EGERVÁRY RESEARCH GROUP
ON COMBINATORIAL OPTIMIZATION

TECHNICAL REPORTS

TR-2018-16. Published by the Egerváry Research Group, Pázmány P. sétány 1/C, H–1117, Budapest, Hungary. Web site: www.cs.elte.hu/egres . ISSN 1587–4451.

Hypergraph polynomials and the Bernardi process

Tamás Kálmán and Lilla Tóthmérész

October 2018

Hypergraph polynomials and the Bernardi process

Tamás Kálmán* ** and Lilla Tóthmérész *** ‡

Abstract

Recently O. Bernardi gave a formula for the Tutte polynomial $T(x, y)$ of a graph, based on spanning trees and activities just like the original definition, but using a fixed ribbon structure to order the set of edges in a different way for each tree. The interior polynomial I is a generalization of $T(x, 1)$ to hypergraphs. We supply a Bernardi-type description of I using a ribbon structure on the underlying bipartite graph G. Our formula works because it is determined by the Ehrhart polynomial of the root polytope of G in the same way as I is. To prove this we interpret the Bernardi process as a way of dissecting the root polytope into simplices, along with a shelling order. We also show that our generalized Bernardi process is a common extension of bijections (and their inverses) constructed by Baker and Wang between spanning trees and break divisors.

1 Introduction

A few years ago, the first named author of this paper introduced a pair of polynomial invariants of hypergraphs, which generalize the valuations $T(x, 1)$ and $T(1, y)$ of the two-variable graph invariant $T(x, y)$ due to Tutte [14]. They are called the interior and exterior polynomials because they are generating functions of 'interior and exterior activity.' In the case of graphs, activities were associated to spanning trees by Tutte himself in his original definiton of T. In the hypergraph case, instead of spanning trees one considers 'hypertrees' and their activities, in a spirit very close to [14]. Hypertrees were introduced in [12] (and so named in [8]). They generalize characteristic vectors of spanning trees of a graph, preserving some nice polyhedral properties.

Both for graphs and hypergraphs, the computation of individual activities requires fixing an order of the set of edges or hyperedges, respectively, albeit temporarily, because the aggregate polynomials do not depend on it. In his remarkable paper [3], O. Bernardi removed the fixed order from the definition (in the case of graphs)

*Tokyo Institute of Technology

**TK was supported by consecutive Japan Society for the Promotion of Science (JSPS) Grant-in-Aids, for Young Scientists B (no. 25800037) and for Scientific Research C (no. 17K05244).

***MTA-ELTE Egerváry Research Group and Cornell University

‡LT was supported by the National Research, Development and Innovation Office – NKFIH, grant no. 109240, and by the NSF grant DMS-14552

and replaced it with another kind of auxiliary data: a ribbon structure and a base point. Loosely speaking, for a given spanning tree, he traced the boundary of the neighborhood of the tree and numbered the edges of the graph along the way. He used this order to compute the internal and external activities of the tree. He repeated this for all spanning trees and organized the information in a two-variable generating function which happens to satisfy the same deletion-contraction formulas as the Tutte polynomial — hence the two agree, regardless of what ribbon structure we use.

In this paper we extend Bernardi's work to the case of the interior polynomial of a hypergraph. A similar formula is conjectured for the exterior polynomial.

Any hypergraph $\mathcal{H} = (V, E)$ naturally yields a bipartite graph $\operatorname{Bip}\mathcal{H}$ in which one color class corresponds to the vertices of the hypergraph, the other color class to the hyperedges, and edges correspond to containment. We assume $\operatorname{Bip}\mathcal{H}$ to be connected and endow it with a ribbon structure and a base point. (The base point can be thought of as a boundary point of the thickened graph.)

A hypertree is essentially a 'possible degree distribution vector' of a spanning tree of $\operatorname{Bip}\mathcal{H}$ taken at the elements of E, cf. Definition 2.2. We note that hypertrees of ordinary graphs are exactly the characteristic functions of their spanning trees.

Our first order of business is to define what it means to 'trace the boundary of the neighborhood of a hypertree,' which turns out to be a process constructing a certain spanning tree that realizes the hypertree. In fact we define two versions of such a 'Bernardi process.' Contrary to the case of graphs, the fact that the Bernardi process results in a tree is not trivial at all. As a byproduct, we also obtain an order on the set of edges of $\operatorname{Bip}\mathcal{H}$ which we then use to order E as well. Now it makes sense to take the interior and exterior activities, just like in [8], of the hypertree with respect to this order. After repeating the procedure for all hypertrees, we write two one-variable generating functions $\tilde{I}$ and $\tilde{X}$ for the two 'embedding activities.' (As to why not a single, two-variable function, see [8, Remark 5.7], cf. [4].)

The main result of the paper is that the generating function $\tilde{I}$ of interior embedding activities coincides with the interior polynomial. The interior and exterior polynomials of a hypergraph, I and X, look similar to each other but their behavior is rather different. For instance, the former is invariant under taking the transpose of the hypergraph but the latter is not. (Here the transpose of the hypergraph $\mathcal{H} = (V, E)$ is the hypergraph $\overline{\mathcal{H}} = (E, V)$ with the roles of vertices and hyperedges interchanged.) In other words, I is an invariant of the bipartite graph $\operatorname{Bip}\mathcal{H}$. This fact is proven in [10] by noting that (essentially) the same polynomial may be obtained as the Ehrhart polynomial of the so called root polytope of $\operatorname{Bip}\mathcal{H}$. This depends, among other things, on the basic observation that spanning trees of a bipartite graph correspond to maximal simplices in its root polytope. We exploit the same connection to prove our main theorem. Since we do not have an analogous description for the exterior polynomial, the exterior version of our result remains, for the time being, a conjecture.

The notion of root polytope (Definition 3.1) is due to Postnikov [12]. In particular, he studied its triangulations to great effect. A consequence of our proof is an unexpected link between Bernardi's and Postnikov's work: when the Bernardi trees for all hypertrees are translated to simplices in the root polytope, they form a dissection. (I.e., the simplices fill the polytope and their interiors are mutually disjoint. They typically do not form a triangulation though, cf. Examples 9.1 and 9.2 — that is, some pairs of simplices may not intersect in a common face.) In Section 9 we will see that this generalizes the well-known triangulation of the product of two simplices by non-crossing trees.

We get an alternative description of the dissection by reinterpreting Bernardi trees as 'Jaeger trees,' in honor of F. Jaeger's beautiful paper [7] in which they appear as the main terms in a certain expansion of the Homfly polynomial. (The overlap between Jaeger's cases and ours is when Bip $\mathcal{H}$ is embedded in the plane so that the so-called median construction can be performed, resulting in a ribbon structure for the graph, as well as in an alternating link. See Figures 2 and 7.) This description does not refer to hypertrees, instead it defines Jaeger trees by a local rule that is obeyed when we trace the boundaries of their neighborhoods. This leads to the definition of a natural order among Bernardi/Jaeger trees and we prove, as a key step to our main theorem, that this order is a shelling order of the dissection.

If $\mathbf{f}\colon E \to \mathbf{Z}_{\geq 0}$ is a hypertree in $\mathcal{H}$, we may view the 'other' degree vector of its Bernardi tree as a hypertree $\bar{\mathbf{f}}\colon V \to \mathbf{Z}_{\geq 0}$ in $\overline{\mathcal{H}}$. The dissection property, just like in [12], implies that this is a bijection between the two sets of hypertrees. In the case of graphs, hypertrees on the vertex set are easily seen to be the duals of the so-called break divisors. In this special case, the above bijection-by-dissection agrees with the bijection between spanning trees and break divisors, defined by Bernardi [3] and studied further by Baker and Wang [1]. In [1] the inverse of Bernardi's bijection is described in a way that is formally different from the original Bernardi algorithm. In hypergraph language (where the transpose is an obvious involution and transpose hypergraphs share the same bipartite graph and root polytope), the bijection and its inverse are revealed to be of the same nature, defined by the same dissection.

Finally, we note that our family of dissections of the root polytope (depending on ribbon structure and base point) contains several previously known triangulations. Namely, in addition to the triangulation by non-crossing trees [6] (which applies to the root polytope of a complete bipartite graph), the triangulation by duals of arborescences [9] (which works in the case of a plane bipartite graph) is also a special case.

We should warn the reader that there are many orders in this paper, of edges, nodes, hyperedges etc. induced by spanning trees, hypertrees etc. This can be cumbersome but we need each for its own technical reason. All of these orders, however, are defined by the same simple principle: some process propagates through the graph and objects are listed in the order in which they are first reached by the process.

The paper is organized as follows. In Section 2 we set our definitions and summarize some of the necessary background, including Bernardi's embedding activities. Section 3 surveys Postnikov's work and describes the connection between h-vectors of shellable dissections of the root polytope of Bip $\mathcal{H}$ and the interior polynomial of $\mathcal{H}$. We define the Bernardi process for hypergraphs in Section 4, and establish its well-definedness. In Section 5 we give the Bernardi-type description of the interior polynomial and state the equivalence of the two definitions (Theorem 5.4, our main result), as well as several conjectures. In Section 6 we define Jaeger trees, prove their basic properties and show that the set of spanning trees arising as outcomes of the Bernardi process is exactly the set of Jaeger trees. We also discuss the connection of our work to that of Baker and Wang [1]. In Section 7 we show that Jaeger trees induce a shellable dissection of the root polytope of Bip $\mathcal{H}$ with a natural shelling order, and we relate the resulting h-vector to Bernardi-type activities. This allows us to prove Theorem 5.4. In Section 8 we observe that for graphs (i.e., hypergraphs $\mathcal{H}$ where the cardinality of each hyperedge is two), the Bernardi process behaves in a special way in that it induces an activity-preserving bijection between the hypertrees of $\mathcal{H}$ and $\overline{\mathcal{H}}$. Finally, in Section 9 we show how certain previously known triangulations arise from the Bernardi process.

2 Preliminaries

2.1 Basics

A *hypergraph* is an ordered pair $\mathcal{H} = (V, E)$, where V is a finite set and E is a finite multiset of non-empty subsets of V. We refer to elements of V as *vertices* and to elements of E as *hyperedges*. For a hypergraph $\mathcal{H} = (V, E)$, let the *underlying bipartite graph* Bip $\mathcal{H}$ be the bipartite graph with vertex classes V and E, where $v \in V$ is connected to $e \in E$ if $v \in e$ in $\mathcal{H}$. In the context of bipartite graphs such as Bip $\mathcal{H}$, instead of vertices, we will call the elements of $V \cup E$ *nodes*. Specifically, the elements of V and E will be called *violet* and *emerald* nodes, respectively. (In our figures, violet appears as blue and emerald appears as green.) Throughout the paper, we assume that $\mathcal{H}$ is *connected*, which means that Bip $\mathcal{H}$ is connected.

We also assume that Bip $\mathcal{H}$ has a ribbon graph structure. Here for a graph G without loop edges, a *ribbon structure* is a family of cyclic permutations: namely for each vertex x of G, a cyclic permutation of the edges incident to x is given. For an edge xy of G, we use the following notations:

- yx^+_G: the edge following yx at y
- yx^-_G: the edge preceding yx at y
- xy^+_G: the edge following xy at x
- xy^-_G: the edge preceding xy at x.

If the graph G is clear from the context, we omit the subscript. We will sometimes need the operation of removing an edge from a ribbon graph. If G is a ribbon graph, and ε is an edge of G, then $G - \varepsilon$ is the ribbon graph with

$$xy^+_{G-\varepsilon} = \begin{cases} xy^+_G & \text{if } xy^+_G \neq \varepsilon \\ (xy^+_G)^+_G & \text{if } xy^+_G = \varepsilon \end{cases}$$

for any edge xy of $G-\varepsilon$. More generally, any subgraph of G inherits a ribbon structure from G in the obvious way, by restrictions of the cyclic orders.

Throughout the paper, when we consider ribbon structures, we will assume that there is a fixed vertex b_0 of G that we call the *base vertex* (or *base node*, in cases when G is bipartite), and a fixed edge b_0b_1 incident to b_0 that we call the *base edge*.

Remark 2.1. Ribbon structures may be equivalently described by *ribbon surfaces*, as follows. See Figure 2 for an example. We consider the graph as a topological space, thicken a small neighborhood of each vertex to a disk, and orient it (hence also orient its boundary) so that the edges incident to the vertex intersect the boundary of the disk in their prescribed cyclic order. (When the vertex has degree three or more, this orientation is uniquely determined once the disk has been constructed.) Then we thicken each edge into a rectangle, attached along two opposite sides to the appropriate disks, so that the orientations extend over the rectangle. Thinking of the two attaching sides of the rectangle as 'short' and the other two, running along the edge, as 'long,' explains the name of the structure. Conversely, if a graph is embedded in an oriented surface, the orientation induces a ribbon structure on it which is equivalent to taking a regular neighborhood of the embedding. Finally, one may equivalently specify the base node and the base edge by placing a *base point* on the boundary of the disk centered at b_0, along the bit running from $b_0b_1^-$ to b_0b_1.

Let G be a graph, T be a spanning tree of G, and $\varepsilon \in T$ be an edge. As T is a spanning tree, $T - \varepsilon$ is a graph with two connected components. We call the set of edges of G connecting two vertices from different components of $T - \varepsilon$ the *base cut* of ε in T, and denote it by $C^*(T, \varepsilon)$. Now for an edge ε of G that is not part of T, adding ε to T creates a unique cycle, which we call the *base cycle* of ε with respect to T and denote with $C(T, \varepsilon)$.

2.2 The interior polynomial

The following definition plays a central role in our paper.

Definition 2.2. Let G be a bipartite graph and E one of its vertex classes. We say that the vector $\mathbf{f}\colon E \to \mathbf{Z}_{\geq 0}$ is a *hypertree* on E if there exists a spanning tree T of G that has degree $d_T(e) = \mathbf{f}(e) + 1$ at each node $e \in E$. We denote the set of all hypertrees on E by B_E.

Disconnected bipartite graphs have no spanning trees and thus no hypertrees, either. It is slightly more natural to call the objects above hypertrees in the hypergraph $\mathcal{H} = (V, E)$, as opposed to in $G = \operatorname{Bip}\mathcal{H}$, and to denote their set with $B_{\mathcal{H}}$. In that sense, hypertrees generalize (characteristic vectors of) spanning trees from graphs to hypergraphs (cf. [8, Remark 3.2]). We will often adopt this point of view, even though the wording of Definition 2.2 suits our current purposes better.

The set B_E is such that $(\operatorname{Conv} B_E) \cap \mathbf{Z}^E = B_E$, where Conv denotes the usual convex hull in $\mathbf{R}^E$, cf. [8, Lemma 3.4]. We will call $\operatorname{Conv} B_E = \operatorname{Conv} B_{\mathcal{H}}$ the *hypertree polytope* of $\mathcal{H}$. In the special case of (characteristic vectors of) spanning trees, hypertrees are exactly the vertices of $\operatorname{Conv} B_E$, which in that case is known as the spanning tree polytope.

We note that for all hypertrees $\mathbf{f}$ on E, it holds that

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

is independent of $\mathbf{f}$ [8, Theorem 3.4]. Hence B_E and $\operatorname{Conv} B_E$ lie along an affine hyperplane of $\mathbf{R}^E$.

For a spanning tree T of $\operatorname{Bip}\mathcal{H}$, let $\mathbf{f}_E(T)$ be the hypertree on E *realized* or *induced* by T, i.e.,

$$\mathbf{f}_E(T)(x) = d_T(x) - 1 \quad \text{for all } x \in E.$$

Similarly, let $\mathbf{f}_V(T)$ be the hypertree on V realized by T.

The definition of the interior polynomial is based on hypertrees and a natural generalization of internal activity used in the case of graphs (and matroids). First, if the hypertree $\mathbf{f} \in B_E$ 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 let us say that $\mathbf{f}$ and $\mathbf{f}'$ are related by a *transfer of valence* from e to e'. Another expression we will use is that $\mathbf{f}$ is such that e *can transfer valence* to e'.

Definition 2.3. Let (V, E) be a hypergraph and let us order the set E arbitrarily. A hyperedge $e \in E$ is *internally active* with respect to the hypertree $\mathbf{f}$ if $\mathbf{f}$ is such that e cannot transfer valence to any smaller hyperedge. Let $\iota(\mathbf{f})$ denote the number of internally active hyperedges with respect to $\mathbf{f}$ and call this value the *internal activity* of $\mathbf{f}$.

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 $\mathbf{f}$. Note that $\iota(\mathbf{f})$ and $\bar{\iota}(\mathbf{f})$ depend on the order used on E.

Definition 2.4. Let $\mathcal{H} = (V, E)$ be a hypergraph so that $\operatorname{Bip}\mathcal{H}$ is connected. For some fixed linear order on E we consider the generating function of internal inactivity, $I_{\mathcal{H}}(\xi) = \sum_{\mathbf{f} \in B_E} \xi^{\bar{\iota}(\mathbf{f})}$, and call it the *interior polynomial* of $\mathcal{H}$. By [8, Theorem 5.4] (see also [10, subsection 2.2]), $I_{\mathcal{H}}$ does not depend on the order.

Example 2.5. Any connected bipartite graph serves as the underlying bipartite graph for two hypergraphs, a transpose pair. For the graph that appears throughout the paper (see Figures 2, 4, 5, 7, 9), both of these hypergraphs have the interior polynomial $1 + 3\xi + 3\xi^2$, as computed in [8, Example 5.6]. In particular, the number of hypertrees on either vertex class is 7.

Note that if we specialize our notion of internal activity to graphs, then external edges of a spanning tree become internally active. This is not the case for the notion used by Tutte and Bernardi, which we review in the next subsection. However, the number of internally inactive edges is the same as the number of 'internal, not internally active' edges in the original definition. This subtlety is hard to avoid because for a hyperedge, there is no natural notion of being inside or outside of a hypertree. Even if we defined 'e being external to $\mathbf{f}$' by $\mathbf{f}(e) = 0$, the number of such hyperedges *would* depend on $\mathbf{f}$.

Also, as opposed to the definition of the Tutte polynomial, in Definition 2.4 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 $\beta_1 = |E| - |V| + 1$ of

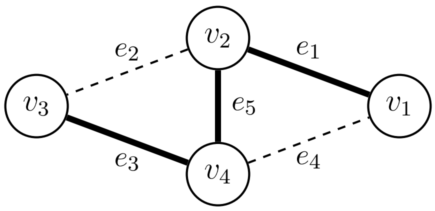

Figure 1: An example of the tour of a spanning tree. Let the ribbon structure be the one induced by the positive orientation of the plane. The edges of the tree are drawn by thick lines, the non-edges by dashed lines. With $b_0 = v_1, b_1 = v_2$, we get the tour $v_1e_1, v_2e_2, v_2e_5, v_4e_3, v_3e_2, v_3e_3, v_4e_4, v_4e_5, v_2e_1, v_1e_4$.

the graph), all these tweaks in the definition just mirror and shift the distribution of the 'classical' internal activity statistic. The precise claim is that if the hypergraph $\mathcal{H}$ happens to be a graph with Tutte polynomial $T(x, y)$, then its interior polynomial is

$$I_{\mathcal{H}}(\xi) = \xi^{|V|-1}T(1/\xi, 1).$$

We will almost always work in the larger context of hypergraphs and use the notions of Definitions 2.3 and 2.4. Because the difference is so minimal, we will not introduce separate terminology, only separate notation for internal activity in the 'usual' sense: for a spanning tree T of the graph G with vertex set V, and an ordering of the set E of edges, we let $i(T) = \iota(T) - \beta_1(G) = |V| - 1 - \bar{\iota}(T)$.

Hypergraphs also have an exterior polynomial invariant. For this paper it is less important but we will indicate its definition and various properties in Section 5.

2.3 The Bernardi process for graphs

In this subsection we recall the Bernardi process for ordinary graphs, as well as the Bernardi-type definition of the Tutte polynomial [2, 3].

The basic notion in [2] is the *tour of a spanning tree*. Let G be a ribbon graph, and T be a spanning tree of G. We specify a base vertex b_0 of G and a base edge b_0b_1 incident to b_0.

The tour of T is the following sequence of vertex-edge pairs: The current vertex at the first step is b_0, and the current edge is b_0b_1. If the current vertex is x, the current edge is xy, and $xy \notin T$, then the current vertex of the next step is x, and the current edge of the next step is xy^+. If the current vertex is x, the current edge is xy, and $xy \in T$, then the current vertex of the next step is y, and the current edge of the next step is yx^+. In the first case we say that the tour *skips* xy and in the second case we say that the tour *traverses* it. The tour stops right before when b_0 would once again become current vertex with b_0b_1 as current edge. See Figure 1 for an example. Bernardi proved the following:

Lemma 2.6 ([2, Lemma 5]). *In the tour of a spanning tree T, each edge xy of G becomes current edge twice, in one case with x as current vertex, and in the other case with y as current vertex.*

In other words, the tour of T lists the pairs (u, ε), where u is a vertex of G and ε is an edge of G incident to u, in a linear order from smallest to largest. We will denote this ordering by $<_T$, and write $(u, \varepsilon) \leq_T (v, \varepsilon')$ if either $u = v$ and $\varepsilon = \varepsilon'$, or $(u, \varepsilon) <_T (v, \varepsilon')$.

Now $\leq_T$ induces an ordering of the edges of G: Let uv be smaller than zw if $\min_{\leq_T}\{(u, uv), (v, uv)\} \leq \min_{\leq_T}\{(z, zw), (w, zw)\}$. We denote this order also by $\leq_T$.

The *internal and external embedding activities* of a spanning tree T of G are defined as the internal and external activities of T with respect to the order $\leq_T$ of edges. Let us denote them by $ie(T)$ and $ee(T)$, respectively. That is,

- $ie(T)$ is the number of edges ε of T so that ε is the $<_T$-minimal element of the base cut $C^*(T, \varepsilon)$
- $ee(T)$ is the number of non-edges ε of T so that ε is the $<_T$-minimal element of the base cycle $C(T, \varepsilon)$.

Bernardi gave the alternative definition

$$T_G(x, y) = \sum_{T \text{ is a spanning tree in } G} x^{ie(T)} y^{ee(T)}$$

for the Tutte polynomial of a graph. In particular, it follows from his results that this expression does not depend on the ribbon structure, base vertex, or base edge.

The tour of a spanning tree can also be used to give a bijection between spanning trees and so-called break divisors. A *break divisor* for a graph $G = (V, E)$ is a vector $\mathbf{z} \in \mathbf{Z}^V$ so that $\mathbf{d} - \mathbf{1} - \mathbf{z}$ is a hypertree on V in $\operatorname{Bip} G$, that is the graph obtained from G by adding a new vertex halfway along each edge. (Here $\mathbf{d} - \mathbf{1}$ denotes the vector whose v-component is the degree of v in G minus one for each $v \in V$). Baker and Wang [1] proved that given a spanning tree T, if one takes its tour and at each vertex, counts the non-edges of T that first become current in conjunction with that vertex, the resulting values give a break divisor. Conversely, given a break divisor $\mathbf{z}$, one may start a walk on the ribbon graph and whenever an edge is encountered which is such that $\mathbf{z}$ remains a break divisor in the smaller graph after the edge's removal, cut that edge. By the end of the walk, the remaining edges form a spanning tree.

Said in another way, the Bernardi process on a graph gives a bijection between the sets of hypertrees on V and on E. Our Bernardi process for hypergraphs generalizes this bijection (see Remark 6.22 for more).

3 The root polytope and its dissections

The root polytope of a bipartite graph was defined by Postnikov [12]. For a detailed list of its properties we refer the reader to [10, Section 3] (and to [12] for many of the proofs). Here we only repeat the most important points. Let G be a bipartite graph with color classes E (as in emerald) and V (as in violet).

Definition 3.1. In the Euclidean space $\mathbf{R}^E \oplus \mathbf{R}^V$ let us write $\mathbf{x}$ for the standard basis vector that corresponds to $x \in E \cup V$. The *root polytope* of G, denoted by Q_G, is the convex hull of the vectors $\mathbf{e} + \mathbf{v}$ for all edges ev of G.

We get an isometric polytope if we replace $\mathbf{e} + \mathbf{v}$ with $\mathbf{e} - \mathbf{v}$ in the construction. Note how the definition is inspired by the standard proof that a graphic matroid is representable. However in the theory that Postnikov built it is important that we specifically use real coefficients and examine Q_G from the point of view of convex geometry. It turns out that the root polytope reflects certain properties of the bipartite graph in a non-trivial and very effective way.

A set of vertices of Q_G is affine independent if and only if the corresponding subgraph of G is cycle-free. (Note that every vertex of Q_G corresponds to an edge of G.) In particular, the dimension of Q_G is one less than the number of edges in a spanning forest of G. In the case when G is connected, which we usually assume, a spanning forest is a spanning tree, and the dimension is $|E|+|V|-2$. The one-to-one correspondence (for connected G)

$$\{\text{ spanning trees of } G\,\} \longleftrightarrow \{\text{ maximal simplices of } Q_G\,\}$$

will be crucial for the rest of the paper. Note that the simplex corresponding to the tree T is none other than its root polytope Q_T.

The maximal simplices of Q_G share the same volume. There is a description [12, Lemma 12.6] for when two maximal simplices intersect in a common face, given in terms of the corresponding spanning trees: two trees are compatible in this sense if and only if there does not exist a cycle in G so that its first, third, fifth etc. edges come from one tree and its second, fourth, sixth etc. edges come from the other. A *triangulation* of the root polytope is a collection of maximal simplices whose union is Q_G and each pair of which do intersect in a common face. When the second condition is weakened to require only that the interiors of the simplices be disjoint, we get the notion of a *dissection*. The observation on volumes implies that the number of maximal simplices in a dissection of Q_G depends only on G. In Example 9.1 we show a dissection that is not a triangulation.

A thorough look into dissections reveals some spectacular properties. Let us fix a dissection of Q_G and consider the corresponding collection of spanning trees in G. We claim that any hypertree (either on E or on V) is realized by exactly one of our chosen trees. (Consequently the numbers of hypertrees on E and on V are the same.) This is proved for triangulations in [12], but the same proof applies in general, as follows.

The polytope Q_G contains the set of *emerald markers* (which form an affine transformation of the set B_E of hypertrees on E)

$$\frac{1}{|V|}B_E + \frac{1}{|E|\cdot|V|}\mathbf{i}_E + \frac{1}{|V|}\mathbf{i}_V$$

and a similarly defined set of *violet markers*. Here $\mathbf{i}_S$ stands for the characteristic function of a subset S of $E \cup V$, viewed as a vector in $\mathbf{R}^E \oplus \mathbf{R}^V$. Any maximal simplex in Q_G contains, in its interior, exactly one marker of each color: these are (essentially) the hypertrees on E and on V realized by the tree that corresponds to the simplex [12, Lemma 14.9]. If some simplices form a dissection, then it is also true that each marker is contained by a unique simplex. Hence the two sets of markers (that is, the two sets of hypertrees B_E and B_V) are equinumerous with each other and with the set of maximal simplices in our (arbitrary) dissection. In particular, the following holds.

Theorem 3.2. *Let $\mathcal{T}$ be a set of spanning trees of a bipartite graph G such that the simplices $\{\, Q_T \mid T \in \mathcal{T} \,\}$ form a dissection of Q_G. Then the mapping assigning $\mathbf{f}_V(T)$ to $\mathbf{f}_E(T)$ for each $T \in \mathcal{T}$ is a bijection between B_E and B_V.*

By [10, Theorem 3.10], the h-vector of any triangulation of Q_G is equivalent to the Ehrhart polynomial of Q_G. This property, too, extends to dissections, at least in those cases when we are able to define an h-vector for such objects. In other words, one is able to generalize [10, Remark 3.12] as follows.

Let us call a dissection *shellable* if it has a *shelling order*, that is, a total order of the maximal simplices so that, starting from the second one, each maximal simplex σ intersects the union of the previous ones in a non-empty union of facets (codimension one faces) of σ. For a dissection with a shelling order, let a_i denote the number of maximal simplices for which the number of said facets is exactly i. We put $a_0 = 1$, accounting for the first simplex in the order. Let us define the *h-vector* of the shelling order to be the finite sequence $(a_0, a_1, \ldots)$.

The Ehrhart polynomial of Q_G is the unique polynomial ε_G such that for nonnegative integers k, we have $\varepsilon_G(k) = |(k \cdot Q_G) \cap (\mathbf{Z}^E \oplus \mathbf{Z}^V)|$. If we have a shellable dissection of Q_G (for a connected G), then, putting $d = \dim Q_G = |E| + |V| - 2$, the Ehrhart polynomial of Q_G can be expressed as

$$\varepsilon_G(k) = a_0\binom{d+k}{d} + a_1\binom{d+k-1}{d} + \cdots + a_i\binom{d+k-i}{d} + \cdots. \tag{3.1}$$

The proof of (3.1) is based on [10, Lemma 3.3] and an easy simplex-by-simplex counting argument. Indeed, $\binom{d+k-i}{d}$ is the number of lattice points in a standard d-dimensional simplex of sidelength $k - i$, and [10, Lemma 3.3] says that for our purposes, all maximal simplices in the root polytope (and all their faces) behave just like standard simplices. Now if, in the k times inflated root polytope, the lattice points along i of the facets of a maximal simplex have already been counted, then what remains to count is the lattice points in a simplex of sidelength not k but $k - i$.

Since the binomial coefficients in (3.1) are linearly independent as polynomials of k [10, Lemma 3.8], from the uniqueness of the Ehrhart polynomial it follows that the h-vectors of all shelling orders of all shellable dissections of Q_G coincide. (Triangulations have h-vectors even when they are not shellable. For triangulations of Q_G by maximal simplices, all their h-vectors are the same, too [10, Theorem 3.10].) Furthermore, [10, eq. (5.1)] (the main theorem of that paper) states that the same coefficient sequence gives the interior polynomial of both hypergraphs (V, E) and (E, V) that are induced by G:

Theorem 3.3. *For a connected hypergraph $\mathcal{H} = (V, E)$, if the Ehrhart polynomial $\varepsilon_{\mathrm{Bip}\,\mathcal{H}}$ of $Q_{\mathrm{Bip}\,\mathcal{H}}$ can be expressed as*

$$\varepsilon_{\mathrm{Bip}\,\mathcal{H}}(k) = a_0\binom{d+k}{d} + a_1\binom{d+k-1}{d} + \cdots + a_i\binom{d+k-i}{d} + \cdots,$$

where $d = \dim Q_G = |E| + |V| - 2$, *then the interior polynomial of* $\mathcal{H}$ *is*

$$I_{\mathcal{H}}(x) = a_0 + a_1 x + \cdots + a_i x^i + \cdots .$$

Both sums in Theorem 3.3 are of course finite. The largest i so that $a_i \neq 0$ is definitely no more than $d+1$; by [8, Proposition 6.1] it is in fact at most $\min\{\,|E|, |V|\,\}-1$. For more on the degree of the interior polynomial, see [5].

K. Kato [11] found a concise formulation of Theorem 3.3, restating it as a connection between $I_{\mathcal{H}}$ and the Ehrhart series $\mathrm{Ehr}_{\mathrm{Bip}\,\mathcal{H}}(x) = \sum_{s=0}^{\infty} \varepsilon_{\mathrm{Bip}\,\mathcal{H}}(s)\, x^s$ of $Q_{\mathrm{Bip}\,\mathcal{H}}$. Namely, if $\mathrm{Bip}\,\mathcal{H}$ is connected, then we have

$$\frac{I_{\mathcal{H}}(x)}{(1-x)^{|E|+|V|-1}} = \mathrm{Ehr}_{\mathrm{Bip}\,\mathcal{H}}(x).$$

4 The Bernardi process for hypergraphs

In this section we describe two processes, both of which generalize what Bernardi defined for ordinary graphs. Sometimes, in order to distinguish them from Bernardi's original algorithm, we will refer to the next them as the *hypergraphical Bernardi processes*. In both cases, the input consists of

(a) a connected hypergraph $\mathcal{H} = (V, E)$

(b) a ribbon structure for $\mathrm{Bip}\,\mathcal{H}$, a base node, and a base edge

(c) a hypertree $\mathbf{f}$ in $\mathcal{H}$, that is, on E.

Here (a) and (c) generalize Bernardi's inputs of a connected graph and a spanning tree. There is a slight difference between the graphs G and $\mathrm{Bip}\,G$ in that the latter is obtained from the former by placing a new vertex halfway along each edge. Bernardi fixed a ribbon structure on G and not on $\mathrm{Bip}\,G$, but as the new vertices of $\mathrm{Bip}\,G$ are of degree 2, the structure extends uniquely to $\mathrm{Bip}\,G$. Therefore (b) above also generalizes what Bernardi used in his approach.

Let us first describe the Bernardi process informally, in terms of a walk on $\mathrm{Bip}\,\mathcal{H}$, traversing or cutting edges as we go. We outline two different processes, depending on whether we are allowed to cut edges at their violet or emerald endpoint. (One of the two will be allowed, the other prohibited.) If the walk reaches an edge from the endpoint where we are not allowed to cut, there is no choice but to traverse the edge and continue on the other side. In the other case we will have a choice: either traverse the edge as above, or cut it (remove it from the graph) and continue with the next edge incident to our current node. What governs this choice is whether the values of $\mathbf{f}$ still define a hypertree after cutting the edge. (That is, whether the smaller graph still has a spanning tree that realizes $\mathbf{f}$.) If they do, we cut; otherwise, we keep and traverse. (Note that this is the same principle with which Baker and Wang define their mapping from break divisors to spanning trees.)

We formalize these ideas as follows. Let the base node be b_0 and the base edge be $b_0 b_1$.

Definition 4.1 (Bernardi process, hypertree on emerald nodes, cut at violet nodes (ht:E, cut:V))**.** Given is a hypertree $\mathbf{f}$ on E. The process maintains a current edge and a current graph at any moment. At the beginning, the current graph is $\mathrm{Bip}\,\mathcal{H}$. If b_0 is a violet node, then at the beginning, the current edge is b_0b_1. If b_0 is an emerald node, then the current edge at the beginning is $b_1b_0^+$, and we say that b_0b_1 was traversed from the emerald direction.

In each step, we check whether for the current graph G and the current edge ve, the vector $\mathbf{f}$ is a hypertree on the emerald nodes of $G - ve$. If the answer is yes, let the current graph of the next step be $G - ve$ and the current edge be ve_G^+. We say that ve was *removed* or *deleted* from the graph. If the answer is no, let the current graph of the next step be G, and let the current edge of the next step be we_G^+, where $ew = ev_G^+$. In this case we say that ve is traversed from the violet direction and ew is traversed from the emerald direction.

The process stops right before when an edge would be traversed for the second time from the same direction.

Definition 4.2 (Bernardi process, hypertree on emerald nodes, cut at emerald nodes (ht:E, cut:E))**.** Given is a hypertree $\mathbf{f}$ on E. The process maintains a current edge and a current graph at any moment. At the beginning, the current graph is $\mathrm{Bip}\,\mathcal{H}$. If b_0 is an emerald node, then at the beginning, the current edge is b_0b_1. If b_0 is a violet node, then the current edge at the beginning is $b_1b_0^+$, and we say that b_0b_1 was traversed from the violet direction.

In each step, we check whether for the current graph G and the current edge ev, the vector $\mathbf{f}$ is a hypertree on the emerald nodes of $G - ev$. If the answer is yes, let the current graph of the next step be $G - ev$ and the current edge be ev_G^+. We say that ev was *removed* or *deleted* from the graph. If the answer is no, let the current graph of the next step be G, and let the current edge of the next step be dv_G^+, where $vd = ve_G^+$. In this case we say that ev is traversed from the emerald direction and vd is traversed from the violet direction.

The process stops right before when an edge would be traversed for the second time from the same direction.

In both cases the current graphs form a decreasing sequence. We say that an edge was *kept* by the process if it was examined as a current edge, and was not removed from the current graph. If an edge ε is kept then '$\mathbf{f}$ cannot be realized without it' in the current graph, that is, ε is part of any spanning tree realizing $\mathbf{f}$ in the current graph and hence in all subsequent current graphs, too. In particular, once an edge is kept, it can never get removed — the decision of keeping is final, just like the decision of removal.

We say that an edge has been *traversed* by the process if it has been traversed from the violet direction or from the emerald direction. Traversed edges form an increasing sequence of subgraphs of $\mathrm{Bip}\,\mathcal{H}$. In fact, the graph of traversed edges will always be a subgraph of the current graph, but for this we have yet to show that traversed edges never get removed (cf. Lemma 4.10). The problem of course is with the edges ew of Definition 4.1 and vd of Definition 4.2, which our walk traverses without examining them as current edges. It turns out that these are always previously unexamined edges and that later they will be examined and kept, but these facts are not obvious from the definition. For the time being let us reinforce that even if a traversed edge was later removed, we would still count it as traversed.

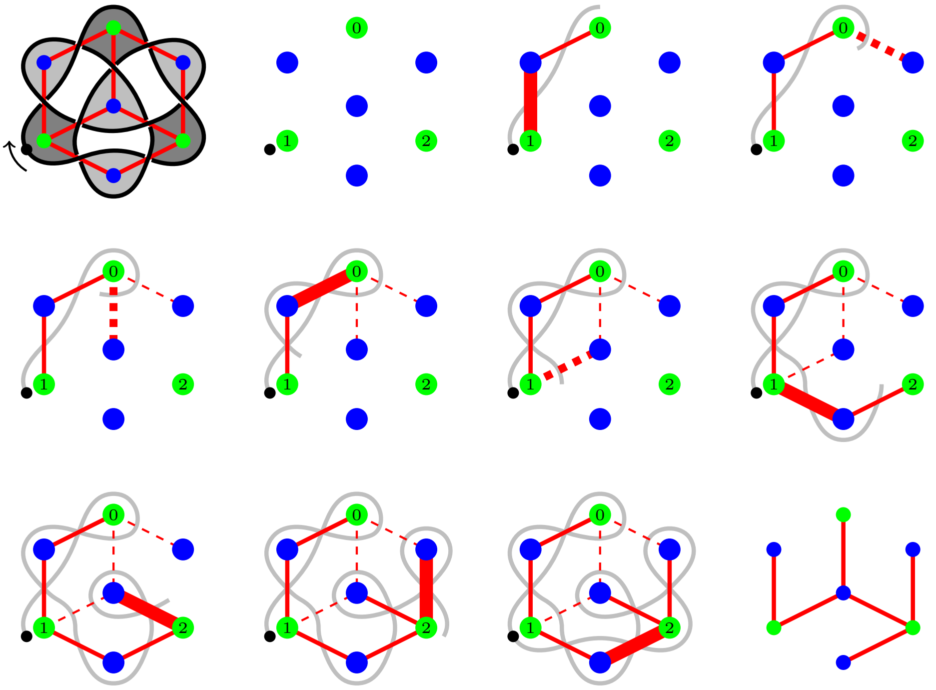

Figure 2: The Bernardi process on a hypergraph.

Example 4.3. Let G be the plane bipartite graph shown in Figure 2, with the three emerald (green) nodes forming the color class E and the four violet (blue) ones the color class V. We will refer to elements of E as top, left, and right. A ribbon structure is chosen so that at emerald nodes the cyclic order of the incident edges is clockwise, whereas at violet nodes it is counterclockwise. In the first panel of Figure 2 we show (using a particular embedding in 3-space) the associated ribbon surface that was described in Remark 2.1. Also indicated (in the lower left) is the base point; equivalently, our base node is the left emerald point and our base edge is the vertical one on the left.

The numbers $0, 1, 2$ written over the emerald nodes are a hypertree on E. In the panels two through eleven we show how the (ht:E, cut:E) Bernardi process operates with this input. At first there is just the hypertree and the process starts from the base point.

In the first step, which is probably the most interesting one in this example, the base edge is current. The hypertree calls for a realization with degree $1 + 1 = 2$ at the base node, so one might expect that the first of the three incident edges will be removed. However if we did that, then in the remaining graph, all four edges of the bottom rhombus would be 'wanted' by the hypertree: the left two to make the base node degree 2, and the right two to make the right emerald node degree $2 + 1 = 3$. In other words, in the remaining graph it would not be possible to realize the given numbers as the degrees (minus 1) of a spanning tree. Thus the Bernardi process will not remove the base edge, rather it will traverse it, which then will force it to traverse the upper left edge of G as well.

In the next eight panels we show the remaining steps of the process. The current edge of each step is highlighted. The gray curve is included to help keep track of the cyclic permutations of the ribbon structure. The decisions are rather straightforward: removal, removal, traversal (in fact for the second time; note that if we did not traverse here, the top node would become isolated so it could no longer have degree $0+1=1$), removal, traversal (of the current edge and one more edge), traversal (here and in the next step, the other edge we are forced to traverse coincides with the current edge), traversal, traversal. Note that each edge was current exactly once and that at the end, the subgraph of those edges that we did not remove coincides with the subgraph of traversed edges, and this subgraph is a spanning tree that realizes the given hypertree.

The last panel of Figure 2 shows the outcome of the (ht:E, cut:V) Bernardi process on the same input hypertree. It is again a spanning tree realization. We leave it to the interested reader to construct the steps leading to it, and to check that again there are nine such steps, with each edge becoming current exactly once. (Note that in this case the first current edge is not the base edge, rather it is the upper left one; the base edge becomes current in the second step.)

We will usually think of part (c) of the input, the hypertree, as a 'variable,' so that the process itself is determined by (a) and (b) only. In this sense we may in fact speak of four Bernardi processes on the ribbon bipartite graph $\operatorname{Bip}\mathcal{H}$, by applying the two definitions above to $\mathcal{H}$ and to $\overline{\mathcal{H}}$. In the latter case the hypertree is given on V and we denote those two processes with (ht:V, cut:E) and (ht:V, cut:V).

Many more instances of the process can be generated by varying part (b) of the input. See for example Lemma 6.5 for the case when the ribbon structure is reversed (even though that process turns out not to be completely new).

Remark 4.4. Both of our processes for hypergraphs do indeed generalize Bernardi's original process for graphs, and in fact, they can be thought of as a common generalization of the Bernardi process and its inverse, given by Baker and Wang [1]. That is, for each spanning tree T of the ribbon graph G, Bernardi's tour of T is basically equivalent to our walk on $\operatorname{Bip}G$, where the latter is defined using the uniquely extended ribbon structure as part (b) of the input and the characteristic function of T as part (c). We sketch the main ideas of an induction proof.

First note that a (current vertex, current edge) pair (x, xy) in G can be equivalently given as a 'current half-edge' xe between x and the node e placed at the center of xy. While xe is a half-edge in G, it is an edge in $\operatorname{Bip}G$. In particular, where Bernardi chose a base vertex b_0 and base edge $e_0 = b_0b_1$, we may speak instead of the base node b_0 and the base edge b_0e_0.

If xy is in the spanning tree T, then the corresponding hypertree has the value 1 at e and hence to realize it in $\operatorname{Bip}G$, both ex and ey are necessary. Therefore both versions of our process will traverse both of those edges. In the first version, having arrived at xe 'near' x, we decide to keep it and then the ribbon structure at e forces us to traverse ey, too. In the second version we have to traverse xe anyway and then at

e we make the decision to keep ey as well. After this we continue with the (half-)edge that follows ey in the cyclic order at y. In both cases this is exactly what the original process would have done, too.

If xy is not in T, that is when the hypertree takes the value 0 at e, then let us consider two sub-cases. If ey has not been cut thus far, then the first version of our process, upon arrival at xe, will cut xe because it is not necessary for a realization of the hypertree. The second version will traverse xe and then cut ey, which forces it to backtrack to x. On the other hand if ey is not in the graph any more, then the first version will decide to keep xe so that e does not become an isolated node; the second version will have to traverse xe anyway but in both versions, since e was already a leaf, the process will bounce back to x and continue with the edge that follows xe in the cyclic order at x. This again matches the behavior of the original Bernardi process.

When applied to a graph G, our versions of the Bernardi process do not only trace a given spanning tree T, they also select one of the two half-edges of each non-edge of T so as to enlarge it into a spanning tree of $\operatorname{Bip} G$. As is clear from the above, the difference between the two generalizations is whether this half-edge is opposite to (first version) or on the same side as (second version) the first endpoint of the edge that Bernardi's tour of T visits.

The main theorem of this section draws identical conclusions for the processes of Definitions 4.1 and 4.2. It is convenient to state and prove it for the process of Definition 4.1 applied to $\mathcal{H}$, coupled with the process of Definition 4.2 applied to $\overline{\mathcal{H}}$, so that in both cases we are allowed to cut edges at their violet endpoints. (We do not lose any generality because $\mathcal{H} = \overline{(\overline{\mathcal{H}})}$.)

Theorem 4.5. *For any hypertree* $\mathbf{f}$ *on E (respectively, on V), the (ht:E, cut:V) Bernardi process (respectively, the (ht:V, cut:V) Bernardi process) takes each edge of* $\operatorname{Bip} \mathcal{H}$ *exactly once as current edge. The current graph at the end of the process is a spanning tree of* $\operatorname{Bip} \mathcal{H}$ *realizing* $\mathbf{f}$.

This theorem generalizes Lemma 2.6, cf. Remark 4.4. In Section 9 we will see that it also generalizes [8, Theorem 10.1]. It is no wonder then that the proof is somewhat lengthy. The key will be Lemma 4.7, which ensures that the subgraph of traversed edges never contains a cycle.

By a *violet Bernardi run* we mean a running of either the (ht:E, cut:V) Bernardi process or the (ht:V, cut:V) Bernardi process on some hypertree.

Lemma 4.6. *In a violet Bernardi run, if until some moment there is no cycle in the subgraph of traversed edges, then until that moment,*

(i) at any node $x \in V \cup E$, the edges of the current graph G incident to x that were traversed from the direction of x are consecutive edges in the cyclic order at x (in G), covering less than one full turn, and were traversed consecutively.

(ii) at any $x \in V$, those edges of $\operatorname{Bip} \mathcal{H}$ *incident to x that have already been current edges, became current edges in a consecutive order compatible with the cyclic order at x (in* $\operatorname{Bip} \mathcal{H}$*), covering less than one full turn.*

Proof. If $x = e \in E$, then, having arrived at e on an edge ve, the Bernardi process next traverses $ev^+_{G'}$ by definition. Here G' is the current graph at the time of these traversals. If there are any edges of $\mathrm{Bip}\,\mathcal{H}$ incident to e that fall between ev and $ev^+_{G'}$ in the cyclic order (at e in $\mathrm{Bip}\,\mathcal{H}$), then they have already been cut at their violet endpoints. Hence in any later current graph, if ev and $ev^+_{G'}$ are still present in it, then they are still consecutive.

The next time the walk associated to the process arrives at e, it arrives on $ev^+_{G'}$ since otherwise it would have traversed a cycle. To get our conclusion we just have to repeat our argument and note that if the sequence of traversed edges covered a full turn, then some edge would be traversed twice from the direction of e, which would cause the process to stop.

For a vertex $v \in V$, it suffices to show (ii) since it implies (i). If the current edge ve is removed from the graph, the next current edge is ve^+_G by definition. Here G is the current graph when ve is the current edge. On the other hand, if ve is traversed, the next time the process arrives at v, it has to arrive on ev, otherwise it would traverse a cycle. But then the next current edge incident to v is ve^+_G. In this case, the G in ve^+_G a priori refers to the current graph at the moment of traversing ev, but we can also take it, as before, to mean the current graph when ve is the current edge. Indeed the edge ve^+_G can only be cut at v, so it cannot be cut while the process is away from v.

We claim that if the Bernardi process does not terminate upon returning to v along ev, then $ve^+_G = ve^+_{\mathrm{Bip}\,\mathcal{H}}$, in other words, $ve^+_{\mathrm{Bip}\,\mathcal{H}}$ is not yet removed from the graph. Take the first moment when for a current edge ve, the edge $ve^+_{\mathrm{Bip}\,\mathcal{H}}$ is already missing from the current graph. Until this moment, the edges incident to v became current edges in an order compatible with the cyclic order at v (in $\mathrm{Bip}\,\mathcal{H}$), and as $ve^+_{\mathrm{Bip}\,\mathcal{H}}$ is already removed from the graph, it has already been a current edge. Hence we conclude that all the edges incident to v have already become current edges. Thus, ve^+_G has also been a current edge. As ve^+_G is still part of the graph, it was kept (and thus traversed) when it first became a current edge. Now when it becomes current edge for the second time, it will be kept once again. Therefore it will be traversed again. But that means that the Bernardi process terminates after traversing ev. □

Lemma 4.7. *During a violet Bernardi run, the subgraph of traversed edges of* $\mathrm{Bip}\,\mathcal{H}$ *never contains a cycle.*

Proof. Suppose that our violet Bernardi run has the hypertree $\mathbf{f}$ as input (now $\mathbf{f}$ is a hypertree on E or on V depending on whether our Bernardi process is a (ht:E, cut:V) or a (ht:V, cut:V) one). Suppose for contradiction that after a while, a cycle appears in the subgraph of traversed edges of $\mathrm{Bip}\,\mathcal{H}$ and stop the process at the first moment when this occurs. Let O be this cycle, and G be the current graph at the moment. (Edges are typically traversed in pairs to form a violet → emerald → violet path. If the first of the two edges completes a cycle, then we stop the process right there, midway through a step.)

In the rest of the proof we refer to this aborted process only, so that we may apply Lemma 4.6 throughout. Note that O is the only cycle in the subgraph of traversed edges, as the addition of an edge can create only one cycle in a cycle-free graph.

We claim that for each edge of O, we can choose an orientation in which it was traversed, such that the chosen orientations give a cyclic orientation of O. Indeed, if this is not the case, then there are two edges xy and zw of O, such that xy was only traversed from the direction of x and zw was only traversed from the direction of w, furthermore these two directions are opposite with respect to the cycle O. Suppose that xy was the first one to be traversed among the two edges. Then after traversing xy the walk associated to the process was at y. Later it needed to reach w to be able to traverse zw from the direction of w. But as neither xy was traversed from the direction of y nor wz was traversed from the direction of z, the process could not go to w using the edges of O, hence by the time the edges of O are traversed, there needs to be another cycle in the subgraph of traversed edges, which is a contradiction.

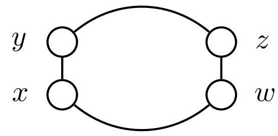

Pick a violet node along O and call it v_0. Then name the other vertices of the cycle $e_0, v_1, e_1, \dots, v_{t-1}, e_{t-1}$ such that $v_ie_i \in O$ and $e_iv_{i+1} \in O$ for each i, and v_ie_i was traversed from the direction of v_i and e_iv_{i+1} was traversed from the direction of e_i for each i (where indices are now understood modulo t). Since this is a violet Bernardi run, all the edges of the form v_ie_i are kept.

We will need the following two technical claims.

Claim 4.8. *An edge $e_{i-1}v_i$ cannot be deleted before v_ie_i becomes current edge.*

Proof. Suppose on the contrary that for some i, the edge $e_{i-1}v_i$ is deleted before v_ie_i becomes current edge. Since $e_{i-1}v_i$ is traversed by the Bernardi process, it is traversed from the emerald direction before it gets deleted. At the time of the traversal of $e_{i-1}v_i$ from the emerald direction, v_ie_{i-1} cannot have been current edge yet (because then it would have gotten deleted before traversed) and the next current edge is $v_ie_{i-1}^+$. By Lemma 4.6, while there is no cycle in the subgraph of traversed edges, the edges incident to a violet node become current edges in an order compatible with the cyclic order. Hence v_ie_i needs to become current edge before v_ie_{i-1}, which is a contradiction. □

Claim 4.9. *An edge e_iv_{i+1} cannot be deleted before v_ie_i becomes current edge.*

Proof. Suppose on the contrary that e_iv_{i+1} is deleted before v_ie_i becomes current edge. Since e_iv_{i+1} is traversed by the Bernardi process, it is traversed from the emerald direction before it gets deleted. Let us call the orientation of O compatible with orienting e_iv_{i+1} from e_i to v_{i+1} the positive orientation. As v_ie_i does not become current edge before e_iv_{i+1} is deleted, it also does not get traversed in the positive direction until that time. Hence there must be an edge xy in O such that the edges along the arc between v_{i+1} and x are all traversed in the positive direction after the traversal and before the deletion of e_iv_{i+1}, but xy is not. Let zx denote the edge of O preceding xy in the positive direction. We separate two cases.

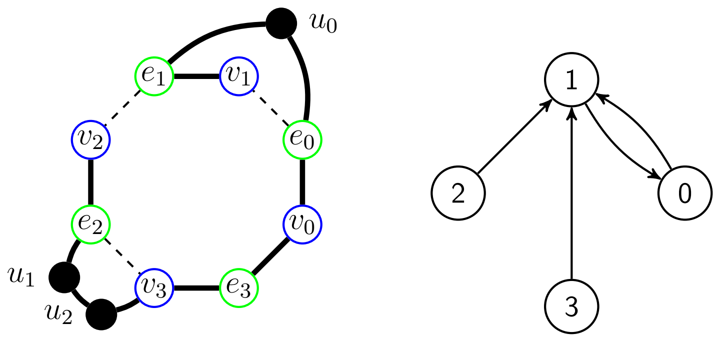

Figure 3: An example for a cycle O and spanning tree T, and the corresponding auxiliary graph in the (ht:E, cut:V) case. The T-tour of O is the following walk: $v_0, e_0, u_0, e_1, v_1, e_1, u_0, e_0, v_0, e_3, v_3, u_2, u_1, e_2, v_2, e_2, u_1, u_2, v_3, e_3, v_0$.

Case 1: $xy = v_{i+1}e_{i+1}$ (hence $z = e_i$). As $v_{i+1}e_{i+1}$ is not deleted during the process, if it is not traversed until e_iv_{i+1} is removed, then it does not become a current edge until $v_{i+1}e_i$ becomes current edge. After the traversal of e_iv_{i+1} from the emerald direction, the next current edge is $v_{i+1}e_i^+$. By Lemma 4.6, while there is no cycle in the subgraph of traversed edges, the edges incident to a violet node become current edges in an order compatible with the cyclic order. Hence $v_{i+1}e_{i+1}$ needs to become current edge before $v_{i+1}e_i$, which is a contradiction.

Case 2: $xy \neq v_{i+1}e_{i+1}$. This means that the Bernardi process reaches x through the edge zx before it deletes e_iv_{i+1}. Hence after reaching x, the walk associated to the process needs to get back to v_{i+1} with $v_{i+1}e_i$ as current edge before traversing xy from the direction of x. If the process does not traverse xz from the direction of x before it returns to v_{i+1}, then necessarily it traverses a cycle, from which we get a contradiction with the fact that there is no cycle in the graph of traversed edges until we traverse all the edges of O. On the other hand, Lemma 4.6 tells us that until all the edges of O are traversed, the edges adjacent to x are traversed (from the direction of x) in an order compatible with the ribbon structure. Hence xy must be traversed before xz, which is again a contradiction. □

Let us now refocus on the moment when, upon the traversal of all edges of O, we stopped the Bernardi process. Take a spanning tree T in the current graph G that realizes $\mathbf{f}$. Such a tree exists by the definition of the Bernardi process and it contains every edge that we so far decided to keep. In particular, T contains each edge of O of the form v_ie_i.

Let us take the following walk on T. Start from v_0, and traverse v_0e_0. If $e_0v_1 \in T$, then traverse e_0v_1 as well. If not, then traverse the unique path in T connecting e_0 to v_1. Continue this way until for each edge of O, the path in T connecting its two endpoints is traversed. At the end we arrive back to v_0. Let us call this walk the T-tour of O.

As T is a tree, in this tour, each traversed edge is traversed in both directions. Each edge of the form v_ie_i is traversed from the direction of v_i, hence each of these edges is also traversed in the reverse direction. By definition, an edge v_ie_i is traversed from the direction of e_i only if it is part of the path in T connecting some nodes e_j and v_{j+1}. Hence each edge of the form e_iv_i is in the base cycle $C(T, e_jv_{j+1})$ for some j where $e_jv_{j+1} \notin T$.

From this point on, we prove Lemma 4.7 separately for the (ht:E, cut:V) and for the (ht:V, cut:V) Bernardi processes. This is only for notational convenience as the two arguments remain very similar.

For the (ht:E, cut:V) Bernardi process, let us take the following auxiliary directed graph D. Let the vertex set of D be $\{0, \dots, t-1\}$. Draw an edge from i to j if $v_ie_i \in C(T, e_jv_{j+1})$. From the above remarks, each vertex has outdegree at least one in D. Hence D contains at least one directed cycle. Take a directed cycle of minimal length. Let the vertex set of this cycle be $x_0, \dots, x_{r-1}$, where there is a directed edge from x_i to x_{i+1} (that is, $v_{x_i}e_{x_i} \in C(T, e_{x_{i+1}}v_{x_{i+1}+1})$) for each i (meant modulo r).

We claim that $T' = \left(T \setminus \bigcup_{i=0}^{r-1} v_{x_i}e_{x_i}\right) \cup \bigcup_{i=0}^{r-1} e_{x_i}v_{x_i+1}$ is a spanning tree realizing the hypertree $\mathbf{f}$ on E. First we show that T' is a spanning tree of $\mathrm{Bip}\,\mathcal{H}$. (The following argument is a special case of a matroid theoretical lemma [13, Theorem 39.13], but we include it for completeness.) Having chosen a minimal cycle in D ensures that $v_{x_j}e_{x_j} \notin C(T, e_{x_i}v_{x_i+1})$ for any i and j such that $i \neq j+1$ (which is now meant modulo r), since otherwise there would be a shortcut in the cycle $x_0, \dots, x_{r-1}$, contradicting its minimality. Now $T_1 = (T \setminus v_{x_0}e_{x_0}) \cup e_{x_1}v_{x_1+1}$ is a spanning tree since $v_{x_0}e_{x_0} \in C(T, e_{x_1}v_{x_1+1})$. Moreover, since $v_{x_0}e_{x_0} \notin C(T, e_{x_j}v_{x_j+1})$ for $j \neq 1$, we have $C(T_1, e_{x_j}v_{x_j+1}) = C(T, e_{x_j}v_{x_j+1})$ for all $j \neq 1$. Hence when we replace the edge $v_{x_1}e_{x_1}$ of T_1 with $e_{x_2}v_{x_2+1}$, the result is another tree T_2 and further, the base cycles of $e_{x_3}v_{x_3+1}, \dots, e_{x_{r-1}}v_{x_{r-1}+1}, e_{x_0}v_{x_0+1}$ with respect to T_2 are still the same as with respect to T. By a trivial induction proof, we may continue switching edges like this until we arrive at T'.

Now to show that T' realizes $\mathbf{f}$, notice that by construction the degree of each emerald node is the same in T as in T'. As T realizes $\mathbf{f}$, so does T'.

Let $I = \{x_0, \dots, x_{r-1}\}$. Consider the first moment when one of the edges $v_\ell e_\ell$ becomes current edge for an index $\ell \in I$. By Claim 4.9, at this moment, for any $j \in I$, the edges v_je_j and e_jv_{j+1} are present in the current graph. Hence T' is a spanning tree in the current graph, realizing $\mathbf{f}$ and not containing $v_\ell e_\ell$. This contradicts the fact that $v_\ell e_\ell$ was kept by the process.

For the (ht:V, cut:V) Bernardi process, take a similar auxiliary directed graph D on the vertex set $\{1, \dots, t\}$. This time we draw an edge from i to j if $v_ie_i \in C(T, e_{j-1}v_j)$. Then again, each vertex has outdegree at least one in D and thus D contains at least one directed cycle. We take a directed cycle of minimal length and let its vertex set be $x_0, \dots, x_{r-1}$.

The fact that $T' = T \setminus \bigcup_{i=0}^{r-1} v_{x_i}e_{x_i} \cup \bigcup_{i=1}^{r} e_{x_i-1}v_{x_i}$ is a spanning tree realizing $\mathbf{f}$ (on V) can be established in the same way as in the previous case.

Let $I = \{x_0, \dots, x_{r-1}\}$. Take the first moment when an edge of the type $v_\ell e_\ell$ becomes current edge for an index $\ell \in I$. By Claim 4.8, at this moment, for any $j \in I$, the edges v_je_j and $e_{j-1}v_j$ are present in the current graph. Hence T' is a spanning tree in the current graph realizing $\mathbf{f}$ and not containing $v_\ell e_\ell$. This again contradicts the fact that $v_\ell e_\ell$ was kept.

Finally, as for both versions of the process the existence of O led to a contradiction, we conclude that the Bernardi process may never traverse a cycle. □

Lemma 4.10. *If an edge is traversed during a violet Bernardi run, then it is not removed later.*

Proof. If the edge ε is traversed from the direction of its violet node, then it is kept and, as we have already argued after Definition 4.2, is never removed later.

Now consider the case when $\varepsilon = ev$ is traversed from the direction of its emerald node e. Suppose for contradiction that later on ε gets deleted. We claim that v was first reached (by the walk associated to the process) through ev. Suppose that v was first reached through another edge $e'v$. Then take the traversed path from the time of the first arrival to v until ev gets traversed. Since ve does not get traversed from the violet direction (as it is due to be deleted), this path arrives back to v so that it only traverses ev in one direction. Hence there is a cycle in the set of traversed edges, contradicting Lemma 4.7.

Stop the Bernardi process immediately after the deletion of ε, and let the current graph of that moment be G. The graph G is connected because $\mathbf{f}$ is a hypertree in it. Let S be the set of nodes that were reached by the Bernardi process until the deletion of ε, and were first reached after reaching v. Let also $v \in S$ which guarantees that $S \neq \varnothing$. Note that e was reached before v, i.e., $e \notin S$ and hence $S \neq E \cup V$.

Take a vertex $x \in S$. At the current moment, the Bernardi process is at v, as the last step was deleting ve. If $x \neq v$, then when we last left x before returning to v, we must have left it along the same edge through which we first reached it, because otherwise there would be two paths between v and x and hence a cycle in the subgraph of traversed edges. By Lemma 4.6, the edges incident to x in G that were traversed from the direction of x are consecutive edges according to the cyclic order at x (in G). Hence all the edges incident to x in G are already traversed from the direction of x. Note that this is also true for $x = v$ as the last current edge was ve which is also the edge through which v was first reached.

Take an edge xy of $\operatorname{Bip}\mathcal{H}$ such that $x \in S$ and $y \in (E \cup V) \setminus S$. If xy is in G, then xy was traversed from the direction of x. Hence y is already reached. As $y \notin S$, the node y was also reached before reaching v. As x was first reached after reaching v (or else $x = v$), there is a path in the subgraph of traversed edges from y to x through v. But then after the traversal of xy from the direction of x, there would be a cycle in the subgraph of traversed edges. We conclude that there can be no edge xy in G such that $x \in S$ and $y \in (E \cup V) \setminus S$. Hence G is disconnected, which is a contradiction. (In other words, right before its deletion ε was the only edge connecting S to its complement and thus it should not have been deleted.) □

Proof of Theorem 4.5. By Lemma 4.10, the subgraph of traversed edges is part of the current graph at all times, and by Lemma 4.7, it is cycle-free. The current graph is always connected as it contains a spanning tree realizing $\mathbf{f}$. Hence if we make sure that at the end of the process, the subgraph H of traversed edges coincides with the current graph G, then we will know that $H = G$ is both cycle-free and connected, thus a spanning tree, and since it contains a realization of $\mathbf{f}$, it itself has to be such a realization.

Now for this last remaining claim, it suffices to prove the assertion in the Theorem that each edge becomes current edge exactly once, because then in particular each edge is either deleted (thus belongs neither to G nor to H) or kept and traversed (hence belongs to both). In fact, it is enough to show that the process examines each edge of $\mathrm{Bip}\,\mathcal{H}$ as a current edge at least once, i.e., by the time an edge would be traversed for the second time from the same direction, each edge has been current edge. This is because if an edge becomes current for a second time, then it has to have been kept the first time and thus traversed both times, ending the process.

We shall first consider where the process may end, and find that it is only possible at the base node b_0.

Suppose that the first edge to be traversed twice from the same direction is ev, and it is traversed from the direction of the emerald node e. If $e \neq b_0$, then take the edge ue through which we first reached e during the process. As the Bernardi process does not traverse a cycle, by Lemma 4.6, the edges incident to e are traversed in an order compatible with the cyclic order at e. Hence before ev is traversed for the second time, eu is traversed from the emerald direction (by Lemma 4.10 it cannot be deleted from the graph, as it was already traversed). But as, again, the Bernardi process does not traverse a cycle, after traversing eu from the emerald direction, we can only get back to e by traversing ue from the violet direction. Hence before ev gets traversed for the second time from the emerald direction, ue gets traversed twice from the violet direction, which is a contradiction.

Now suppose that the first edge to be traversed twice from the same direction is ve, and it is traversed from the violet direction, that is, from v. If $v \neq b_0$, then there must be an edge $e'v$ through which we first reached v. As $e'v$ is traversed, by Lemma 4.10, it is not deleted from the graph during the process. Now we can repeat the argument of the previous case: The Bernardi process does not traverse a cycle, and by Lemma 4.6, the edges incident to v are traversed in an order compatible with the cyclic order at v. Hence before ve is traversed for the second time from the violet direction, ve' is traversed from the violet direction (as it is not deleted). But as the Bernardi process does not traverse a cycle, after traversing ve' from the violet direction, it can only get back to v by traversing $e'v$ from the emerald direction. Hence before ve gets traversed for the second time from the violet direction, $e'v$ gets traversed twice from the emerald direction, which is a contradiction.

We conclude that the first edge to be traversed twice by the Bernardi process can only be an edge incident to b_0, and it is traversed from the direction of b_0.

Next we claim that if a violet node $v \in V - \{b_0\}$ is reached by the Bernardi process, then all the edges incident to it get examined as current edges. Indeed, as the Bernardi process does not traverse a cycle, the last time the process leaves v before returning to b_0, it must leave v on the same edge on which it was first reached. Hence from Lemma 4.6, by this time each edge incident to v is examined as a current edge. This also holds for b_0, if $b_0 \in V$, since the current edge at the end of the process is an edge that has already been a current edge. We also claim that if a node $e \in E$ is reached by the Bernardi process at all, then all the edges incident to it get either removed or traversed. Indeed, if $e \neq b_0$, then the last time the process leaves e before returning to b_0, it must leave e on the same edge on which it was first reached, and by Lemma

4.6, by this time the traversed edges incident to e have been traversed in an order compatible with the cyclic order at e in the current graph of that moment. Hence if an edge incident to e has not yet been deleted, then it has necessarily been traversed. Again, the same holds if $e = b_0$.

To complete the proof, we need to show that each node in V is reached by the process. If that was not the case, then the vertex set X of the subgraph H of traversed edges (at the end of the process) would be a proper subset of $E \cup V$, giving rise to a cut K in $\operatorname{Bip}\mathcal{H}$. As $\operatorname{Bip}\mathcal{H}$ is connected, K is a non-empty set of edges. By the previous paragraph, all edges in K get examined, and hence removed, by the process. We claim that K is such that each of its edges has its violet end in X. Indeed, edges can only be cut by the process at their violet ends and violet nodes in $E \cup V \setminus X$ are never reached. But this means that when the last edge of K was removed, the current graph became disconnected, which contradicts the basic principle of the Bernardi process. □

By interchanging the roles of E and V, we see that the statement of Theorem 4.5 holds also for the (ht:V, cut:E) and the (ht:E, cut:E) Bernardi processes.

Now that we have made sure that the Bernardi process converges to a spanning tree T, we may spell out the relationship between the walk associated to our process and Bernardi's original tour of T. As the process traverses exactly the edges of T, we see that the two are essentially the same, with the only difference being that we only consider edges as current edge once, whereas Bernardi's tour does so twice, at both endpoints. More precisely, we have the following.

Lemma 4.11. *Let T be a spanning tree of the ribbon graph* $\operatorname{Bip}\mathcal{H}$ *(with fixed base node and base edge), resulting from a violet Bernardi run on some hypertree (on E or on V). If we list the edges of* $\operatorname{Bip}\mathcal{H}$ *in the order in which they become current in the tour of T and delete*

- *the second occurrences of the edges not in T and*
- *for each edge of T, its occurrence in conjunction with its emerald endpoint,*

then we get the list in which the edges of $\operatorname{Bip}\mathcal{H}$ *become current in our hypergraphical Bernardi process.*

5 Hypergraph polynomials à la Bernardi

Using the process of Definition 4.2, that is the version that cuts edges of the bipartite graph near where hypertree values (spanning tree degrees) are assigned, we are able to give an alternative definition (analogous to the definition of the Tutte polynomial by Bernardi) for the interior polynomial of a hypergraph. In section 7, we show that this definition is equivalent to the original one; in particular, the polynomial does not depend on the choice of ribbon structure and base point. For the time being, we will denote this Bernardi-type interior polynomial by $\tilde{I}$ to avoid confusion. We will also set up an exterior version $\tilde{X}$.

First we define embedding activities. Note that this would not work without Theorem 4.5, in particular the claim that the Bernardi process reaches every part of the bipartite graph.

Definition 5.1 (internal and external embedding inactivity of a hypertree)**.** Let $\mathcal{H} = (V, E)$ be a hypergraph and $\mathbf{f}$ be a hypertree on E. Fix a ribbon structure on $\operatorname{Bip}\mathcal{H}$, base node, and base edge and run the (ht:E, cut:E) Bernardi process. Take the order in which the edges of $\operatorname{Bip}\mathcal{H}$ become current edges during the process. This induces the following order on the elements of E:

$$e_1 <_{\mathbf{f}} e_2 \iff \text{some edge incident to } e_1 \text{ becomes current edge earlier than any edge incident to } e_2. \tag{5.1}$$

The *internal embedding inactivity* of $\mathbf{f}$ is the number of internally inactive (that is, not internally active) hyperedges with respect to $\mathbf{f}$ and the above order, defined exactly as in Definition 2.3. We denote this quantity by $\overline{ie}(\mathbf{f})$.

The *external embedding inactivity* of $\mathbf{f}$ is the number of externally inactive hyperedges with respect to $\mathbf{f}$ and the above order, defined as in [8, Definition 5.1]. We denote this quantity by $\overline{ee}(\mathbf{f})$. I.e., $\overline{ee}(\mathbf{f})$ is the number of hyperedges which, with respect to $\mathbf{f}$ and $<_{\mathbf{f}}$, may receive a transfer of valence from a smaller hyperedge.

Example 5.2. We return to Example 4.3 (and its notation), where we ran the (ht:E, cut:E) Bernardi process on a hypertree $\mathbf{f}$ on E. From the order in which the edges of $\operatorname{Bip}\mathcal{H}$ became current (cf. Figure 2), we see that the induced order on E is left $<_{\mathbf{f}}$ top $<_{\mathbf{f}}$ right, which is just the order in which the walk associated to the process reached the elements of E.

The smallest node is always active, both internally and externally. As the top node has the minimal possible $\mathbf{f}$-value, it is internally active (cannot transfer valence to left). Externally, however, it is inactive: it is easy to see (for instance by symmetry) that left $\mapsto 0$, top $\mapsto 1$, right $\mapsto 2$ is also a hypertree, i.e., with respect to $\mathbf{f}$, top may receive a transfer of valence from left. Notice that it does not concern us at all whether the new hypertree has some realization related to the just-constructed realization of $\mathbf{f}$. Finally the right node is internally inactive (can transfer valence to both left and top) but externally active because its $\mathbf{f}$-value is the largest possible and hence may not receive any transfers of valence.

Thus, in this case we find the embedding inactivity values $\overline{ie}(\mathbf{f}) = \overline{ee}(\mathbf{f}) = 1$.

Definition 5.3. For a hypergraph $\mathcal{H} = (V, E)$ with a set of hypertrees B_E and with a ribbon structure on $\operatorname{Bip}\mathcal{H}$, we let

$$\tilde{I}_{\mathcal{H}}(\xi) = \sum_{\mathbf{f}\in B_E} \xi^{\overline{ie}(\mathbf{f})} \quad \text{and} \quad \tilde{X}_{\mathcal{H}}(\eta) = \sum_{\mathbf{f}\in B_E} \eta^{\overline{ee}(\mathbf{f})}. \tag{5.2}$$

We prove the following theorem in Section 7.

Theorem 5.4. *For any connected hypergraph $\mathcal{H}$, ribbon structure on* $\operatorname{Bip}\mathcal{H}$*, base node, and base edge, we have $\tilde{I}_{\mathcal{H}} = I_{\mathcal{H}}$. Here $I_{\mathcal{H}}$ is the interior polynomial of subsection 2.2, whereas $\tilde{I}_{\mathcal{H}}$ is defined in* (5.2) *using the (ht:E, cut:E) Bernardi process. In particular, $\tilde{I}_{\mathcal{H}}$ does not depend on the ribbon structure.*

When applied to a (ribbon) graph $\mathcal{H} = G = (V, E)$, the (ht:E, cut:E) and (ht:E, cut:V) processes induce the same order on the set of edges E (cf. Section 8). This is not true for general hypergraphs. Thus when we use, in the exact same way as above, the (ht:E, cut:V) process to define interior embedding inactivities (for hypertrees on E), the individual values of the statistic will change. Also, the order on E (induced by a hypertree on E) is much less natural in this case: nodes in E may 'get numbered' when the walk associated to the process *cuts* some edge adjacent to them, which may be well before the walk actually *reaches* the node. Yet computer simulations suggest that the distribution of the interior inactivity statistic remains the same. Unfortunately, at the moment, we are only able to state this as a conjecture.

Conjecture 5.5. *Let $\mathcal{H} = (V, E)$ be a connected hypergraph with a ribbon structure and base point for* Bip $\mathcal{H}$*. For a hypertree* $\mathbf{f}$ *on E, run the (ht:E, cut:V) Bernardi process and define the order $<_{\mathbf{f}}$ on E as in* (5.1)*. Compute the internal (embedding) inactivity of* $\mathbf{f}$ *with respect to $<_{\mathbf{f}}$ and define a one-variable polynomial just like in the first equation of* (5.2)*. The result coincides with the interior polynomial of $\mathcal{H}$.*

We also conjecture that both Theorem 5.4 and Conjecture 5.5 hold true for the exterior polynomial, too.

Conjecture 5.6. *For any connected hypergraph $\mathcal{H}$, ribbon structure on* Bip $\mathcal{H}$*, base node, and base edge, we have $\tilde{X}_{\mathcal{H}} = X_{\mathcal{H}}$. In particular, the polynomial $\tilde{X}_{\mathcal{H}}$ of* (5.2) *does not depend on the ribbon structure. The same two claims remain true if we re-define $\tilde{X}_{\mathcal{H}}$ using the (ht:E, cut:V) Bernardi process.*

We are probably much farther away from the proof of Conjecture 5.6 than that of 5.5. This is because we do not know what geometric object should play the same role, in relation to exterior activities, as the root polytope plays for interior activities in the argument that we are about to present.

6 Jaeger trees

In this section we give an alternative characterization of the set of spanning trees of Bip $\mathcal{H}$ that the Bernardi process outputs. Here $\mathcal{H} = (V, E)$ is a connected hypergraph and just like in the previous two sections, we work with a fixed ribbon structure, base node, and base edge.

6.1 Definitions

Since the notion of the 'walk associated to the Bernardi process' of Section 4 only applies in the presence of a hypertree, and Jaeger trees will be defined without referring to hypertrees, we will revert to the definition of the tour of a spanning tree from subsection 2.3. For a spanning tree T of the graph Bip $\mathcal{H}$, we can take the tour of T and hence obtain the order $\leq_T$ on the set of incident node-edge pairs of Bip $\mathcal{H}$. If the edge xy is not in T, then by Lemma 2.6, it is skipped twice by the tour, once with each endpoint as current node. Let us consider the first of the two events and

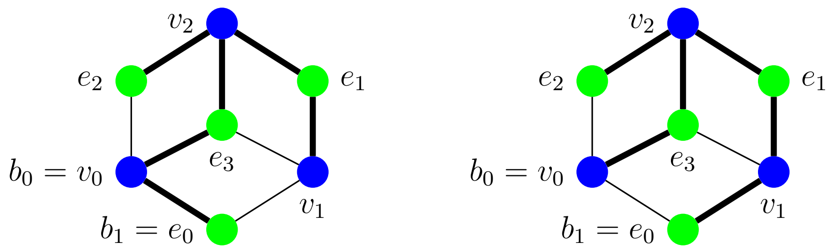

Figure 4: Let the ribbon structure of the bipartite graph be induced by the counterclockwise orientation of the plane. The spanning tree on the left is an E-cut Jaeger-tree, while the spanning tree on the right is not a Jaeger tree.

say that xy is *cut* at x during the tour of T if $(x, xy) <_T (y, xy)$. In this case we will also think of the 'time' when (x, xy) is current as the 'moment when xy is cut from the graph' (on our way of trimming Bip $\mathcal{H}$ down to T). We also make the following definitions.

Definition 6.1 (V-cut and E-cut Jaeger trees)**.** A spanning tree T of Bip $\mathcal{H}$ is called a *V-cut* (resp. *E-cut*) *Jaeger tree* if in the tour of T, each non-edge of T is cut at its violet (resp. emerald) endpoint.

We say that a spanning tree of Bip $\mathcal{H}$ is a *Jaeger tree* if it is either a V-cut or an E-cut Jaeger tree. I.e., Jaeger trees are distinguished by the property that in their tours, each non-edge is 'skipped for the first time at the same color.' See Figure 4 for an example. We will see that Jaeger trees are exactly the outcomes of the Bernardi process. Let us prove the easier implication first.

Proposition 6.2. *The (ht:E, cut:V) and the (ht:V, cut:V) Bernardi processes produce V-cut Jaeger trees for any hypertree (on E or on V, respectively).*

Proof. Let T be one of the spanning trees in question. In Lemma 4.11 we spelled out the correspondence between the tour of T and the generalized Bernardi process which produced T. To put it simply, since in the Bernardi process an edge needs to be traversed if we arrive at it from the emerald direction, in the tour of T, each edge not in T has to be first reached, and cut, at its violet endpoint. Hence T is indeed a V-cut Jaeger tree. □

With the same proof, we also obtain the following.

Proposition 6.3. *The (ht:V, cut:E) and the (ht:E, cut:E) Bernardi processes produce E-cut Jaeger trees for any hypertree.*

Remark 6.4. The sequence of current vertex-edge pairs of the tour of a spanning tree can also be described using the topology of the ribbon surface of Remark 2.1. Namely, we start from the base point (meant as a boundary point of the surface) and proceed along the boundary in the positive direction. When we reach a vertex of one of the rectangle pieces, then

- the midpoint of the disk, together with the core of the rectangle, become current, and
- if the edge is in the tree, we proceed along the long side of the rectangle to the next adjacent disk (traversal), and continue along its boundary, or
- if the edge is not in the tree, then we proceed along the short side (skipping the edge), and then on along the boundary of the same disk. The first time this happens at a non-edge, we may imagine that it has been physically cut from the graph.

Now if the graph is bipartite, and the tree is an outcome of the Bernardi process run on some hypertree, then from the sequence above we obtain the sequence of current edges in the hypergraphical Bernardi process of Section 4 as in Lemma 4.11. Cf. Example 4.3, where in Figure 2 the gray curves can be understood as portions of the continuous path we have just described.

Our next aim is to prove the converse of Proposition 6.2, namely that each Jaeger tree is obtained as the outcome of the Bernardi process on some hypertree. Until we finish proving this (in Theorem 6.17) we will continue to rely on the notions of tour and current node-edge pair of Subsection 2.3, which apply to all spanning trees. First let us establish some useful properties of Jaeger trees.

By the *reverse* of a ribbon structure, we mean the one where all cyclic orders are turned opposite. This is equivalent to reversing the orientation of the ribbon surface. Below we consider the reversed ribbon structure with the same base point, i.e., the same boundary point of the ribbon surface.

Lemma 6.5. *A V-cut Jaeger tree with base node b_0 and base edge b_0b_1 is an E-cut Jaeger tree for the reversed ribbon structure, base node b_0, and base edge $b_0b_1^-$ (where the latter is meant in the original ribbon structure).*

Proof. Let T be a spanning tree of $\mathrm{Bip}\,\mathcal{H}$. In the tour of T, each non-edge ε of T is visited twice, once at each endpoint. The tree T is V-cut if and only if the violet endpoint is always reached before the emerald one. For the reversed ribbon structure, the order of the two occurrences of ε in the tour becomes opposite. Hence T is V-cut for the original structure if and only if it is E-cut for the reversed structure. □

Note that by symmetry, an E-cut Jaeger tree with base node b_0 and base edge b_0b_1 is also a V-cut Jaeger tree with base node b_0, base edge $b_0b_1^-$ and the reversed ribbon structure. As the reversal of the ribbon structure is an involution, we get the following.

Lemma 6.6. *The set of E-cut Jaeger trees with base node b_0 and base edge b_0b_1 is equal to the set of V-cut Jaeger trees with base node b_0, base edge $b_0b_1^-$ and the reversed ribbon structure.*

Definition 6.7. If T is a V-cut Jaeger tree, let us call the tour of T (with respect to the original ribbon structure) the *violet tour of T*. Let us call the tour of T with base point b_0, base edge $b_0b_1^-$ and the reversed ribbon structure the *emerald tour of T*.

By Lemma 6.5, for a V-cut Jaeger tree, in the emerald tour, each non-edge is cut at its emerald endpoint. The violet and the emerald tours of a Jaeger tree both induce orderings of the edges of $\operatorname{Bip}\mathcal{H}$, as follows.

Definition 6.8. For a V-cut Jaeger tree T, let the *violet T-order of the edges of* $\operatorname{Bip}\mathcal{H}$ be the order in which the edges of $\operatorname{Bip}\mathcal{H}$ appear in the violet T-tour with their violet endpoint as current node. Let us denote this ordering by $<_{T,V}$. That is, $ev <_{T,V} e'v'$ if $(v, ev) <_T (v', e'v')$.

Similarly, let the *emerald T-order of the edges of* $\operatorname{Bip}\mathcal{H}$ be the order in which the edges of $\operatorname{Bip}\mathcal{H}$ appear in the emerald T-tour with their emerald endpoint as current node. Let us denote this ordering by $<_{T,E}$; i.e., $ev <_{T,E} e'v'$ if $(e, ev) <_T (e', e'v')$.

Definition 6.9. Let $T \subset \operatorname{Bip}\mathcal{H}$ be a Jaeger tree. The *order induced on E by the emerald T-order* of the edges of $\operatorname{Bip}\mathcal{H}$, or simply the *emerald T-order on E*, is by the smallest of the incident edges. That is, e_1 is less than e_2 if there is an edge, incident to the node e_1, that is smaller in the emerald T-order than any edge incident to e_2. There is a similar order of the nodes in V, induced by the violet T-order of the edges.

Example 6.10. For the Jaeger tree on the left side of Figure 4, the emerald T-order is $e_0v_1, e_0v_0, e_3v_1, e_3v_2, e_1v_1, e_1v_2, e_2v_0, e_2v_2, e_3v_0$, while the violet T-order is $v_0e_2, v_0e_3, v_2e_2, v_2e_1, v_1e_0, v_1e_3, v_1e_1, v_2e_3, v_0e_0$. The former induces the order e_0, e_3, e_1, e_2 of the emerald nodes and the latter induces the order v_0, v_2, v_1 of the violet nodes.

Remark 6.11. The previous two definitions are almost superfluous, in that they describe orderings that we have already seen — we just cannot prove this yet. But once we know that our Jaeger tree is an outcome of the Bernardi process, in fact in two different senses by Lemma 6.6, Lemma 4.11 will imply that the orders of Definition 6.8 are the same in which the edges of $\operatorname{Bip}\mathcal{H}$ become current in those processes.

More precisely, if a V-cut Jaeger-tree T is the outcome of the (ht:V, cut:V) or the (ht:E, cut:V) Bernardi process, then in the run producing T, the edges of $\operatorname{Bip}\mathcal{H}$ become current edges in the violet T-order. Also, if an E-cut Jaeger-tree T is the outcome of the (ht:E, cut:E) or the (ht:V, cut:E) Bernardi process, then in the run producing T, the edges of $\operatorname{Bip}\mathcal{H}$ become current edges in the emerald T-order.

Similarly, we will eventually see that the orders of Definition 6.9 are instances of (5.1) in Definition 5.1.

6.2 Base cuts of Jaeger trees

Base cuts of Jaeger trees have a nice interplay with the orderings of the edges induced by the trees. The following lemma is proved in [3]. We specialize it to our case and translate it into our notation.

Lemma 6.12. *[3, Lemma 5] Let T be a spanning tree of* $\operatorname{Bip}\mathcal{H}$ *and $xy \in T$. Let T_0 and T_1 be the two subtrees of $T - xy$ so that T_0 contains the base node. If $(x, xy) <_T (y, xy)$, then x is in the node set of T_0. Moreover,*

$$\{(z, zw) \mid (x, xy) <_T (z, zw) \leq_T (y, xy)\} = \{(z, zw) \mid z \textit{ is in the node set of } T_1\}.$$

For a subgraph of Bip $\mathcal{H}$, we let its *base component* be the connected component that contains the base node.

Lemma 6.13 (Base cuts of Jaeger trees). *Let T be a V-cut Jaeger tree, $\varepsilon \in T$, and $\varepsilon_1, \varepsilon_2 \in C^*(T, \varepsilon)$. If ε_1 has its violet endpoint in the base component of $T - \varepsilon$ and ε_2 has its emerald endpoint in the base component of $T - \varepsilon$, then*

(i) $\varepsilon_1 <_{T,V} \varepsilon_2$, *and*

(ii) $\varepsilon_1 \leq_{T,V} \varepsilon$.

In other words, the tour of T first visits those edges of $C^*(T, \varepsilon)$ that have their violet endpoint in the base component of $T - \varepsilon$, the last one of which, and in that case the only one not to be cut, may be ε. If ε has its emerald endpoint in the base component of $T-\varepsilon$, then it is the first such edge to be visited by the tour, but because it is traversed from the emerald direction, it does not become smallest among them with respect to $<_{T,V}$. (In fact, somewhat counterintuitively, it will be the largest.) In both cases, by traversing ε the tour leaves the base component and visits, and cuts at their violet endpoints, the remaining edges of $C(T, \varepsilon)$ having their emerald endpoint in the base component.

Proof. Let $\varepsilon = xy$. The subgraph $T - \varepsilon$ is a forest made up of the base component T_0 and another component T_1. Suppose that $(x, \varepsilon) <_T (y, \varepsilon)$. Take the violet tour of T. By Lemma 6.12, we have

$$\{(z, zw) \mid (x, xy) <_T (z, zw) \leq_T (y, xy)\} = \{(z, zw) \mid z \text{ is in the node set of } T_1\}.$$

For $i = 1, 2$, let $\varepsilon_i = v_i e_i \in C^*(T, \varepsilon)$, where $v_i \in V$ and $e_i \in E$. Since ε_1 has its violet endpoint v_1 in T_0, its emerald endpoint e_1 is from the node set of T_1. Hence $(x, \varepsilon) <_T (e_1, \varepsilon_1) \leq_T (y, \varepsilon)$. Since this is the violet tour of T, we either have $\varepsilon_1 = \varepsilon$ (and then $x = v_1$) or the edge ε_1 is cut at its violet endpoint, i.e., $(v_1, \varepsilon_1) <_T (e_1, \varepsilon_1)$. As v_1 is not in the node set of T_1, we must have $(v_1, \varepsilon_1) \leq_T (x, \varepsilon)$. From this, (ii) immediately follows, for no matter if x or y is the violet endpoint of ε, we have $(v_1, \varepsilon_1) \leq_T (x, \varepsilon) <_T (y, \varepsilon)$.

Since ε_2 has its emerald endpoint e_2 in T_0, its violet endpoint v_2 is from the node set of T_1. Hence $(x, \varepsilon) <_T (v_2, \varepsilon_2) \leq_T (y, \varepsilon)$. We conclude that $(v_1, \varepsilon_1) <_T (v_2, \varepsilon_2)$, implying $\varepsilon_1 <_{T,V} \varepsilon_2$. □

Lemma 6.14. *If the violet tours of the V-cut Jaeger trees T and T' coincide until ε becomes current edge, furthermore $\varepsilon \in T$ but $\varepsilon \notin T'$, then ε has its violet endpoint in the base component of $T - \varepsilon$. Moreover, for any $\varepsilon' \in T'$ such that $\varepsilon' \in C^*(T, \varepsilon)$, the emerald endpoint of ε' lies in the base component of $T - \varepsilon$.*

Proof. Up to the time when the two tours diverge, ε has been neither cut (for otherwise it could not be an edge in T) nor traversed (because then it would be an edge in T'). So this is when the tour of T' will cut it and since T' is a V-cut Jaeger tree, the current node v has to be violet. As v is connected to b_0 in T by a path not traversing ε, the first claim follows.

From part (ii) of Lemma 6.13 we know that all the edges in $C^*(T,\varepsilon)\setminus\{\varepsilon\}$ that have their violet endpoints in the base component had already been current edges, in conjunction with their violet endpoints, in the violet tour of T before reaching ε, and they were not included into T. As the violet tours of T and T' coincide until reaching ε, these edges are not included in T' either. Since we also have $\varepsilon \notin T'$, all the edges of T' from $C^*(T,\varepsilon)$ have their emerald endpoints in the base component of $T-\varepsilon$. □

6.3 All Jaeger trees are outcomes of the Bernardi process

In this subsection we prove that each hypertree is realized by a unique Jaeger tree, hence Jaeger trees are exactly the trees obtained as outcomes of the Bernardi process.

Theorem 6.15. *Let $\mathcal{H}=(V,E)$ be a connected hypergraph with a ribbon structure on* $\operatorname{Bip}\mathcal{H}$*, a base node, and a base edge fixed. For each hypertree $\mathbf{f}\in B_E$ there is exactly one E-cut Jaeger tree T such that $\mathbf{f}=\mathbf{f}_E(T)$.*

Proof. We proved in Proposition 6.3 that for each hypertree $\mathbf{f}\in B_E$, the (ht:E, cut:E) Bernardi process produces an E-cut Jaeger tree realizing $\mathbf{f}$. Hence it is enough to show that for each hypertree, there is at most one suitable Jaeger tree. Suppose for a contradiction that for some hypertree $\mathbf{f}\in B_E$ there are two E-cut Jaeger trees $T\neq T'$ such that $\mathbf{f}=\mathbf{f}_E(T)=\mathbf{f}_E(T')$.

Consider the emerald tours of T and T', and suppose that the two tours first differ when an edge ve is included into T but not included into T'. As T and T' are E-cut Jaeger trees, this means that the current node at this time is e. Cf. Lemma 6.14, which also says (note the roles of the colors changed) that each edge in $T'\cap C^*(T,ve)$ has its violet endpoint in the base component of $T-ve$.

Let V_0 be the set of violet nodes and E_0 be the set of emerald nodes in the base component of $T-ve$. We will compute $\sum_{u\in E_0}\mathbf{f}(u)$ twice, using the realizations T and T'.

As edges of T between elements of $V_0\cup E_0$ form a tree and other than those, T only has one edge, ve, incident to an element of E_0, we have $\sum_{u\in E_0}\mathbf{f}(u)=(|V_0\cup E_0|-1)+1-|E_0|=|V_0|$.

On the other hand, as T' is a tree, it can have at most $|V_0\cup E_0|-1$ edges between elements of $V_0\cup E_0$. All other edges of T' incident to E_0 belong to $C^*(T,ve)$ but we have already seen that there are no such edges. Hence this time we obtain $\sum_{u\in E_0}\mathbf{f}(u)\leq(|V_0\cup E_0|-1)-|E_0|=|V_0|-1$, which is a contradiction. □

Corollary 6.16. *Under the same assumptions as in Theorem 6.15, for each hypertree $\mathbf{f}\in B_E$ there is exactly one V-cut Jaeger tree T such that $\mathbf{f}=\mathbf{f}_E(T)$. Similarly, hypertrees in B_V have unique E-cut and V-cut Jaeger tree representatives.*

Proof. Note that $\mathbf{f}$ does not depend on the ribbon structure. By Lemma 6.5, the trees we are considering are exactly the E-cut Jaeger tree representatives of $\mathbf{f}$ with respect to the reverse ribbon structure. According to Theorem 6.15, there is a unique such tree. The second claim follows by interchanging the roles of colors. □

In Proposition 6.2 we saw that each Bernardi process of Section 4 provides a function from the set of hypertrees (on an arbitrarily fixed color class of Bip $\mathcal{H}$) to the appropriate set of Jaeger trees. The map is an injection because the same tree cannot represent different hypertrees. By Theorem 6.15 or Corollary 6.16, the map must also be a surjection. Hence we obtain the following two conclusions.

Theorem 6.17. *For any connected bipartite graph with a fixed ribbon structure, base node, and base edge, the outcomes of the (ht:E, cut:V) Bernardi process, as well as the outcomes of the (ht:V, cut:V) Bernardi process, are exactly the V-cut Jaeger trees.*

In the same way, the outcomes of the (ht:E, cut:E) Bernardi process, as well as the outcomes of the (ht:V, cut:E) Bernardi process, are exactly the E-cut Jaeger trees.

Corollary 6.18. *For any connected bipartite graph with color classes E, V and a fixed ribbon structure, base node, and base edge, the (say) V-cut Jaeger trees give a bijection between the sets of hypertrees B_E and B_V, namely for each hypertree on E, there is exactly one V-cut Jaeger tree that realizes it and the same holds for each hypertree on V.*

This bijection generalizes the bijection between spanning trees and break divisors of a graph found by Bernardi [3] and studied further by Baker and Wang [1].

In section 7 we will give a geometric interpretation of the above facts by proving that the simplices corresponding to Jaeger trees form a dissection of the root polytope. That result also implies the statements of the present subsection, and in particular shows that the Jaeger trees give a bijection between B_E and B_V as discussed in Theorem 3.2.

6.4 The four types of Bernardi processes

In our treatment of hypergraphs $\mathcal{H}$ through their associated bipartite graphs Bip $\mathcal{H}$, the roles of the emerald and violet color classes are inherently symmetric. In particular, there are only two different versions of the Bernardi process, as described in Section 4. But if one is solely interested in graphs G, as opposed to hypergraphs, the construction of Bip G becomes somewhat unnatural and the symmetry is easy to miss. In this subsection we formulate statements on four Bernardi processes: the two above and their transposes. Capitalizing on the nice structure of Jaeger trees, we deduce several relationships between them.

Theorem 6.19. *Run the (ht:E, cut:V) Bernardi process on the hypertree* $\mathbf{f}$ *(on E), and let the resulting spanning tree of* Bip $\mathcal{H}$ *be T. Then by running the (ht:V, cut:V) Bernardi-process on the induced (on V) hypertree* $\mathbf{f}_V(T)$*, we again get T as resulting spanning tree.*

Proof. Let the spanning tree produced by the (ht:V, cut:V) Bernardi-process on $\mathbf{f}_V(T)$ be T'. Then by Proposition 6.2, T and T' are both V-cut Jaeger trees, and they both realize $\mathbf{f}_V(T)$ on V. Hence by Theorem 6.15, we have $T' = T$. □

Theorem 6.20. *Run the (ht:E, cut:V) Bernardi process on the hypertree* $\mathbf{f}$ *(on E), and let the resulting spanning tree of* $\operatorname{Bip}\mathcal{H}$ *be T. Then by running the (ht:E, cut:E) Bernardi process on* $\mathbf{f}$*, with base node b_0, base edge $b_0b_1^-$, and the reversed ribbon structure, we again get T as resulting spanning tree.*

Proof. Let T' be the spanning tree produced by the (ht:E, cut:E) Bernardi process on $\mathbf{f}$ with base node b_0, base edge $b_0b_1^-$ and the reversed ribbon structure. Then T' is an E-cut Jaeger tree for the reversed ribbon structure. Hence by Lemma 6.5, T' is a V-cut Jaeger tree with base node b_0, base edge b_0b_1 and the original ribbon structure. Thus T and T' are both V-cut Jaeger trees, and they both realize $\mathbf{f}$ on E. Hence by Corollary 6.16, we have $T' = T$. □

Theorem 6.21. *Run the (ht:E, cut:V) Bernardi process on the hypertree* $\mathbf{f}$ *(on E), and let the resulting spanning tree of* $\operatorname{Bip}\mathcal{H}$ *be T. Then by running the (ht:V, cut:E) Bernardi process on $\mathbf{f}_V(T)$, with base node b_0, base edge $b_0b_1^-$ and the reversed ribbon structure, we again get T as resulting spanning tree.*

Proof. Let T' be the spanning tree produced by the (ht:V, cut:E) Bernardi process on $\mathbf{f}_V(T)$, with base node b_0, base edge $b_0b_1^-$, and the reversed ribbon structure. Then T' is an E-cut Jaeger tree in that setup and by Lemma 6.5, T' is a V-cut Jaeger tree with base node b_0, base edge b_0b_1 and the original ribbon structure. Hence T and T' are both V-cut Jaeger trees, and they both realize $\mathbf{f}_V(T)$ on V. Thus by Theorem 6.15, we have $T' = T$. □

Remark 6.22. The original Bernardi process for graphs is most easily identified with the (ht:E, cut:V) version. In [1], Baker and Wang define a right inverse and a left inverse for the bijection between spanning trees and break divisors given by the Bernardi process on graphs. We note that the (ht:V, cut:V) Bernardi process is a generalization of their right inverse and the (ht:V, cut:E) Bernardi process with base node b_0, base edge $b_0b_1^-$, and the reversed ribbon structure is a generalization of their left inverse. Hence the theorems of this subsection extend the results of Baker and Wang. This hinges on the realization that both spanning trees and break divisors are special cases of hypertrees.

7 Shellability

Let us start by introducing a natural order on the set of Jaeger trees. As always, we assume that a ribbon structure, base node, and base edge are fixed for $\operatorname{Bip}\mathcal{H}$, where $\mathcal{H} = (V, E)$ is a connected hypergraph. Then we actually have two sets of Jaeger trees, the E-cut set and the V-cut set. We will order them both in two different ways. (The exact relationship between the two orders on the same set is somewhat unclear. In some sense they should be opposites but that is not literally true.)

Definition 7.1. Let T_1 and T_2 be V-cut Jaeger trees. Consider the last time when their (violet) tours are identical. This happens when a violet node is current[1] together

[1]It is easy to see that the current node v has to be violet, cf. the first paragraph of the proof of Lemma 6.14, but in fact the definition can be made without this piece of information.

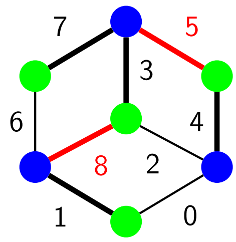

Figure 5: Internally semi-passive edges, for the given edge-order (numbering) and the spanning tree represented by thick lines, are shown in red.

with an edge ε incident to v, where, say, $\varepsilon \in T_2$ but $\varepsilon \notin T_1$. In this case we put $T_1 <_V T_2$. This defines a total order on the set of V-cut Jaeger trees. Let us call this order the *violet order of V-cut Jaeger trees.*

We also define the *emerald order of V-cut Jaeger trees.* This time $T_1 <_E T_2$ if in their emerald tours, the first edge that behaves differently is part of T_2 but not part of T_1. By symmetry of color classes, we also have *violet and emerald orders of V-cut Jaeger trees.*

The goal of this section is to show that the simplices corresponding to the V-cut Jaeger trees form a dissection of the root polytope, and moreover, the violet order of V-cut Jaeger trees translates to a shelling order for this dissection. This will be a crucial ingredient in the proof of Theorem 5.4, i.e., that our alternative definition for the interior polynomial agrees with the original one. The other key to the proof will be that the interior polynomial coincides with the h-vector of any shellable dissection of the root polytope.

The following notion is independent of ribbon structures and will greatly help to describe the terms in the h-vector. We invented it based on 'external semi-activity,' which, in turn, we heard about from Alexander Postnikov in private communication.

Definition 7.2 (internal semi-passivity)**.** Given a bipartite graph G, a total order on its edges, and a spanning tree T of G, we say that an edge $\varepsilon \in T$ is *internally semi-passive* in T, if for the smallest edge $\varepsilon' \in C^*(T, \varepsilon)$ with respect to the order of the edges, ε' "stands opposite to ε," that is, ε and ε' have endpoints from different color classes in each component of $T - \varepsilon$.

See Figure 5 for an example. 'Passive' just stands for 'not active.' In Tutte's original sense, 'internally active' means 'being smallest in the base cut.' In 'internally semi-active' that notion is weakened to 'standing parallel to the smallest element of the base cut.'

The next lemma plays an important role in the proof of shellability, as well as in relating the terms of the h-vector to the coefficients in the Bernardi-type definition of the interior polynomial. As always, we fix a ribbon structure, base node, and base edge.

Lemma 7.3. *Let T be a V-cut Jaeger tree, and $\varepsilon \in T$ be an edge. The following statements are equivalent.*

(i) *ε arises as a "first difference" between T and some Jaeger tree preceding T in the violet ordering of Jaeger trees, i.e., there exists a V-cut Jaeger tree T' such that the violet tours of T and T' coincide until reaching ε, where ε is included into T but not included into T'.*

(ii) *ε is internally semi-passive in T with respect to the emerald T-order (Definition 6.8) of the edges of* $\mathrm{Bip}\,\mathcal{H}$.

(iii) *$\varepsilon = ve$ has its violet endpoint v in the base component of $T-\varepsilon$ and e is internally inactive in $\mathbf{f}_E(T)$ with respect to the order induced on E by the emerald T-order of the edges of* $\mathrm{Bip}\,\mathcal{H}$ *(Definition 6.9).*

(iv) *ε is not the largest element in $C^*(T,\varepsilon)$ according to the violet T-order of the edges of* $\mathrm{Bip}\,\mathcal{H}$.

(v) *ε has its violet endpoint in the base component, and there exists an edge in $C^*(T,\varepsilon)$ with its emerald endpoint in the the base component.*

Proof. (i) $\Rightarrow$ (v) is straightforward from Lemma 6.14. Note that as T' is also a spanning tree of $\mathrm{Bip}\,\mathcal{H}$, it includes some edge ε' from the cut $C^*(T,\varepsilon)$, which then has the required property by the Lemma.

(v) $\Rightarrow$ (i). Let T_0 and T_1 be the two subtrees of $T-\varepsilon$, with T_0 containing the base node. Let G_0 be the subgraph of $\mathrm{Bip}\,\mathcal{H}$ spanned by the node set of T_0, and G_1 be the subgraph spanned by the node set of T_1. We build up a V-cut Jaeger tree T' together with its violet tour, such that ε is the first difference between T' and T. Follow the violet tour of T until reaching ε, but at that moment, do not include ε into T'. Instead, stay in T_0 and continue with the part of the tour of T that would be traversed after traversing ε from the emerald direction. Stop at the first moment when an edge $\varepsilon' \in C^*(T,\varepsilon)$, with its emerald endpoint e' in T_0, becomes current edge. The assumption (v) guarantees the the existence of ε'. Let $\varepsilon' = e'v'$ where v' is in the node set of G_1. Take an arbitrary V-cut Jaeger tree T_1' of G_1 with base point v' and base edge $v'e'^{+}$. Such a Jaeger tree can be constructed by running, say, the (ht:V, cut:V) Bernardi process on an arbitrary hypertree of G_1. We claim that $T' = T_0 \cup \{\varepsilon'\} \cup T_1'$ is a V-cut Jaeger tree. Once we prove this, it becomes immediate from the construction that the first difference in the violet tours of T and T' is that ε is included into T while not included into T'.

Until reaching ε, the violet tours of T and T' coincide. Then ε is cut at its violet endpoint in the tour of T'. Next we continue the traversal of T_0 until we arrive at ε'. During this time, any edges that we cut are edges that are cut in the tour of T as well. As ε' is traversed, the tour of T' thus far did not cut any edge at its emerald endpoint. Next we start the traversal of T_1', which is a V-cut Jaeger tree in G_1. Compared to the tour of T_1' with regard to G_1, the only difference in the tour of T' is that any edges in $C^*(T,\varepsilon)$ that we encounter have to be skipped. But by Lemma 6.13, all the edges from $C^*(T,\varepsilon)$ that have their emerald endpoint in G_1 have already been cut, hence none of them will be cut at its emerald endpoint. Then when we arrive back at e' after traversing ε' for the second time, the set of node-edge pairs that have not been current is the same as in the violet tour of T after $(e', e'v')$ serves as current pair. Moreover, the orders in which they become current are the same, too. Hence during the traversal of the remaining part of T_0, each edge that is cut is still cut at its violet endpoint.

(v) $\Rightarrow$ (ii). The tree T is an E-cut Jaeger-tree for the reversed ribbon structure by Lemma 6.5. Applying Lemma 6.13 with "emerald" instead of "violet," we see that those edges in $C^*(T, \varepsilon)$ that have their emerald endpoint in the base component all precede, in the emerald T-order, those which have their violet endpoint there. Since now there exists an edge in $C^*(T, \varepsilon)$ that has its emerald endpoint in the base component, necessarily the smallest edge of $C^*(T, \varepsilon)$ according to the emerald T-order also has its emerald endpoint in the base component. As ε has its violet endpoint in the base component, the smallest edge in $C^*(T, \varepsilon)$ according to the emerald T-order stands opposite to ε.

(ii) $\Rightarrow$ (iii). First we prove that if $\varepsilon = ve$ is internally semi-passive in T with respect to the emerald T-order of the edges of $\operatorname{Bip}\mathcal{H}$, then ε has its violet endpoint v in the base component. The condition means that the smallest edge $v'e'$ in $C^*(T, \varepsilon)$, according to the emerald T-order, stands opposite to ve. By Lemma 6.13, among two edges of $C^*(T, \varepsilon)$ standing opposite to each other, the smaller one according to the emerald T-order has its emerald endpoint in the base component. Hence ε has its violet endpoint in the base component.

Now we show that e is internally inactive in $\mathbf{f}_E(T)$ with respect to the emerald T-order. Since $v'e' \in C^*(T, \varepsilon)$, replacing ve with $v'e'$ in T yields another spanning tree T' of $\operatorname{Bip}\mathcal{H}$. In particular $\mathbf{f}_E(T')$ is a hypertree. As $\mathbf{f}_E(T')$ is obtained from $\mathbf{f}_E(T)$ by a transfer of valence from e to e' (letting of course e' be the emerald endpoint of $e'v'$), it is enough to show that e' precedes e in the emerald T-order. We already know that ve has its violet endpoint in the base component, implying that e is reached through ve in the emerald tour of T. From Lemma 6.13 we also know that $v'e'$ becomes current edge, with e' as current node, earlier during the emerald tour of T than the first traversal of ev. Hence any edge incident to e comes later in the emerald T-order than $v'e'$. Thus, e comes later in the emerald T-order than e'.

(iii) $\Rightarrow$ (v). As both (iii) and (v) claim that $\varepsilon = ve$ has its violet endpoint v in the base component, it is enough to prove that if e is internally inactive in $\mathbf{f}_E(T)$ with respect to the emerald T-order, then there is an edge in $C^*(T, \varepsilon)$ that has its emerald endpoint in the base component.

Suppose for contradiction that there is no edge in $C^*(T, \varepsilon)$ with its emerald endpoint in the base component T_0 of $T - \varepsilon$. Let V_0 be the set of violet nodes of T_0, and let E_0 be the set of its emerald nodes. Our assumption implies that for each $x \in E_0$, the set of edges of T incident to x coincides with the set of edges of T_0 incident to x. Moreover, all edges of $\operatorname{Bip}\mathcal{H}$ incident to x have their violet endpoint in V_0.

Since T_0 is a tree spanning $V_0 \cup E_0$, we have $\sum_{x \in E_0} \mathbf{f}_E(T)(x) = |V_0| + |E_0| - 1 - |E_0| = |V_0| - 1$. Recall that for any hypertree $\mathbf{f}$ on E, we have $\sum_{x \in E_0} \mathbf{f}(x) \le |V_0| - 1$ [8, Theorem 3.4]. In other words, with respect to $\mathbf{f}_E(T)$, the set E_0 is 'tight,' meaning that no element of it may receive a transfer of valence from outside of E_0. Now this is a contradiction because all the emerald nodes that are smaller than e in the emerald

T-order lie in E_0, which effectively blocks e from being internally inactive.

(iv) $\Rightarrow$ (v) follows directly from Lemma 6.13. Let ε' be an edge of $C^*(T,\varepsilon)$ such that $\varepsilon <_{T,V} \varepsilon'$. Then by part (ii) of Lemma 6.13, ε' must have its emerald endpoint in the base component, for otherwise it would precede ε in the violet T-order.

(v) $\Rightarrow$ (iv) also follows directly from Lemma 6.13. Let ε' be an edge in $C^*(T,\varepsilon)$ having its emerald endpoint in the base component. Since ε has its violet endpoint in the base component, by part (i) of Lemma 6.13 we have $\varepsilon \le_{T,V} \varepsilon'$, preventing ε from being the largest edge of $C^*(T,\varepsilon)$ with respect to the violet T-order. □

We can further deduce the following characterization of internally inactive emerald nodes.

Corollary 7.4. *Let T be a Jaeger tree in* $\mathrm{Bip}\,\mathcal{H}$. *An emerald node $e \in E$ is internally inactive with respect to the emerald T-order on E if and only if there is an internally semi-passive edge (with respect to the emerald T-order of the edges of* $\mathrm{Bip}\,\mathcal{H}$*) incident to it. Moreover, in this case there is a unique such edge, namely the edge through which e is reached in the emerald tour of T.*

Proof. If $e = b_0$, then e is the smallest element in the emerald T-order of E and hence it cannot be inactive. Also in this case, edges incident to b_0 have their emerald endpoint in the base component, which prevents them from being internally semi-passive by Lemma 7.3.

If $e \neq b_0$, then there is a unique edge $\varepsilon = ve$ such that e is reached through ε in the emerald T-tour. This is the only edge incident to e that has its violet endpoint in the base component of its own base cut — that is, ε is the only edge incident to e that may be internally semi-passive. Now by Lemma 7.3, ε is internally semi-passive with respect to the emerald T-order if and only if e is internally inactive with respect to the emerald T-order on E. □

Now we are in a position to explain the connection between internal activities of hypertrees and internal semi-activities of Jaeger trees.

Corollary 7.5. *For a Jaeger tree T, the number of internally semi-passive edges of T with respect to the emerald T-order (of the edges of* $\mathrm{Bip}\,\mathcal{H}$*) equals the number of internally inactive emerald nodes in $\mathbf{f}_E(T)$ with respect to the emerald T-order (of E).*

Next we turn to proving that simplices corresponding to Jaeger trees form a shellable dissection. First we show that they form a dissection.

Theorem 7.6. *For any ribbon structure (and base node and base edge) on the connected bipartite graph G, of color classes E and V, the simplices in the root polytope Q_G corresponding to the V-cut Jaeger trees form a dissection of Q_G.*

It turns out that we have already done the hard part of the proof and only the following easy part remains.

Lemma 7.7. *For any two V-cut Jaeger-trees, their corresponding maximal simplices in Q_G have disjoint interiors.*

Proof. Take two V-cut Jaeger trees T_1 and T_2, and suppose that their violet tours coincide until reaching ε, where $\varepsilon \notin T_1$ and $\varepsilon \in T_2$. It suffices to show that the maximal simplices that correspond to the pair of trees are separated by a hyperplane. Such hyperplanes may be constructed using cuts in G. This is described in [10, Sec. 3] for 'directed' cuts, when the result is a supporting hyperplane of Q_G; here we will work with undirected cuts.

Let $E_1 \sqcup E_2$ be a partition of E and $V_1 \sqcup V_2$ a partition of V. Let us associate the real number -1 to those basis vectors of $\mathbf{R}^E \oplus \mathbf{R}^V$ that correspond to elements of $E_1 \cup V_2$. Similarly, we associate 1 to elements of $E_2 \cup V_1$. This has a unique linear extension $\kappa\colon \mathbf{R}^E \oplus \mathbf{R}^V \to \mathbf{R}$. The value of κ at the vertices of Q_G (which correspond to edges of G) is then

- 0 for edges outside the cut defined by our partitions (when we split $E \cup V$ into $E_1 \cup V_1$ and $E_2 \cup V_2$), i.e., for edges between E_1 and V_1 or between E_2 and V_2,
- -2 for edges of the cut that are between E_1 and V_2, and
- 2 for edges of the cut between E_2 and V_1.

Thus the kernel of κ is a hyperplane that contains all vertices of Q_G except for the ones that correspond to edges of the cut; the remaining vertices fall on the two sides of the hyperplane depending on whether their emerald or violet endpoint lies on an arbitrarily fixed side of the cut.

Now to construct the required hyperplane, we consider the cut $C^*(T_2, \varepsilon)$. More precisely, let $E_1 \cup V_1$ be the node set of the connected component of $T_2 - \varepsilon$ which contains the violet endpoint v of ε. Let $E_2 \cup V_2$ be the node set of the other component. Then the corresponding functional κ takes only non-negative values at the vertices of the maximal simplex defined by T_2: the value is 2 for (the vertex corresponding to) ε and 0 for all other edges of G/vertices of Q_G.

We claim that the opposite is true for T_1: at the vertices of Q_G corresponding to those, all values of κ end up non-positive. To show this, we need to ascertain that all edges of $C^*(T_2, \varepsilon) \cap T_1$ are connecting E_1 and V_2. But that is exactly the statement of Lemma 6.14. □

Proof of Theorem 7.6. By Proposition 6.3, each hypertree is realized by at least one Jaeger tree and hence the number of V-cut Jaeger trees is at least the number of hypertrees in B_V. (In fact in subsection 6.3 we have showed that these numbers are equal, but the present proof does not rely on that.) By Lemma 7.7, the (maximal) simplices in Q_G corresponding to the V-cut Jaeger trees have disjoint interiors. Moreover, each maximal simplex in Q_G has the same volume [12, Section 12] and the volume of Q_G itself is the number of hypertrees times this common volume. Hence the simplices corresponding to V-cut Jaeger trees fill Q_G, in other words, they form a dissection. □

Note that this reasoning also implies that each hypertree is realized by at most one Jaeger tree, which we have already proven in Theorem 6.15.

Theorem 7.8. *The violet ordering of the V-cut Jaeger trees induces a shelling order of the dissection given in Theorem 7.6. For each V-cut Jaeger tree T, the number of facets of the corresponding simplex Q_T, that lie in the union of previous simplices of the shelling, equals the number of internally semi-passive edges in T with respect to the emerald T-order of the edges.*

In other (more precise) words, the edges of T described in five equivalent ways in Lemma 7.3 correspond exactly to those vertices of Q_T whose opposite facets make up the intersection of T with the union of the previous simplices.

Proof. Let us fix a V-cut Jaeger tree T. We have to show that for an edge $\varepsilon \in T$, if ε is internally semi-passive in T with respect to the emerald T-order, then the facet $Q_{T-\varepsilon}$ of the simplex Q_T lies in the union of the simplices corresponding to V-cut Jaeger trees preceding T in the violet order. We also need to prove that if ε is internally semi-active with respect to the emerald T-order, then the interior of $Q_{T-\varepsilon}$ is disjoint from $\bigcup_{T' <_V T} Q_{T'}$.

Take an edge $\varepsilon \in T$ such that ε is internally semi-passive in T with respect to the emerald T-order. In the same way as in the proof of Lemma 7.7, we consider the base cut and the linear functional defined by T and ε. Let us partition the edges in $C^*(T, \varepsilon)$ into the sets P and N, where elements of P have their violet endpoint in the root component of $T - \varepsilon$ and elements of N have their emerald endpoint there. (This is the same partition as the one induced by the two sides, positive and negative, of the hyperplane H that corresponds to the cut.) By Lemma 7.3 we have $\varepsilon \in P$ and $N \neq \varnothing$. Let ε' be the edge in N that is reached first among the edges of N during the violet tour of T. Let T_0 and T_1 be the two subtrees of $T - \varepsilon$, so that T_0 contains b_0. Let G_0 and G_1 be the subgraphs spanned by the node sets of T_0 and T_1, respectively. Let $\varepsilon' = xy$ such that x is in the node set of G_1. Let us denote the set of V-cut Jaeger trees of G_1 with base node x and base edge xy^+ by $\mathcal{J}_1$. Let $\mathcal{T}$ be the following set of trees:

$$\mathcal{T} = \{\, T_0 \cup \varepsilon' \cup T_1' \mid T_1' \in \mathcal{J}_1 \,\}.$$

From the part (v) $\Rightarrow$ (i) of the proof of Lemma 7.3, we know that all of these trees are V-cut Jaeger trees preceding T in the violet order of Jaeger trees. We claim that

$$Q_{T-\varepsilon} \subset \bigcup_{T' \in \mathcal{T}} Q_{T'}. \tag{7.1}$$

Take a point $\mathbf{p}$ from $Q_{T-\varepsilon}$. Then

$$\mathbf{p} = \sum_{e'v' \in T-\varepsilon} \lambda_{e'v'}(\mathbf{e}' + \mathbf{v}'),$$

where $\lambda_{e'v'} \geq 0$ and $\sum_{e'v' \in T-\varepsilon} \lambda_{e'v'} = 1$. Now as $T - \varepsilon = T_0 \cup T_1$, we have

$$\mathbf{p} = \sum_{e'v' \in T_0} \lambda_{e'v'}(\mathbf{e}' + \mathbf{v}') + \sum_{e'v' \in T_1} \lambda_{e'v'}(\mathbf{e}' + \mathbf{v}').$$

Let $\lambda_0 = \sum_{e'v' \in T_0} \lambda_{e'v'}$ and $\lambda_1 = \sum_{e'v' \in T_1} \lambda_{e'v'}$. If $\lambda_0 = 0$ then $\mathbf{p} \in Q_{G_1}$, which by Theorem 7.6 is dissected by the faces $Q_{T_1'}$ of the simplices $Q_{T'}$, making (7.1) obvious. If $\lambda_1 = 0$, then $\mathbf{p} \in Q_{T_0}$, which in turn is part of all of the $Q_{T'}$. Otherwise, write

$$\mathbf{p} = \lambda_0 \cdot \sum_{e'v' \in T_0} \frac{\lambda_{e'v'}}{\lambda_0}(\mathbf{e}' + \mathbf{v}') + \lambda_1 \cdot \sum_{e'v' \in T_1} \frac{\lambda_{e'v'}}{\lambda_1}(\mathbf{e}' + \mathbf{v}')$$

and let $\mathbf{p}_0 = \sum_{e'v' \in T_0} \frac{\lambda_{e'v'}}{\lambda_0}(\mathbf{e}' + \mathbf{v}')$ and $\mathbf{p}_1 = \sum_{e'v' \in T_1} \frac{\lambda_{e'v'}}{\lambda_1}(\mathbf{e}' + \mathbf{v}')$. That is, $\mathbf{p} = \lambda_0 \mathbf{p}_0 + \lambda_1 \mathbf{p}_1$, where $\lambda_0, \lambda_1 \geq 0$ and $\lambda_0 + \lambda_1 = 1$. Furthermore, we have $\mathbf{p}_1 \in Q_{T_1} \subset Q_{G_1}$ and $\mathbf{p}_0 \in Q_{T_0}$. Combining this with the convexity of $Q_{G_1} = \bigcup_{T_1' \in \mathcal{J}_1} Q_{T_1'}$, we obtain

$$\bigcup_{T' \in \mathcal{T}} Q_{T'} \supset \bigcup_{T' \in \mathcal{T}} Q_{T' - \varepsilon'} = \bigcup_{T_1' \in \mathcal{J}_1} \operatorname{Conv}(Q_{T_0} \cup Q_{T_1'}) = \operatorname{Conv}\left(Q_{T_0} \cup \bigcup_{T_1' \in \mathcal{J}_1} Q_{T_1'}\right)$$
$$= \operatorname{Conv}(Q_{T_0} \cup Q_{G_1}) \ni \mathbf{p}.$$

Now if ε is internally semiactive with respect to the emerald T-order, then by Lemma 7.3, ε cannot be obtained as the first difference between T and some Jaeger tree preceding T in the violet ordering of Jaeger trees. Moreover, by the proof of Lemma 7.7, for any Jaeger tree T' preceding T in the violet ordering of Jaeger trees, Q_T is separated from $Q_{T'}$ by the hyperplane containing the facet $Q_{T-\varepsilon'}$ where ε' is the first difference between T and T'. Hence the interior of $Q_{T-\varepsilon}$ is indeed disjoint from $Q_{T'}$. □

Now it is finally time to summarize our findings and prove our main theorem.

Proof of Theorem 5.4. By Theorem 3.3 and the discussion preceding it, the coefficients of the interior polynomial $I_{\mathcal{H}}$ are the entries in the h-vector of any shellable dissection (together with any shelling order) of $Q_{\operatorname{Bip}\mathcal{H}}$. Hence it suffices to show that there exists a shellable dissection of $Q_{\operatorname{Bip}\mathcal{H}}$, with a shelling order yielding the h-vector $(a_0, a_1, \dots)$, so that a_i is equal to the number of hypertrees of $\mathcal{H}$ with internal embedding inactivity i (defined using our fixed ribbon structure and the (ht:E, cut:E) Bernardi process).

We choose the dissection of $Q_{\operatorname{Bip}\mathcal{H}}$ by E-cut Jaeger trees. By Lemma 6.6, these are the same as the V-cut Jaeger trees of the *reversed* ribbon structure (with base node b_0 and base edge $b_0 b_1^-$). The *violet* order of Jaeger trees will be our shelling order, as in Theorem 7.8.

By Remark 6.11 (whose statements we already *can* prove because we have Theorem 6.17 in hand) and Corollary 7.5, a hypertree on E has internal embedding inactivity i with respect to the (ht:E, cut:E) Bernardi process if and only if the E-cut Jaeger tree T realizing it (i.e., the outcome of the process) has i internally semi-passive edges with respect to the emerald T-order. Now the statement follows from the second sentence in Theorem 7.8. □

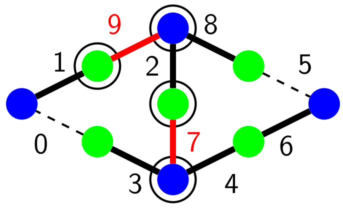

Figure 6: The ribbon structure is the positive orientation of the plane at each node. The thick edges form the Jaeger tree T. The numbers on the edges show the violet T-order. The circled violet and emerald nodes are the internally inactive elements with respect to the order induced by the violet T-order on the violet and emerald nodes, respectively. The red edges are the internally semi-passive edges in T with respect to the violet T-order.

8 The Bernardi bijection is activity-preserving for graphs

Let us now take a look at the case of ordinary graphs $G = (V, E)$, i.e., hypergraphs so that $d(e) = 2$ for each $e \in E$. We claim that in this case, the Bernardi process gives an internal-activity-preserving bijection between the hypertrees on E and on V. This implies Conjecture 5.5 in the case of graphs. The property however does not remain true for all hypergraphs and thus cannot be used to prove the Conjecture in general. The precise statement is as follows.

Theorem 8.1. *For a ribbon graph G, and a V-cut Jaeger tree T of $\mathrm{Bip}\, G$, the internal inactivity of $\mathbf{f}_E(T)$ with respect to the violet T-order on E equals the internal inactivity of $\mathbf{f}_V(T)$ with respect to the violet T-order on V.*

We also have the following refinement. See Figure 6 for an illustration.

Theorem 8.2. *For a graph G, and a V-cut Jaeger tree T on $\mathrm{Bip}\, G$, the internally semi-passive edges with respect to the violet T-order of the edges of $\mathrm{Bip}\, G$ give a matching between the internally inactive elements of V in $\mathbf{f}_V(T)$ with respect to the violet T-order on V and the internally inactive elements of E in $\mathbf{f}_E(T)$ with respect to the violet T-order on E.*

Proof. Let T be a V-cut Jaeger tree and recall (Lemma 6.5) that T is an E-cut Jaeger tree with respect to the reversed ribbon structure. Throughout the proof we will often use Lemma 7.3 with interchanging the roles of the colors violet and emerald.

By Corollary 7.4, a violet node v is internally inactive in $\mathbf{f}_V(T)$ with respect to the violet T-order on V if and only if there is an edge ε incident to it which is internally semi-passive with respect to the violet T-order, moreover, in this case there is exactly one such edge.

For an emerald node $e \in E$, if $e \neq b_0$, then T has at most one edge incident to e that has its emerald endpoint in the base component, since for any $e \in E$ there are at most two incident edges in T, and if e is not the base point, then one of them necessarily connects e to the base component. Hence the latter edge has its violet endpoint in the base component.

Now it is enough to show that an emerald node $e \in E$ is internally inactive in $\mathbf{f}_E(T)$ with respect to the violet T-order on E if and only if $e \neq b_0$ and there is an edge of T incident to it that has its emerald endpoint in the base component, and is internally semi-passive with respect to the violet T-order.

Lemma 8.3. *If an emerald node $e \in E$ is internally inactive in $\mathbf{f}_E(T)$ with respect to the violet T-order, then $e \neq b_0$ and there is an edge $\varepsilon \in T$ incident to e that has its emerald endpoint in the base component, and ε is internally semi-passive in T with respect to the violet T-order.*

Proof. If e is internally inactive in $\mathbf{f}_E(T)$ with respect to the violet T-order, then $\mathbf{f}_E(T)(e) > 0$, i.e., both edges of $\operatorname{Bip} G$ incident to e are in T. Moreover, there is an emerald node $e' \in E$ such that e' comes before e in the violet T-order, and $\mathbf{f}_E(T)$ is such that a transfer of valence is possible from e to e'. Hence necessarily $\mathbf{f}_E(T)(e') = 0$, that is, T contains only one edge incident to e'. Let $v'e'$ be the edge not in T. As T is a V-cut Jaeger tree and $v'e' \notin T$, it necessarily has its violet endpoint in the base component.

As the hypertrees on E in $\operatorname{Bip} G$ determine their realizing spanning trees at nodes $\tilde{e} \in E$ with $\mathbf{f}(\tilde{e}) = 1$, the transfer of valence from e to e' can be achieved by removing an edge incident to e from T, and adding $v'e'$. Hence, one of the edges incident to e is in the base cycle $C(T, v'e')$. As the degree of e is two, this means that both edges incident to e must be part of that base cycle. Hence $v'e'$ is in the base cut of both edges incident to e in $\operatorname{Bip} G$. One of these edges has its emerald endpoint in the base component. Now by Lemma 7.3, this edge is internally semi-passive with respect to the violet T-order. (To apply the lemma note that a V-cut Jaeger tree is also an E-cut Jaeger tree with the reversed ribbon structure.)

Also, by Lemma 6.13, $v'e'$ comes earlier in the violet T-tour than either edge incident to e. Thus, by Lemma 6.12, in the violet tour of T, the edge $v'e'$ becomes current with current vertex v earlier than any edge incident to e with current vertex e. Hence we cannot have $e = b_0$ as that would mean that one of the edges incident to e is traversed first. ☐

Lemma 8.4. *If an edge $\varepsilon = ve$ is internally semi-passive in T with respect to the violet T-order, then e is internally inactive in $\mathbf{f}_E(T)$ with respect to the violet T-order on E.*

Proof. By Lemma 7.3, if ε is internally semi-passive, then it has its emerald endpoint in the base component, and there exists an edge $\varepsilon' = e'v' \in C^*(T, \varepsilon)$ that has its violet endpoint in the base component of $T - \varepsilon$. By Lemma 6.13 and Lemma 6.12, ε' is deleted before the first traversal of ε. Since there are only two edges in $\operatorname{Bip} G$ incident to e, the node e receives its number in the violet T-order immediately before the first traversal of ε. Thus, e' precedes e in the violet T-order on E. Since $v'e' \in C^*(T, ve)$, the set of edges $T' = T - ve + v'e'$ is also a spanning tree of $\operatorname{Bip} G$, inducing the hypertree $\mathbf{f}_E(T')$. We get $\mathbf{f}_E(T')$ by decreasing the value at e by one, and increasing the value at e' by one. By this we have proved that there is a transfer of valence from e to an emerald node preceding e in the violet T-order, hence e is internally inactive in $\mathbf{f}_E(T)$ with respect to the violet T-order. □

The two Lemmas together imply the Theorem. □

Theorem 8.2 is not true for gereral hypergraphs, as in general, the internally semi-passive edges of a Jaeger tree do not necessarily form a matching.

9 Examples

In this section we recall two constructions for triangulating a root polytope and show that they occur as special cases of the dissections of Section 6. But before that, let us work out two concrete examples where the dissection fails to be a triangulation.

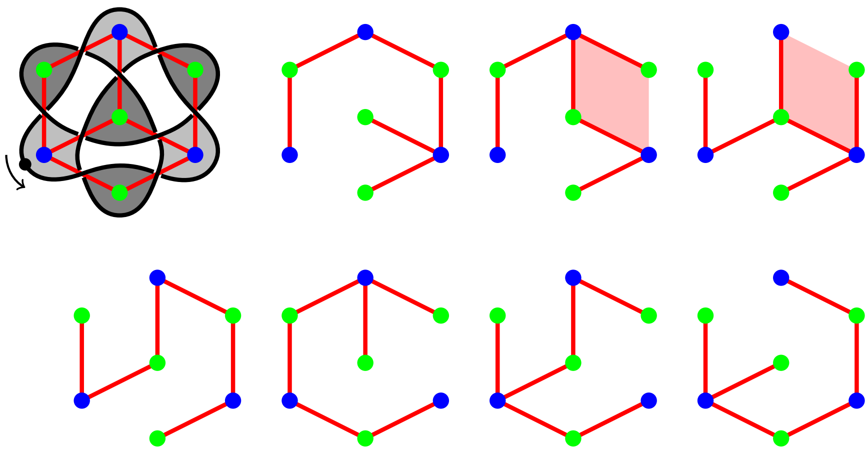

Figure 7: A plane bipartite graph with a ribbon structure that rotates in *opposite* directions at its two color classes. The V-cut Jaeger trees corresponding to the indicated base point are listed in the associated shelling order.

Example 9.1. Let us consider the plane bipartite graph of Figure 4 again, with the same base node and base edge as before, but this time let us use a ribbon structure very similar to that of Figure 2. That is, we use counterclockwise rotations about violet nodes and clockwise ones about emerald nodes. (This is the same rule as in Example 4.3, except that the colors traded places.) In Figure 7 we show again an embedding in $\mathbf{R}^3$ of the corresponding ribbon surface with the base point along its boundary. The seven spanning trees of the Figure are the resulting V-cut Jaeger trees, listed in the shelling order of Section 7. One may check that (since they form a dissection of the root polytope) the trees do realize all hypertrees on both color classes. On the other hand, they do not form a triangulation of the root polytope. For example, the second and third trees are such that, with regard to the upper right quadrangular region of the embedding, one tree contains one pair of opposite edges and the other tree contains the other pair. That violates Postnikov's compatibility condition [12, Lemma 12.6], i.e., the simplices corresponding to the two trees do not intersect in a common face.

The bipartite graph of the previous example had points of degree at least three in both color classes, in other words it was not of the form Bip G for any graph G. But that was not the reason for its 'bad behavior;' rather, the ribbon structure was. In the next example we show that the dissection can fail to be a triangulation even in Bernardi's original family of cases.

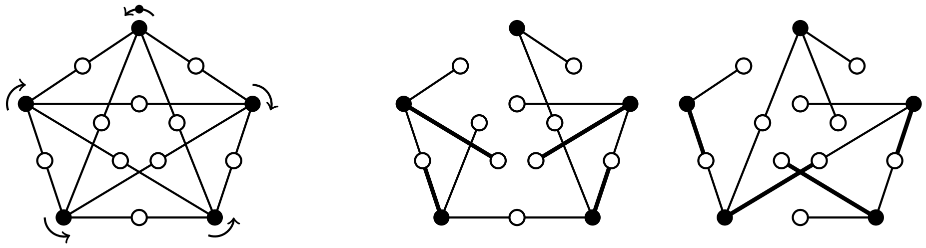

Figure 8: A ribbon structure for the complete graph K_5 and two of its Jaeger trees.

Example 9.2. Let us consider the complete graph K_5 and its ribbon structure indicated in Figure 8. That is, we refer to the planar drawing shown and let edges be ordered clockwise around two of the vertices and counterclockwise around the other three. This extends uniquely to the degree two nodes of Bip K_5, as discussed in Section 4. For K_5, let the base vertex be the one on top, and let the top left edge be the base edge. Equivalently, for Bip K_5, let the base node be the same point and let the base edge be the upper half of the previous.

Now if cutting edges of Bip K_5 is allowed near the five vertices of K_5, then the two spanning trees in Figure 8 are Jaeger trees. The simplices in $Q_{\mathrm{Bip}\,K_5}$ that correspond to these trees do not intersect in a common face: indeed, the thickened edges (four from each tree) form a cycle in Bip K_5 and that violates the condition in [12, Lemma 12.6].

Next, let us consider the 'triangulation by non-crossing trees' [6], see also [10, Example 5.3]. This applies in the case of a complete bipartite graph K, say on $(m+1)+(n+1)$ vertices, whose root polytope is the product $Q_K = \Delta_m \times \Delta_n$ of an m- and an n-dimensional unit simplex. The idea is to draw the vertices of K on two parallel lines in the plane, separated by color, and then consider those spanning trees whose edges do not cross each other in the drawing. (See Figure 9 for an example.) The maximal simplices corresponding to these form a triangulation of Q_K.

Now let us imagine the two lines as horizontal, with emerald vertices on the lower one and violet vertices on the upper. Let us define a ribbon structure by rotating counterclockwise around each vertex. Our base node will be the lower left (emerald) one and the base edge will be the one connecting the base node diagonally to the upper right violet node. (I.e., in the sense of Remark 2.1, we place our base point slightly below the lower left vertex.) Then it is easy to see that non-crossing trees

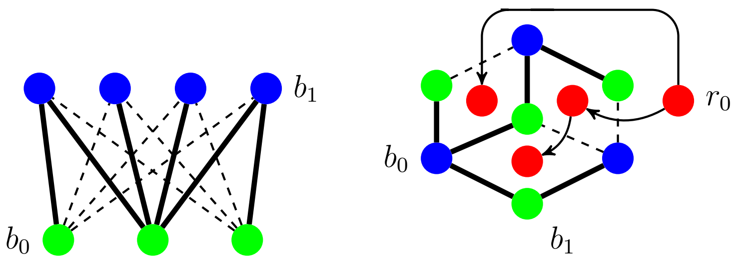

Figure 9: Left: A non-crossing tree in $K_{3,4}$. Right: A V-cut Jaeger tree in a plane bipartite graph and the corresponding arborescence of the dual graph.

are E-cut Jaeger trees. The converse is not hard to check either but since we are not aware of an elegant proof, we just let it follow from the fact that all dissections of a root polytope consist of the same number of maximal simplices.

Our final class of triangulations was introduced in [9] and it applies whenever the connected bipartite graph G comes with an embedding in the plane. (That is, unlike in the previous situation, edges of G are not allowed to cross.) In that case, the planar dual graph G^* has a natural orientation by the rule that every edge of G^* is oriented so that the violet endpoint of the corresponding edge of G is on its left. (See Figure 9.) Let us fix a vertex r_0 of G^*. Then the spanning trees of G dual to spanning arborescences of G^* rooted at r_0 form a triangulation of Q_G. (Here an arborescence is a tree so that all of its edges point away from the root. A spanning tree of G is dual to a spanning arborescence of G^* if it contains precisely the edges corresponding to the non-edges of the arborescence.)

We claim that the trees given above are exactly the V-cut Jaeger trees of G, where the ribbon structure is induced by the positive orientation of the plane (so that, contrary to Example 9.1, both colors spin in the *same* direction). To be more specific, let the base node b_0 be violet and incident to the face r_0, and let the base edge b_0b_1 be also incident to r_0 so that when we orient it from b_0 to b_1, then r_0 is on the right.

To justify our claim it is enough to show that the dual A^* of each arborescence $A \subset G^*$ is a V-cut Jaeger tree of G. Tracing the boundary of a neighborhood of A or of A^* are equivalent. (The common base point should be thought of as an interior point of the arc connecting b_0 and r_0.) When we do the latter, because A is an arborescence, we first walk along each edge of A in the direction of its orientation and on its left side. But by definition, this means that the corresponding dual edge of G is cut at its violet endpoint.

Conversely, each V-cut Jaeger tree is the dual of an arborescence rooted at r_0. Indeed, consider the tour of a V-cut Jaeger tree T. Notice that until cutting the first edge, we move along the boundary of the face r_0, with r_0 on the right. When we first cut an edge ve, we start to see a new face r on the right, and r_0r is an oriented edge of G^* which 'starts' the dual arborescence. Continuing this argument inductively, we see that cutting an edge at its violet endpoint corresponds to inserting a new edge pointing away from r_0 in the dual graph.

This in particular implies that Jaeger trees of a plane bipartite graph with the above, perhaps most natural, ribbon structure, form a shellable triangulation (and not just a dissection) of the root polytope. Let us point out a nice consequence of this property.

Proposition 9.3. *If G is a plane bipartite graph and T and T' are V-cut Jaeger trees of G for the ribbon structure described above (all nodes spinning in the same direction), then there exist edges $\varepsilon_1, \dots, \varepsilon_r$ and $\varepsilon'_1, \dots, \varepsilon'_r$ of G such that $T_i = (T \setminus \{\, \varepsilon_1, \dots, \varepsilon_i \,\}) \cup \{\, \varepsilon'_1, \dots, \varepsilon'_i \,\}$ is a Jaeger tree for each i and $T' = T_r$. (Edges might occur multiple times in the sequences.)*

Proof. Fix a shelling order $<$ of the triangulation. By the definition of shelling, the intersection of the simplex Q_T with $\bigcup_{\Gamma<T} Q_\Gamma$ is a union of facets. Let us choose a facet $Q_{T-\varepsilon_1}$ of Q_T that lies in the intersection. As we have a triangulation, $Q_{T-\varepsilon_1}$ is also a face of Q_{T_1} for some Jaeger tree $T_1 < T$. Hence in this case $T - \varepsilon_1 \subseteq T_1$, that is, $T_1 = (T \setminus \{\varepsilon_1\}) \cup \{\varepsilon'_1\}$ for some edge ε'_1. We can further apply this argument to obtain a sequence of Jaeger trees $T > T_1 > \cdots > T_k$ where T_k is the first Jaeger tree in the shelling order, and we get each Jaeger tree from the previous one by switching just one edge.

We can apply the same procedure for T'. The chains of trees for T and for T' have a common element at some point (if no other then T_k). Concatenating the two chains gives us the required sequence of edge switches. □

References

[1] M. Baker and Y. Wang, *The Bernardi process and torsor structures on spanning trees*, International Mathematics Research Notices (2017) doi: 10.1093/imrn/rnx037

[2] O. Bernardi, *A characterization of the Tutte polynomial via combinatorial embedding*, Ann. Combin. 12 (2008), no. 2, 139–153.

[3] O. Bernardi, *Tutte Polynomial, Subgraphs, Orientations and Sandpile Model: New Connections via Embeddings*, Electron. J. Combin. 15 (2008), no. 1, R109

[4] 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.

[5] A. Frank and T. Kálmán, *Root polytopes and strong orientations of bipartite graphs*, in preparation.

[6] 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.

[7] F. Jaeger, *A Combinatorial Model for the Homfly Polynomial*, Europ. J. Combin. 11 (1990), no. 6, 549–558.

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

[9] T. Kálmán and H. Murakami, *Root polytopes, parking functions, and the HOMFLY polynomial*, Quantum Topol. 8 (2017), no. 2, 205–248.

[10] T. Kálmán and A. Postnikov, *Root polytopes, Tutte polynomials, and a duality theorem for bipartite graphs*, Proc. London Math. Soc. 114 (2017), no. 3, 561–588.

[11] K. Kato, *Interior polynomial for signed bipartite graphs and the HOMFLY polynomial*, arXiv:1705.05063.

[12] A. Postnikov, *Permutohedra, Associahedra, and Beyond*, Int. Math. Res. Not. 2009, no. 6, 1026–1106.

[13] A. Schrijver, *Combinatorial Optimization - Polyhedra and Efficiency*, Springer-Verlag, Berlin, 2003

[14] W. T. Tutte, *A contribution to the theory of chromatic polynomials*, Canadian J. Math. 6 (1954), 80–91.