbf{i}_{\{x_j\}})$ is not. The latter implies that there is a set K of hyperedges that is tight at $\mathbf{u}$, contains a_k but does not contain x_k. Now at $\mathbf{f}$, there exist tight sets $L_1, \ldots, L_{k-1}$ of hyperedges so that L_i contains a_i but does not contain x_k. By Lemma 2.8, $L = L_1 \cup \cdots \cup L_{k-1}$ is also tight at $\mathbf{f}$. Since $a_1, \ldots, a_{k-1} \in L$, we have to have $x_1, \ldots, x_{k-1} \in L$ too for the first $k-1$ of our assumed transfers to be possible. That implies that L is also a tight set at $\mathbf{u}$ and hence the same is true for $L \cup K$. But since $\mathbf{u}$ and $\mathbf{f}$ produce the same sum of values over elements of $L \cup K$, we obtain that $L \cup K$ is a tight set at $\mathbf{f}$. That is a contradiction because $L \cup K$ separates x_k from a_k. □

5. *The main result*

To prove Theorem 1.1, it remains to show that the interior polynomial I of the connected hypergraph (V, E) satisfies

$$I(x) = a_0 + a_1 x + \cdots + a_d x^d. \tag{5.1}$$

First we give the reduction of this claim to a somewhat technical result, Theorem 5.1, and then proceed to prove the latter.

The sequence $a_0, a_1, \ldots, a_d$ of (5.1), where first of all $d = |E| + |V| - 2$ is the dimension of the root polytope Q_G, was defined in (1.2) (and then again in (3.3)) in terms of the Ehrhart polynomial of Q_G. However, in Theorem 3.10 we have already equated $\sum_{k=0}^{d} a_k x^k$ to the right hand side of (1.1) so that from now on we may think of $a_0, a_1, \ldots, a_d$ as the coefficient sequence of the h-vector h_G. In particular, as we have seen in Remark 3.12, if a triangulation of Q_G has a shelling order then a_k is the number of maximal simplices that are adjacent (through a facet) to exactly k 'previous' maximal simplices.

Let us fix a shellable triangulation $\mathcal{T}$ (such as a regular triangulation†, cf. [**13**]) of Q_G with f-vector $f(y)$ as in (3.1). Let the number of interior simplices of $\mathcal{T}$ of dimension m be $\tilde{f}_m$ for $m = 0, 1, \ldots, d$. (Note that $\tilde{f}_0 = 0$ and $\tilde{f}_d = f_d$.) Similar to the proof of Theorem 3.10 (and especially to Remark 3.12), we may use $\mathcal{T}$ in conjunction with Lemma 3.3 to count *interior* lattice points in the dilation $s \cdot Q_G$ (where s is a positive integer). This time we will count points along interior facets when the *second* maximal simplex adjacent to the facet is encountered in the shelling order. By doing so we find that the number of such points is

$$\tilde{\varepsilon}_G(s) = a_0 \binom{s-1}{d} + a_1 \binom{s}{d} + \cdots + a_k \binom{s-1+k}{d} + \cdots + a_d \binom{s-1+d}{d}.$$

This is because of the following reason: in the interior of each maximal simplex there are $\binom{s-1}{d}$ lattice points (cf. Remark 3.9) whose convex hull is a d-dimensional simplex of sidelength $s-d-1$. (We may assume here that $s \geq d+1$ since from Ehrhart theory we know that the result of our lattice point count is a polynomial in s, determined uniquely by just finitely many of its values.) If the interior (to $s \cdot Q_G$) lattice points along k facets of the maximal simplex are to be counted as well, the sidelength of the convex hull (which is still a simplex) increases by k so that the number of lattice points becomes $\binom{s-1+k}{d}$.

---

†Roughly speaking, a regular triangulation is one obtained by assigning a (generic) new coordinate to each vertex, thereby lifting them to one dimension higher, then taking their convex hull and projecting 'back down' its upper boundary. Not all triangulations of a root polytope are regular, not even for (small) complete bipartite graphs [**11**].

On the other hand, using the obvious partition of the interior of Q_G and Remark 3.9, the same quantity can also be expressed as

$$\tilde{\varepsilon}_G(s) = \tilde{f}_d\binom{s-1}{d} + \tilde{f}_{d-1}\binom{s-1}{d-1} + \cdots + \tilde{f}_{d-l}\binom{s-1}{d-l} + \cdots + \tilde{f}_1\binom{s-1}{1}.$$

Let us manipulate the first expression using Vandermonde's identity[†]

$$\binom{s-1+k}{d} = \binom{k}{0}\binom{s-1}{d} + \binom{k}{1}\binom{s-1}{d-1} + \cdots + \binom{k}{l}\binom{s-1}{d-l} + \cdots + \binom{k}{k}\binom{s-1}{d-k}$$

and equate coefficients of $\binom{s-1}{d-l}$ (noting that these polynomials have different degrees in s and hence they are linearly independent) to find that

$$\tilde{f}_{d-l} = \binom{l}{l}a_l + \binom{l+1}{l}a_{l+1} + \cdots + \binom{d}{l}a_d.$$

Notice that the right hand side is exactly the coefficient of $(x-1)^l$ in the Taylor expansion, centred at 1, of $\sum_{k=0}^{d} a_k x^k$. Therefore to show (5.1), it suffices to prove that $I(x)$ has the same expansion, that is that the coefficient of $(x-1)^k$ in the Taylor expansion of $I(x)$ centred at 1 is $\tilde{f}_{d-k}$ for all $k \geq 0$ (here for negative u we define $\tilde{f}_u = 0$). In other words, what is left to prove is that

$$I^{(k)}(1) = k! \cdot \tilde{f}_{d-k} \tag{5.2}$$

holds for all $k \geq 0$. We note that the case when $k = 0$ is settled in [**9**].

Now if we recall Definition 2.4 of the interior polynomial, we see that the Taylor coefficient in question is $\sum_{i=k}^{|E|-1}\binom{i}{k}\chi_i$, where χ_i is the number of those hypertrees in (V, E) whose internal inactivity is exactly i. Put another way, it is the number of tuples consisting of a hypertree and exactly k of its internally inactive hyperedges. Therefore Theorem 1.1 is a corollary of the following statement. (Note that from here on we will not rely on the condition of shellability.)

THEOREM 5.1. *Let $\mathcal{T}$ be a triangulation of the root polytope Q_G of the connected bipartite graph G, and let $k \geq 0$ be an integer. Then the number of codimension k interior simplices of $\mathcal{T}$ agrees with the number of pairs $(\mathbf{f}, S)$, where $\mathbf{f}$ is a hypertree in the hypergraph (V, E) induced by G and S is a set of k distinct internally inactive hyperedges with respect to $\mathbf{f}$. Here inactivity is defined with respect to an arbitrarily fixed order on E.*

Proof. As the case $k = 0$ follows from [**9**, Lemma 12.8 and Theorem 12.9], we may assume that $k \geq 1$. We shall define a map σ from the second named set to the first and show that it is one-to-one and onto.

If the pair $(\mathbf{f}, \{e_1, \ldots, e_k\})$ is as described in the Theorem, then first we find the unique maximal simplex $\gamma_{\mathbf{f}}$ in $\mathcal{T}$ which belongs to a spanning tree $\Gamma_{\mathbf{f}}$ of G that induces $\mathbf{f}$. (See [**9**] or subsection 3.4.) For each $1 \leq i \leq k$, we construct the feeler for the pair $(\mathbf{f}, e_i)$. By Lemma 4.3, the i'th feeler leaves $\gamma_{\mathbf{f}}$ through a facet that corresponds to (removing) an edge ε_i of $\Gamma_{\mathbf{f}}$ which is adjacent to e_i. Hence these k facets are mutually different and thus their intersection is a codimension k face $\sigma(\mathbf{f}, \{e_1, \ldots, e_k\})$ of $\mathcal{T}$.

Well-definedness of σ. We need to show that the face $\sigma(\mathbf{f}, \{e_1, \ldots, e_k\})$ lies in the interior of Q_G. This is because one is able to find enough hypertrees near $\mathbf{f}$ (namely the endpoints of the k feelers and, using Lemma 4.5, some others) so that the convex hull of the corresponding emerald markers, which is k-dimensional and interior to Q_G, meets $\sigma(\mathbf{f}, \{e_1, \ldots, e_k\})$ in just one point and that point is in the relative interior of both sets. A detailed computation follows.

[†]This step is also found in one of Andrews's proofs [**1**] of the Saalschütz formula.

Let the favourite of e_i be e'_i ($1 \le i \le k$). (Favourites were defined in Section 4.) The sets $\{ e_1, \dots, e_k \}$ and $\{ e'_1, \dots, e'_k \}$ are disjoint by [**5**, Lemma 5.2], which says that if a hypertree is such that the hyperedge c can transfer valence to b and b can transfer valence to a, then c can transfer valence to a as well. As $\{ e'_1, \dots, e'_k \}$ may have less than k elements, let us also write it as $S' = \{ a_1, a_2, \dots, a_m \}$ where we assume $a_1 < a_2 < \dots < a_m$. There is then an obvious partition of $S = \{ e_1, \dots, e_k \}$ into m parts $S_1, \dots, S_m$ so that the favourite of each element of S_i is a_i.

For each $i = 1, \dots, m$, we may write an arbitrary linear combination $\mathbf{v}_i$ of the vectors $\mathbf{i}_{\{a_i\}} - \mathbf{i}_{\{e\}}$ for $e \in S_i$ so that the coefficients are positive and their sum is less than 1. Then, by Lemma 4.5, we may add the sum of these vectors to the emerald marker $\mathbf{f}^+$ (cf. (3.11)) and obtain an interior point of Q_G that way. It is enough to show that $\sigma(\mathbf{f}, \{ e_1, \dots, e_k\})$ contains one of these points.

Since for all $i = 1, \dots, m$, the hypertree $\mathbf{f}$ is such that no hyperedge from $K_i = S_{i+1} \cup \dots \cup S_m$ can transfer valence to a_i, there exist sets of hyperedges that are tight at $\mathbf{f}$ and separate a_i from the elements of K_i. The intersection T'_i of these sets is also tight at $\mathbf{f}$ (by Lemma 2.8), contains a_i (hence it contains S_i, too), and is disjoint from K_i. Lemmas 4.2 and 2.9 allow us to assume that $\Gamma_{\mathbf{f}}\big|_{T'_i}$ is connected: indeed, all hyperedges along paths in $\Gamma_{\mathbf{f}}$ between elements of S_i and a_i are such that they can transfer valence to a_i and thus they have to be in T'_i. By letting $T_i = T'_1 \cup \dots \cup T'_i$, we get a nested sequence $T_1 \subset \dots \subset T_m$ of sets, each of which is tight (again by Lemma 2.8) at $\mathbf{f}$. Let us reiterate that

$$\text{for each } e_j \in S, \text{ the path } p_{e_j} \text{ in } \Gamma_{\mathbf{f}} \text{ from } e_j \text{ to } e'_j \text{ lies in } \Gamma_{\mathbf{f}}\big|_{T_i}, \tag{5.3}$$

where i is the index so that $e'_j = a_i$, i.e., that $e_j \in S_i$.

We will construct the vectors $\mathbf{v}_i$ mentioned above inductively, starting from $\mathbf{v}_m$. (One way to think about this next part of the proof is that we will build a polygonal path in P_E from $\mathbf{f}^+$ to $\sigma(\mathbf{f}, \{ e_1, \dots, e_k\})$, that is from an interior point of the Minkowski cell $M_{\Gamma_{\mathbf{f}}}$ of (3.10) to a codimension k stratum of its boundary. We will choose the length of each segment carefully so that the path stays within the cell.) The set S_m has a partial order $<$ defined by the rooted tree $(\Gamma_{\mathbf{f}}, a_m)$, namely we let $x < y$ if the path in $\Gamma_{\mathbf{f}}$ from y to a_m passes through x. The coefficients of the vectors $\mathbf{i}_{\{a_m\}} - \mathbf{i}_{\{e\}}$ in $\mathbf{v}_m$ (where $e \in S_m$) are again defined inductively starting from the maximal elements of S_m. If e is one of those and the weight, in the description of $\mathbf{f}^+$ provided by Lemma 3.14, is u on the edge ε that starts the path p_e from e to a_m in $\Gamma_{\mathbf{f}}$, then let the coefficient of $\mathbf{i}_{\{a_m\}} - \mathbf{i}_{\{e\}}$ be u. Because (by the definition of the partial order and the observation (5.3) above) none of the paths $p_{e'}$ for $e \neq e' \in S$ passes through e, in particular none of them contains ε, this is the last time in the process that the weight on ε changes (cf. (4.1) in the proof of Lemma 4.3). Hence our choice u for the coefficient of $\mathbf{i}_{\{a_m\}} - \mathbf{i}_{\{e\}}$ is the unique way we can reduce the weight on ε to 0.

As discussed in the proof of Lemma 4.3, edge weights (barycentric coordinates) representing the point $\mathbf{f}^+ + u(\mathbf{i}_{\{a_m\}} - \mathbf{i}_{\{e\}})$ of $\gamma_{\mathbf{f}}$ are positive except for that of ε. As explained in Remark 4.4, this is 'demonstrated' by applying the formula of Lemma 3.14 to a rearranged tree, where we erase ε and transport the part C of $\Gamma_{\mathbf{f}}$ that 'lies beyond' e to the favourite a_m of e. Since paths $p_{e'}$ for $e \neq e' \in S$ (along which weight changes are due to occur) are either disjoint from, or contained in C, this transporting operation does not disrupt the structure of the tree that is relevant for the rest of the construction.

Next, we iterate the process that was described in the previous two paragraphs. There will be one step for each element e of S_m and it will occur after all elements of S_m above e in the partial order have been dealt with. During the step we determine the coefficient u_e of the vector $\mathbf{i}_{\{a_m\}} - \mathbf{i}_{\{e\}}$ in the linear combination $\mathbf{v}_m$, namely we define u_e to be the current weight on the edge ε_e which is adjacent to e and points toward a_m. In light of (4.1) and the fact that ε_e is not part of the path p_h for any element $h \in S$ that we will address later (which is, again, guaranteed by (5.3) and the partial order), this is the unique choice that results in the weight on ε_e becoming 0. At the end of each step we update the weights on the edges of $\Gamma_{\mathbf{f}}$ using (4.1), or rather on the edges of the tree that results from our previous transports, and note that these new weights are the same as what the recipe of Lemma 3.14 produces for the rearranged tree where we erase ε_e and identify e with a_m, denoting the new point by a_m and increasing its multiplicity by 1. In particular, all edge weights in the rearranged tree are positive.

After we have assigned coefficients to each vector $\mathbf{i}_{\{a_m\}} - \mathbf{i}_{\{e\}}$ where $e \in S_m$, noting that they are all positive, we also have to make sure that the sum of the coefficients is less than 1. That is because it is equal to the proportion, among all elements of E, of the emerald vertices of all the (disjoint) subtrees that have been transported to a_m so far in the process. This set includes the vertices that were merged with a_m but it does not include the 'original' vertex a_m itself, and hence its proportion in E is less than 1.

After this we continue performing the same operations as above (to the rearranged tree and the updated weights) for a_{m-1}, a_{m-2} and so forth down to a_1. By the exact same arguments as above, each vector $\mathbf{i}_{\{a_i\}} - \mathbf{i}_{\{e\}}$, for $e \in S_i$, will receive a positive coefficient in the linear combination $\mathbf{v}_i$. The sum of the coefficients, for each given i, is strictly less than 1. These coefficients are also the unique choices that result in a weight of 0 on each edge ε_i, $1 \le i \le k$, i.e., which guarantee that $\mathbf{f}^+ + \sum_{i=1}^k \mathbf{v}_i$ is a point of the simplex $\sigma(\mathbf{f}, S)$. We also see that it is an interior point of said simplex because the weights on edges other than the ε_i remain positive at the end. This completes the proof of well-definedness for our correspondence σ.

Let now σ be a codimension k interior face of $\mathcal{T}$. Let the edges of G which correspond to the vertices of σ form the $(k+1)$-component forest Σ. From the point of view of Σ, the statement that $\sigma = \sigma(\mathbf{f}, \{\, e_1, \dots, e_k \})$ for some pair $(\mathbf{f}, \{\, e_1, \dots, e_k \})$ can be described as follows. We add k distinct edges $\varepsilon_1, \dots, \varepsilon_k$ of G to Σ to obtain a spanning tree Γ which is a realisation of $\mathbf{f}$, and which is also such that the corresponding maximal simplex is in $\mathcal{T}$. For each i, the emerald endpoint of ε_i is e_i; furthermore and most crucially, ε_i is the first edge along the unique path in Γ that connects e_i to its favourite (in relation to $\mathbf{f}$) hyperedge. In particular, the emerald endpoints of the 'extra' edges $\varepsilon_1, \dots, \varepsilon_k$ are all distinct and, obviously, Σ and the extra edges determine the pair $(\mathbf{f}, \{\, e_1, \dots, e_k \})$. Let us call the collection $\varepsilon_1, \dots, \varepsilon_k$ *good* with respect to σ if it fits the description above (with the given triangulation $\mathcal{T}$ and for the pair $(\mathbf{f}, \{\, e_1, \dots, e_k \})$ that the collection determines).

We need to show that for the given σ, there exists a unique collection $\varepsilon_1, \dots, \varepsilon_k$ of edges not in Σ that is good in the sense of the previous paragraph. We start with uniqueness.

Injectivity of σ. Let $\varepsilon_1, \dots, \varepsilon_k$ and $\delta_1, \dots, \delta_k$ be two different good collections. If we enlarge Σ with the edges common to them, their remaining parts are still good collections for this larger forest/interior cell. (Note that favourites of hyperedges are decided not by Σ but by the actual spanning tree containing it and the latter remains the same even as we declare Σ to be bigger.) Hence we may assume that the two collections are disjoint.

Let the hyperedge a be the smallest favourite that occurs among simplices of $\mathcal{T}$ containing σ. (By this we mean the favourites of those hyperedges that received an extra edge in the extension and became internally inactive with respect to the hypertree induced by the extension. That set is non-empty because $k \ge 1$ and the emerald endpoint of ε_1 has a favourite different from itself.) Let Γ_a be the spanning tree in G provided by Lemma 4.1 applied to Σ and its component containing a. That is, Γ_a represents a maximal simplex in $\mathcal{T}$, contains Σ, and has all its edges beyond Σ face a with their violet endpoint. We are going to argue that Γ_a and the spanning trees $\Gamma^\varepsilon = \Sigma \cup \{\, \varepsilon_1, \dots, \varepsilon_k \,\}$ and $\Gamma^\delta = \Sigma \cup \{\, \delta_1, \dots, \delta_k \,\}$ cannot all be compatible (i.e., it is not possible for all three pairs to satisfy the condition given in Lemma 3.4), which will contradict their coexistence in $\mathcal{T}$.

Let the emerald endpoints of the ε_i form the set $S^\varepsilon \subset E$, and let the emerald endpoints of the δ_j form the set $S^\delta \subset E$. Within these, let S^ε_a and S^δ_a denote the sets of those hyperedges whose favourite is a. If $e \in S_a^\varepsilon$, then by Lemma 4.2, all elements of S^ε along the path p_e in Γ^ε from e to a also belong to S_a^ε, and hence the edges ε_i adjacent to them belong to p_e so that their violet endpoints face a. Similar observations hold true for S_a^δ. If a component of Σ contains an element of S_a^ε as well as an element of S_a^δ then, just like in the proof of Lemma 4.1, it is easy to see that Γ^ε and Γ^δ are not compatible trees. In the same way we obtain that

$$\text{the edges } \varepsilon_i \text{ adjacent to elements of } S_a^\varepsilon\text{, as well as the edges } \delta_j \text{ adjacent to elements of } S_a^\delta\text{, are edges in } \Gamma_a\text{, too.} \quad (5.4)$$

Let now $\mathbf{f}^\varepsilon$ be the hypertree induced by Γ^ε. For each element $e \in S^\varepsilon \setminus S_a^\varepsilon$, there is a set of hyperedges that is tight at $\mathbf{f}^\varepsilon$, contains a, and does not contain e. The intersection T^ε of these sets is also tight by Lemma 2.8, it is disjoint from $S^\varepsilon \setminus S_a^\varepsilon$, and it contains S_a^ε for otherwise those hyperedges could not transfer valence to a. By the observation (5.4), we see that $\sum_{e\in T^\varepsilon} \mathbf{f}^\varepsilon(e) \le \sum_{e\in T^\varepsilon} \mathbf{f}_a(e)$ but since $\sum_{e\in T^\varepsilon} \mathbf{f}^\varepsilon(e)$ is already at the maximal value $\mu(T^\varepsilon)$ allowed for any hypertree, it follows that

- T^ε is also a tight set at the hypertree $\mathbf{f}_a$ induced by Γ_a and
- $\Gamma_a\big|_{T^\varepsilon} = \Gamma^\varepsilon\big|_{T^\varepsilon}$ consists only of edges of Σ and the ε_i adjacent to elements of S_a^ε.

By the same token there is a set $T^\delta \subset E$ that is tight at $\mathbf{f}_a$, contains a, and intersects S^δ in exactly S_a^δ. By Lemma 2.8, $T = T^\varepsilon \cap T^\delta$ is also tight at $\mathbf{f}_a$. It contains a and since $S_a^\varepsilon \cap S_a^\delta = \varnothing$, the forest $\Gamma_a\big|_T = \big(\Gamma_a\big|_{T^\varepsilon}\big) \cap \big(\Gamma_a\big|_{T^\delta}\big)$ consists only of edges of Σ. However that means that T is tight not just at $\mathbf{f}_a$ but at any hypertree induced by an extension of Σ (so that all these hypertrees are such that a cannot receive transfers of valence from outside T) and furthermore that hyperedges in T may not get extra edges in such extensions. This contradicts the definition of a.

Surjectivity of σ. It remains to show that any $(k+1)$-component forest $\Sigma \subset G$, provided that it represents an interior face of $\mathcal{T}$, can be extended to a spanning tree by a good collection of edges so that the corresponding maximal simplex occurs in $\mathcal{T}$. It is tempting to think that the tree Γ_a defined above (i.e., the set of its edges beyond Σ) would satisfy the requirement but, unfortunately, that is not always the case (see Example 5.2 below). Instead, we devise the following iterative procedure.

Let a_1 be the smallest element of E. Let Γ_1 be the spanning tree of G that results from applying Lemma 4.1 to Σ and its component containing a_1. In the set S_1 of the k emerald endpoints of the edges in $\Gamma_1 \setminus \Sigma$, let us consider those that cannot transfer valence to a_1. These are separated from a_1 by sets that are tight at the hypertree $\mathbf{f}_1$ induced by Γ_1. The intersection T_1 of these sets is also tight at $\mathbf{f}_1$, it contains a_1, and hence it intersects S_1 in exactly those of its elements that can transfer valence to a_1 at $\mathbf{f}_1$. By Lemmas 4.2 and 2.9, we may assume that $\Gamma_1\big|_{T_1}$ is connected. We add to Σ the edges of $\Gamma_1 \setminus \Sigma$ that are adjacent to elements of $S_1 \cap T_1$. During the rest of the construction we will keep the subgraph $\Gamma_1\big|_{T_1}$ fixed so that T_1 remains a tight set at subsequent hypertrees, too. This serves to 'protect' the transfers of valence that we have found from elements of $S_1 \cap T_1$ to a_1, as well as to maintain the situation that the extra edges adjacent to those emerald vertices point toward a_1. Indeed, if an element of $S_1 \cap T_1$ got separated from a_1 by a tight set at a later stage, then the intersection of that set with T_1 would also be tight at that hypertree; but then the same set would be tight at $\mathbf{f}_1$ as well, contradicting the possibility of the transfer at $\mathbf{f}_1$.

Next, let a_2 be the smallest hyperedge not in T_1. Let us apply Lemma 4.1 to the forest $\Sigma \cup (\Gamma_1\big|_{T_1})$ and its component containing a_2 to obtain the spanning tree Γ_2 that induces the hypertree $\mathbf{f}_2$. There is again a tight set T_2' of hyperedges that contains a_2 and all emerald endpoints of the edges in $\Gamma_2 \setminus \Sigma \setminus (\Gamma_1\big|_{T_1})$ that can transfer valence to a_2 at $\mathbf{f}_2$ but does not contain any of those that cannot. We may assume that $\Gamma_2\big|_{T_2'}$ is connected. We let $T_2 = T_1 \cup T_2'$; by Lemma 2.8, this set is also tight at $\mathbf{f}_2$, as well as at later hypertrees. The set T_2 will protect the transfers of valence that we have just found to a_2. The set T_1, on the other hand, will guarantee that the givers of those transfers will not find receivers smaller than a_2 at any time during the rest of the process.

We iterate our construction, always taking the smallest hyperedge a_{i+1} outside of the last element T_i of our sequence of nested tight sets. In each step a_{i+1}, possibly with other hyperedges, will be added to T_i to form the next tight set T_{i+1}. Thus the sequence $\{T_i\}$ is strictly increasing, and eventually the tight set will be E itself. At that stage, all emerald vertices that have an extra edge adjacent to them have it point in the direction of their favourite. This completes the proof of surjectivity and hence the proof of the Theorem, as well as the proof of Theorem 1.1. □

Example 5.2. We use the plane bipartite graph G of Example 2.6 to illustrate some of the subtlety of the proof of Theorem 5.1. Let us consider one of the triangulations provided by [**6**, Theorem 1.1]. Namely, we orient the edges of the dual graph G^* so that each has an emerald point to its right, we fix the root of G^* in the outside region, construct all (say, outgoing) spanning arborescences relative to it, and then consider the spanning trees of G that are dual to those. The maximal simplices of Q_G that correspond to the latter form a triangulation $\mathcal{T}$.

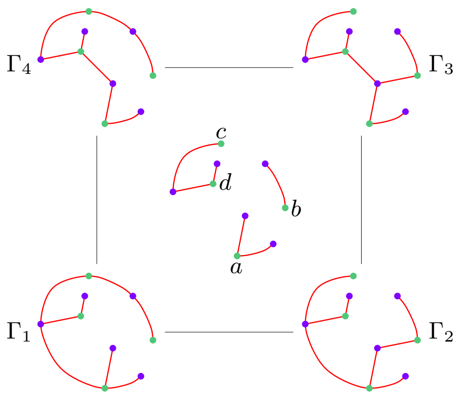

Figure 5. *A codimension 2 interior simplex and the four maximal simplices containing it, each represented by the corresponding cycle-free subgraph.*

We let Σ be the three-component forest shown in the middle of Figure 5. The corresponding codimension 2 simplex σ lies in the interior of Q_G and it is contained in exactly four maximal simplices. These are represented by the corresponding spanning trees Γ_i, $1 \le i \le 4$. The four black line segments indicate adjacency through a facet. Let us also use the symbol $\mathbf{f}_i$ to denote the hypertree (in the hypergraph with emerald hyperedges) induced by Γ_i.

Let us choose the order $a < b < c < d$ on the emerald colour class E. First we illustrate the last (surjectivity) part of the proof. The hypertree $\mathbf{f}_2$ is the one that extends the valence distribution of Σ at the lexicographically smallest multiset, namely $\{\,a, b\,\}$. That hypertree does not in general have to be realised by a maximal simplex adjacent to Σ and in any case, only d is internally inactive with respect to $\mathbf{f}_2$ and d did not even receive an extra edge when Σ was extended to Γ_2. The spanning tree that results if we apply Lemma 4.1 to the component of Σ containing a is Γ_3. The only internally inactive hyperedge with respect to $\mathbf{f}_3$ is d. Indeed, the set $\{\,a, d\,\}$ is tight at $\mathbf{f}_3$ so that b cannot transfer valence to a, regardless of the fact that it received an extra edge pointing toward a. The unique good collection of extra edges for Σ is in fact the one that extends it to Γ_4. With respect to $\mathbf{f}_4$, the hyperedge d is internally inactive and its favourite is a, but since the set $\{\,a, d\,\}$ is still tight, the favourite of the other internally inactive hyperedge, c, ends up being b. In the case of Γ_4 the extra edges are received exactly by d and c and they indeed point toward their respective favourites.

This same example is a good illustration of the first (well-definedness) part of the proof, too. If we think of the black line segments of Figure 5 as connections between the appropriate emerald markers $\mathbf{f}_i^+$, then in that sense they form a proper regular square. Each side is in fact a feeler (in this case neither crosses into more than two maximal simplices). For example with respect to $\mathbf{f}_1$, the hyperedge c is internally inactive and its favourite is b. Since a transfer of valence at $\mathbf{f}_1$ from c to b results in $\mathbf{f}_2$, the corresponding feeler connects $\mathbf{f}_1^+$ to $\mathbf{f}_2^+$. In S_E (cf. (3.8)), the a-coordinates of the emerald markers dilated from $\mathbf{f}_1^+$ and $\mathbf{f}_2^+$ are both $\frac{1}{5}\left(2+\frac{1}{4}\right)=0.45$, and the a-coordinates of the points corresponding to $\mathbf{f}_3^+$ and $\mathbf{f}_4^+$ are both $\frac{1}{5}\left(1+\frac{1}{4}\right)=0.25$, cf. (3.11) and (3.12), whereas any point along $\sigma\cap S_E$ has a-coordinate $\frac{1}{5}+\frac{1}{5}=0.4$ (σ is parametrised by weights on the edges of Σ and along S_E, the sum of the weights on edges adjacent to the same violet vertex has to be $1/5$). This and a similar computation of b-coordinates show that σ intersects the square formed by the four feelers, viewed as a subset of $S_E\subset Q_G$, at the midpoint of the southeast quarter square. (The intersection point is interior to σ because its barycentric coordinates are $1/10$ on the two edges connecting c and d and $1/5$ on the other four edges.) In particular, as the triangle spanned by the two feelers adjacent to $\mathbf{f}_4^+$ is disjoint from σ, we need $\mathbf{f}_2^+$ (i.e., an application of Lemma 4.5) as well to demonstrate that σ is an interior simplex.

EXAMPLE 5.3. Let G be the complete bipartite graph with colour classes of size $m+1$ and $n+1$, respectively. It is not hard to verify (see [**5**, Example 7.2]) that in the interior polynomial $I(\xi)$ of either hypergraph that G induces, the coefficient of ξ^k is $\binom{m}{k}\binom{n}{k}$. On the other hand, the root polytope Q_G is the product of an m- and an n-dimensional unit simplex, so that the number of lattice points in $s\cdot Q_G$ is $\binom{s+m}{m}\binom{s+n}{n}$ and thus the formula (1.3) of Saalschütz follows from (1.2) and (5.1).

It is in fact possible to give a proof of the classical Saalschütz formula using the point of view of this paper but without relying on the interior polynomial and Theorem 1.1. Let us consider the following concrete triangulation of Q_G [**3**]. We denote the colour classes of G by $E=\{\,e_0,e_1,\dots,e_m\,\}$ and $V=\{\,v_0,v_1,\dots,v_n\,\}$. We fix two horizontal lines on the plane and write the symbols for the elements of E and V, respectively, on them from left to right in the indicated order. We call a subgraph of G *non-crossing* if the straight line segments in our figure that represent its edges only intersect at endpoints. The maximal simplices that correspond to two non-crossing spanning trees share a common facet if and only if the two trees differ by a single (obviously defined) 'N $\leftrightarrow$ И transition.' A non-crossing spanning tree can be uniquely described by a *zigzag*, that is a non-crossing path in G containing the leftmost edge e_0v_0 and the rightmost edge e_mv_n. Here the zigzag is a subgraph of the corresponding tree and the degree 2 vertices of the zigzag are exactly the degree ≥ 2 vertices of the tree.

It turns out [**3**] that the collection of simplices in $Q_G=\Delta_E\times\Delta_V$ that correspond to non-crossing spanning trees is a shellable triangulation. For example, it is not hard to verify that the hypertree order (2.2) induces a shelling order. It is also not hard to show that in order for a simplex to have k adjacent (through a facet) simplices that are smaller in the shelling order, its zigzag has to have exactly k degree 2 vertices among $e_1,\dots,e_m$. Once we have fixed those, the zigzag will be uniquely determined by the choice of its k degree 2 vertices among $v_0,\dots,v_{n-1}$. Hence we see that there are $\binom{m}{k}\binom{n}{k}$ such trees and therefore (1.3) follows by the argument outlined in Remark 3.12.

The version (1.4) of Saalschütz's identity can be derived in essentially the same way, by counting interior lattice points in $(q+1)Q_G$ with the help of the same triangulation as above, using the process explained at the beginning of this section.

The following obvious corollary of Theorem 1.1 was stated as Conjecture 7.1 in [**5**] and was also mentioned in [**4**, **6**].

Corollary 5.4. Any pair $\mathcal{H}, \overline{\mathcal{H}}$ of abstract dual (transpose) hypergraphs satisfies

$$I_{\mathcal{H}}(x) = I_{\overline{\mathcal{H}}}(x).$$

Now that we know that the interior polynomial is not just an invariant of a hypergraph but in fact an invariant of the underlying bipartite graph, it becomes an interesting problem to express individual coefficients directly in terms of the graph (i.e., without either breaking the symmetry of the colour classes or using the root polytope). As we recalled (and re-proved) in Remark 3.11, such formulas are known [**5**] for the constant term (which is always 1) and the linear coefficient (which is the first Betti number). Here we present one for the quadratic coefficient.

Proposition 5.5. *Let the connected bipartite graph G have first Betti number (nullity) b_1 and let it have N cycles of length four. Then the quadratic coefficient a_2 in the common interior polynomial of the hypergraphs induced by G is $\binom{b_1+1}{2} - N$.*

Proof. Let the colour classes of G be E and V. The formula (3.7) gives $a_2 = \binom{d+1}{2} - df_0 + f_1$, where $d = |E| + |V| - 2$ is the dimension of the root polytope Q_G, f_0 is the number of edges in G, and f_1 is the number of 1-dimensional simplices in an arbitrary triangulation of Q_G. Since $b_1 = f_0 - d - 1$, it suffices to show that $N = \binom{f_0}{2} - f_1$. In other words, we wish to prove that in any triangulation $\mathcal{T}$ of Q_G, the set of pairs of vertices that are not connected by an edge in $\mathcal{T}$ is in a one-to-one correspondence with the set of four-cycles of G.

Any pair ε, δ of edges of G is disjoint from $|E| + |V| - 4$ (or $|E| + |V| - 3$, if ε and δ share an endpoint) star-cuts. The supporting hyperplanes of these cuts (as in Lemma 3.5, but this time let us take their intersections with the affine hull of Q_G) intersect in an affine subspace[†], which in turn intersects Q_G in the convex hull of 2, 3, or 4 vertices. Indeed, if ε and δ are adjacent then there are no other edges connecting their endpoints. Otherwise, there may be one or two such additional edges, the latter case being when ε and δ are opposite edges along a four-cycle. (As G is bipartite and does not have multiple edges, there may be at most one such four-cycle.) Therefore the segment in Q_G determined by ε and δ is either an edge of Q_G and hence part of $\mathcal{T}$, or a diagonal of a quadrilateral (in fact, square) face of Q_G. Since $\mathcal{T}$ induces triangulations of all faces, exactly one of the two diagonals will be a one-simplex in $\mathcal{T}$. Finally, square faces of Q_G are in a bijection with four-cycles of G: we saw how a four-cycle gives rise to a square face and conversely, the four vertices of a square are affinely dependent, i.e., the four corresponding edges form a cycle. □

Example 5.6. Regarding our two running examples, $K_{2,3}$ has nullity $b_1 = 2$ with 3 four-cycles, implying $a_2 = 0$. The graph of Figure 1 also has $N = 3$ but with $b_1 = 4$, yielding $a_2 = 7$. These values match the results of Examples 2.5 and 2.6. The complete bipartite graph on $(m+1) + (n+1)$ vertices has $b_1 = mn$ and $N = \binom{m+1}{2}\binom{n+1}{2}$, yielding $\binom{b_1+1}{2} - N = \binom{mn+1}{2} - \binom{m+1}{2}\binom{n+1}{2} = \binom{m}{2}\binom{n}{2}$, which of course coincides with the result of [**5**, Example 7.2].

[†] For some very small graphs we may be talking about an intersection of zero hyperplanes, which we take to be the affine hull of Q_G.

Corollary 5.4 implies a new formula for $T_G(x,1)$ of an ordinary (not necessarily bipartite) graph $G=(V,E)$. This can be used, for instance, to write the so-called reliability polynomial of an arbitrary connected plane graph as a generating function of activities associated to its regions. To set up the formula, enumerate all possible valence distributions $\mathbf{h}\colon V\to\mathbf{N}$ of spanning trees in the graph G' that is obtained from G by adding a new vertex in the middle of each edge. (Say that we subtract 1 from the actual valence at each $v\in V$, although this is not important right now.) All spanning trees of G' are obtained from spanning trees of G by adding one of the two halves for each of the $b_1(G)$ external edges. Despite this multitude of choices, the set $B_{(E,V)}$ of our valence distributions ends up being equinumerous with the set $B_{(V,E)}$ of spanning trees.

Corollary 5.7. Order the set V arbitrarily and compute the internal activity $\iota(\mathbf{h})$ for each $\mathbf{h}\in B_{(E,V)}$ as in Definition 2.3. Then we have $T_G(x,1)=\sum_{\mathbf{h}} x^{\iota(\mathbf{h})-1}$.

In spite of the fact that the case of a plane bipartite graph is often simpler than the general one, due mainly to the presence of the dual graph of the planar embedding which is naturally directed and balanced, the following two corollaries eluded proof until now. (Plane bipartite graphs also form triples called trinities. See the seminal paper [**14**] by Tutte that contains the proof of the so-called Tree Trinity Theorem, or [**5**, Sections 8-10].) The first one was predicted in [**5**, Subsection 10.3], where it is also explained why the planar duality formula [**5**, Theorem 8.3] and Corollary 5.4 suffice to establish it.

Corollary 5.8. The interior and exterior polynomials of the six hypergraphs induced by the plane bipartite graphs in the same trinity altogether form a three-element set.

Shapiro and the second author defined so-called parking functions for an arbitrary directed graph [**10**]. These objects have a natural enumerator p which is a one-variable polynomial with positive integer coefficients. Such a parking function enumerator can in particular be associated to the dual graph of a plane bipartite graph (see, e.g., [**6**] for details). Combining Theorem 1.1 with [**6**, Corollary 1.5] yields our last result. In the case when one of the hypergraphs induced by the bipartite graph is in fact a graph, this easily follows from a formula of Merino [**7**] via the connection noted in [**10**] between parking functions and the chip firing game (a.k.a. abelian sandpile model). The general case is however new.

Corollary 5.9. Let G be a connected plane bipartite graph so that the common interior polynomial of its induced hypergraphs is I. Then $I=p$, where p is the parking function enumerator associated to the dual graph G^*.

References

1. G. E. Andrews, 'Identities in combinatorics, I: On sorting two ordered sets', *Discr. Math.* 11 (1975), 97–106.

2. A. Cameron and A. Fink, 'A lattice point counting generalisation of the Tutte polynomial', arXiv:1604.00962, to appear in *Discrete Math. Theor. Comput. Sci.*

3. I. M. Gelfand, M. I. Graev, and A. Postnikov, 'Combinatorics of hypergeometric functions associated with positive roots', in "Arnold–Gelfand Mathematical Seminars: Geometry and Singularity Theory," Birkhäuser, Boston, 1996, 205–221.

4. A. Juhász, T. Kálmán, and J. Rasmussen, 'Sutured Floer homology and hypergraphs', *Math. Res. Lett.* 19 (2012), no. 6, 1309–1328.

5. T. Kálmán, 'A version of Tutte's polynomial for hypergraphs', *Adv. Math.* 244 (2013), no. 10, 823–873.

6. T. Kálmán and H. Murakami, 'Root polytopes, parking functions, and the HOMFLY polynomial', arXiv:1305.4925, to appear in Quantum Topology.
7. C. Merino, 'Chip firing and the Tutte polynomial', *Ann. Comb.* 1 (1997), no. 3, 253–259.
8. H. R. Morton, 'Seifert circles and knot polynomials', *Math. Proc. Camb. Phil. Soc.* 99 (1986), 107–109.
9. A. Postnikov, 'Permutohedra, Associahedra, and Beyond', *Int. Math. Res. Not.* 2009, no. 6, 1026–1106.
10. A. Postnikov and B. Shapiro, 'Trees, parking functions, syzygies, and deformations of monomial ideals', *Trans. Amer. Math. Soc.* 356 (2004), no. 8, 3109–3142.
11. F. Santos, 'The Cayley trick and triangulations of products of simplices', in "Integer Points in Polyhedra – Geometry, Number Theory, Algebra, Optimization," A. Barvinok, M. Beck, C. Haase, B. Reznick, and V. Welker (eds.), *Contemp. Math.* 374, AMS, Providence, 2005.
12. A. Schrijver, 'Combinatorial optimization. Polyhedra and efficiency.' *Algorithms and Combinatorics* 24, Springer-Verlag, Berlin, 2003.
13. E. Swartz, 'From polytopes to enumeration', `http://www.math.cornell.edu/~ebs/math455.pdf`.
14. W. T. Tutte, 'The dissection of equilateral triangles into equilateral triangles', *Proc. Cambridge Philos. Soc.* 44 (1948), 463–482.

Tamás Kálmán
Tokyo Institute of Technology
2-12-1 Ookayama, Meguro-ku
Tokyo 152-8551
Japan

kalman@math.titech.ac.jp

Alexander Postnikov
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139-4307
U.S.A.

apost@math.mit.edu