

Root polytopes, parking functions, and the HOMFLY polynomial

Tamás Kálmán and Hitoshi Murakami

Abstract. We show that for a special alternating link diagram, the following three polynomials are essentially the same: a) the part of the HOMFLY polynomial that corresponds to the leading term in the Alexander polynomial; b) the h -vector for a triangulation of the root polytope of the Seifert graph and c) the enumerator of parking functions for the planar dual of the Seifert graph. These observations yield formulas for the maximal z -degree part of the HOMFLY polynomial of an arbitrary homogeneous link as well. Our result is part of a program aimed at reading HOMFLY coefficients out of Floer homology.

Mathematics Subject Classification (2010). 57M27, 57M25, 57M50.

Keywords. Alternating link, HOMFLY polynomial, Seifert graph, root polytope, h -vector, parking function.

1. Introduction

In this paper we report on a new kind of combinatorial phenomenon in knot theory. Our research was motivated by a desire to understand what the HOMFLY polynomial of an oriented link ‘measures,’ i.e., to find a natural (for example, diagram-independent) definition for it. While that problem remains wide open (with the most promising approach being the Gopakumar–Ooguri–Vafa conjecture [3], [13]), we feel we did carry out an interesting case study with some surprising results.

The HOMFLY polynomial $P(v, z)$ [6] is an invariant of oriented links that specializes to the Conway polynomial $\nabla(z)$ via the substitution $\nabla(z) = P(1, z)$. The latter is equivalent to the Alexander polynomial $\Delta(t)$ through $\Delta(t) = \nabla(t^{1/2} - t^{-1/2})$. Note that ∇ and Δ share the same leading coefficient. In P , on the other hand, one finds several terms that contribute to the leading monomial of ∇ when we set $v = 1$. We will collectively refer to these as the *top* of the HOMFLY polynomial. For homogeneous links [2] (which include all alternating and positive

links), the top can also be described¹ as the sum of those terms that realize the z -degree of P .

A third, perhaps most useful description first associates the coefficient $T(v)$ of z^{n-s+1} in P to an arbitrary link diagram, where n is the number of crossings and s is the number of Seifert circles, respectively. (Here we suppress the common z -power and write only the v -powers.) In general, $T(v)$ may be the empty sum, i.e., 0 (cf. Theorem 3.2), but not for homogeneous diagrams. For those, $T(v)$ provides an equivalent definition of the top of the HOMFLY polynomial of the link. (In particular, $n-s$ is an invariant of the homogeneous link [2].) We recall that *Seifert circles* are the simple closed curves that result when we smooth every crossing of a link diagram in the orientation-preserving way. Seifert circles are the vertices of the *Seifert graph*, in which there is an edge between two of them for every crossing where they meet. A Seifert graph is always bipartite, which is the reason why it can be used in the standard construction of an *oriented* spanning surface for a link.

Let us briefly recall some more terminology. A *block sum* of two connected graphs is a one-point union and under this, each connected graph has a unique decomposition into irreducible (a.k.a. 2-connected) ones, called its *blocks*. The *median construction* on a plane graph² gives a link (diagram) as the boundary of a surface, which in turn is obtained by taking a small neighborhood of the embedding and twisting it (using one of two choices) once over each edge. A priori the link is unoriented. A *special alternating link* is an oriented link that has an alternating diagram obtained by first applying the median construction to a connected plane bipartite graph, and then orienting the link by letting the two color classes spin in opposite directions (cf. [7, Figure 3]). See Figure 2 for an example. Note that alternatingness and connectedness imply that one choice of twisting direction determines all others. Furthermore, the orientation is such that all crossings share the same sign, i.e., special alternating links are either positive or negative. (Here we use the standard convention of \nearrow being a positive crossing and \searrow a negative one.) Moreover, the bipartite graph is recovered as the Seifert graph of the diagram.

A special alternating link is a *Murasugi atom* if the graph in the description above is 2-connected. An oriented link is called *homogeneous* if it has a diagram D that decomposes as a *star product* (a.k.a. *Murasugi sum*) of Murasugi atoms [2]. On the level of Seifert graphs, a star product is a block sum. When constructing the star product, first we draw the circle C that corresponds to the merging vertices and then reproduce the two diagrams either on the same side of C (in which case the result is a connected sum) or on opposite sides (when we may or may not get a connected sum). Note the additional choices that are involved with the latter case. Figure 1 shows two of several possible star products of the same factors.

When we decompose special alternating links, all star products follow the first

¹This second definition is probably the most appropriate for the top of P of an arbitrary link. We will remain vague on this because in this paper we only consider homogeneous links.

²A *plane graph* is an isotopy class of embeddings of a graph into the plane \mathbb{R}^2 or the sphere S^2 . A *planar graph* is an abstract graph that admits such an embedding.

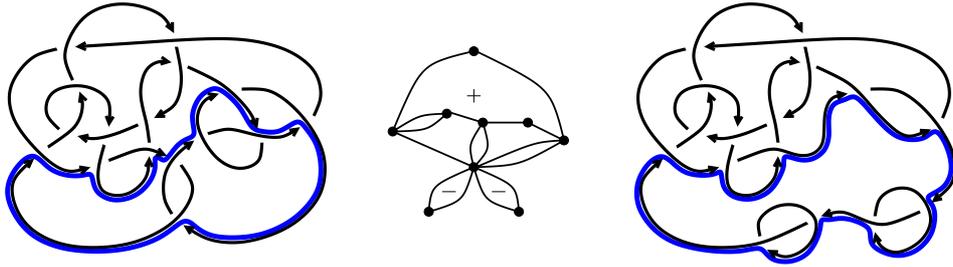


Figure 1. Two alternating links with the same Seifert graph. The signs mean the signs of the crossings in the Murasugi atoms.

(‘same side’) pattern. This is because in the other pattern, being alternating forces opposite signs for crossings on opposing sides of C . I.e., special alternating links are connected sums of Murasugi atoms of uniform sign. (For example, the granny knot is the connected sum of two positive trefoils).

Let D be a homogeneous link diagram. The aim of this paper is to describe the top of the associated HOMFLY polynomial P_D in terms of the Seifert graph of D . A theorem of Murasugi and Przytycki [12] says that the top of the HOMFLY polynomial (more precisely, the polynomial $T(v)$ of our third description above) behaves multiplicatively under star product. Thus it suffices for us to describe the top of P for Murasugi atoms; we will in fact do so for special alternating links. As for any link L and its mirror image L^* we have $P_{L^*}(v, z) = P_L(-v^{-1}, z)$, we may, without loss of generality, concentrate on positive special alternating links only.

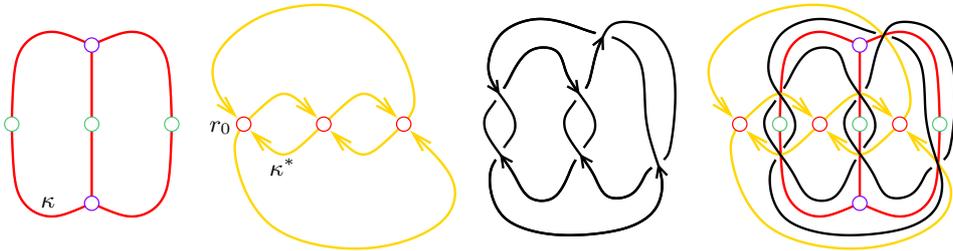


Figure 2. A plane bipartite graph G , its planar dual G^* , and associated special alternating link L_G . The last panel shows the three objects together.

Hence, much of our discussion focuses on a connected plane bipartite graph G which gives rise (via the median construction) to the positive special alternating link L_G . Alexander Postnikov [14] has recently developed a beautiful theory of (not necessarily planar) bipartite graphs. The gist of this paper is the realization that some of his ideas are closely related to knot theory.

Postnikov associates a *root polytope* Q_G to any bipartite graph G , constructed

as follows. Denote the color classes of G by E and³ V and take the convex hull, in $\mathbb{R}^E \oplus \mathbb{R}^V$, of the vectors $\mathbf{e} + \mathbf{v}$ for all edges ev of G . Here $\mathbf{e} \in \mathbb{R}^E$ and $\mathbf{v} \in \mathbb{R}^V$ are the standard generators associated to $e \in E$ and $v \in V$, respectively. If G is connected then the result is an $(|E| + |V| - 2)$ -dimensional polytope so that, of course, edges of G translate to vertices of Q_G . It is also not hard to show that a set of vertices of Q_G is affinely independent if and only if the corresponding edges form a cycle-free subgraph of G . In particular, there is a one-to-one correspondence between maximal simplices formed by vertices of Q_G and spanning trees of G .

A *triangulation* of Q_G is a collection of maximal simplices so that their union is the entire root polytope and any two of them intersect in a common face. (Note how it is not allowed to introduce new vertices when we triangulate a polytope.)

Our first result describes a specific triangulation of Q_G in the case when G is a plane graph. Let us orient the dual graph G^* so that each of its edges has an element of E to the right and an element of V to the left. After fixing a root r_0 (a vertex of G^*) arbitrarily, we consider *spanning arborescences* rooted at r_0 . These are those spanning trees of G^* in which each edge points away from r_0 . Each spanning arborescence has a dual spanning tree in G and we claim the following.

Theorem 1.1. *Let G be a connected plane bipartite graph. Fix a root r_0 and consider all spanning arborescences of G^* rooted at r_0 , as well as the spanning trees of G dual to them. Then, the collection of those simplices in the root polytope Q_G that correspond to the latter forms a triangulation of Q_G .*

A concrete example of this phenomenon is shown in Figure 5.

A triangulation of a polytope is an instance of a pure simplicial complex, i.e., one in which all maximal simplices have the same dimension. To any d -dimensional simplicial complex, it is customary to associate the *f-vector*⁴

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

where f_k , for $k \geq 0$, is the number of k -dimensional simplices in the complex. The *h-vector* of the same complex is defined as $h(x) = f(x-1)$. The latter notion becomes significant (for example, it is guaranteed to have positive coefficients) for so-called *shellable complexes*, that is complexes with a shelling order. Here a *shelling order* of a pure simplicial complex, $\sigma_1 < \sigma_2 < \cdots < \sigma_{f_d}$, lists the maximal simplices in such a way that each σ_i , $i \geq 1$, intersects the set $\sigma_1 \cup \cdots \cup \sigma_{i-1}$ in a union of c_i codimension one faces. We always have $c_1 = 0$ but assume as part of the definition that $c_i \geq 1$ for $i \geq 2$. Whether such an order exists is a subtle question, but when it does, it is not hard to show [16] that

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

³The unusual choice of symbols is motivated by hypergraph considerations, see after Corollary 1.5. For the edge set of G , we will either write C or use G itself as a stand-in.

⁴The *f-vector*, as well as the *h-vector* below, may be more appropriately called a polynomial. However the vector terminology is so common in combinatorics that we decided to keep it.

The HOMFLY polynomial, like most knot polynomials, is usually computed from a diagram of the link via successive applications of a skein relation. The process is captured by a so-called *computation tree*. The nodes of the tree are link diagrams, with the original diagram playing the role of root. Edges in the tree correspond to simple local modifications of the diagrams. For us, the relevant skein relation is

$$P_{\searrow} = v^2 P_{\nearrow} + vz P_{\smile}, \quad (2)$$

coupled with the initial condition $P_{\bigcirc} = 1$ for the HOMFLY polynomial of the unknot. Thus the computation tree is a rooted binary tree in which the two descendants of a non-leaf node result from either *changing* or *smoothing* a crossing. A priori, the crossings that we operate on can be chosen quite freely. The only restriction is that the leaves of the tree should be diagrams of the unknot, or of other links whose HOMFLY polynomials are known.

Now, the main idea of the paper is to use the spanning arborescences above to construct a computation tree \mathcal{T} for P_{LG} . Smoothing a crossing is equivalent to removing the corresponding edge from the Seifert graph. We can also keep track of crossing changes by, say, making the corresponding edge dotted. Thus each vertex of the tree will be described by a subgraph of G with some dotted edges. To build \mathcal{T} , first we will use a backtrack algorithm to enumerate all arborescences (including the non-spanning ones) of G^* . The subgraphs giving rise to the nodes of \mathcal{T} will be their duals and the tree structure (as well as the dotted edges) will reflect the steps in the algorithm. See Figures 3 and 4 for an example.

The leaves of the tree \mathcal{T} arise from two kinds of subgraph: either a spanning tree of G , which corresponds to an unknot diagram, or a subgraph so that along its ‘outside contour’ (if the root of G^* is placed in the ‘outside region’ of G), every other edge is dotted. We will not compute the HOMFLY polynomials associated to the latter (so as far as the ‘full’ HOMFLY polynomial is concerned, \mathcal{T} remains incomplete), but we will prove that these leaves do not contribute to the top of P_{LG} . By contrast, the spanning trees only contribute to the top. In fact, each gives a single monomial in which the exponent of v is determined by the number of dotted edges.

The tree \mathcal{T} has a natural embedding in the plane by always drawing, as we work downward from the root, the result of smoothing to the right and the result of crossing change to the left. In particular, the leaves of \mathcal{T} have a natural order from right to left. As to the leaves that belong to spanning trees,

- (i) by Theorem 1.1, they correspond to the maximal simplices in a triangulation of Q_G and
- (ii) we claim that the right-to-left order is a shelling order for the triangulation.

See Figure 4 for an illustration. From this, the following is immediate.

Proposition 1.2. *The triangulation described in Theorem 1.1 is shellable.*

Furthermore, it turns out that as we build Q_G simplex-by-simplex using the shelling order, the number c_i of facets along which the simplex σ_i is attached is

exactly the number of dotted edges in the corresponding tree. From this our first main result follows:

Theorem 1.3. *For any connected plane bipartite graph G with s vertices and n edges, the top of the HOMFLY polynomial $P_{L_G}(v, z)$ (of the positive special alternating link L_G) is*

$$v^{n+s-1} h(v^{-2}),$$

where h is the h -vector of the triangulation of the root polytope Q_G in Theorem 1.1.

In general, different triangulations of the same polytope can have different h -vectors. Even the number of maximal simplices, i.e., the sum of the coefficients in the h -vector, may vary. That can not occur for a root polytope because its maximal simplices share the same volume [14]. In fact, much more is true: any two triangulations of Q_G have the same f -vector and hence the same h -vector [9]. (The proof is not hard. The theorems of this paper do not depend on it.)

Our other main theorem gives a third description of the two quantities that are equated in Theorem 1.3. It is given in terms of *parking functions* associated to G^* , as defined by Postnikov and Shapiro [15]. Parking functions are also known as superstable chip-firing configurations, see for example [4]. Having fixed a root r_0 in G^* , parking functions are of the form $R \setminus \{r_0\} \rightarrow \mathbb{N} = \{0, 1, 2, \dots\}$, where R is the vertex set of G^* . See Definition 6.1 for the details. Let us associate the *index*

$$i(\pi) = \sum_{r \in R \setminus \{r_0\}} \pi(r) \tag{3}$$

to the parking function π . Let $\Pi = \Pi(G^*, r_0)$ denote the set of parking functions (which is easily seen to be finite), and let the *parking function enumerator* be

$$p(u) = \sum_{\pi \in \Pi} u^{i(\pi)}. \tag{4}$$

We find the following connection between knot theory and the chip-firing model.

Theorem 1.4. *For any connected plane bipartite graph G with s vertices and n edges, the top of the HOMFLY polynomial $P_{L_G}(v, z)$ is equal to*

$$v^{n-s+1} p(v^2),$$

where p is the parking function enumerator of the directed graph G^* .

A consequence of the Theorem is the (known) fact that G^* is such that $p(u)$ is independent of the choice of r_0 . As a byproduct of our arguments, we obtain a bijection between spanning arborescences of G^* rooted at r_0 and parking functions defined on $R \setminus \{r_0\}$. It has been known [15] that those two sets have the same cardinality. Our bijection is similar to, but appears not to be a special case of, those in the literature [1], [4]. We also obtain

Corollary 1.5. *Let G be a connected plane bipartite graph on s vertices with planar dual G^* . Then the h -vector $h(x)$ of any triangulation of the root polytope Q_G and the parking function enumerator $p(u)$ of G^* satisfy*

$$u^{s-1} h(u^{-1}) = p(u).$$

The combinatorial setup used in this paper yields yet another description of the top of P_{L_G} , this time as the *interior polynomial* [8] of the hypergraph (V, E) . This last claim is only conjecturally true⁵, but once it is proved, it will provide a hitherto unknown connection between the HOMFLY polynomial and Floer homology (cf. [7, Conjecture on p. 4]). Namely, the interior polynomial can be computed from the so-called hypertree polytope (see [8] for definitions) and the latter can be thought of as the Floer homology of a certain sutured manifold⁶ [7]. Due to Theorem 1.3, the problem of reading (some) HOMFLY coefficients out of Floer homology is reduced to the following conjecture of Postnikov and the first author:

Conjecture 1.6. *The h -vector h of a triangulation of the root polytope Q_G of the connected bipartite graph G is equivalent to the interior polynomial I of the hypergraph (V, E) , where V and E are the color classes of G . Namely, we have $u^{|E|+|V|-1} h(u^{-1}) = I_{(V,E)}(u)$.*

The sutured manifold mentioned above is the complement of a Seifert surface. However this surface is bounded not by L_G but by a related link. One may wish to consider instead the minimum genus Seifert surface F_G for L_G (produced by the median construction) that deformation retracts to G . The sutured Floer homology S_G of the complement of F_G is also a hypertree polytope (in the sense of footnote 6) but of the wrong hypergraph, whose interior polynomial is different from the top of P_{L_G} . On the other hand, the set $\Pi(G^*)$ of parking functions may be thought of as a rearrangement⁷ of S_G and thus Theorem 1.4 also becomes a way of obtaining information on the HOMFLY polynomial from Floer homology.

We end the introduction with a statement of our main result for homogeneous links. As explained above, this follows directly from our other claims via Murasugi–Przytyczki’s product formula.

Theorem 1.7. *Let D be a homogeneous link diagram that is composed of k positive and l negative Murasugi atoms. Let $p_i(v)$, $1 \leq i \leq k$ and $p'_j(v)$, $1 \leq j \leq l$ be the parking function enumerators of the dual of the Seifert graph of each atom. By Corollary 1.5, these polynomials can also be interpreted as h -vectors.*

Now if the Seifert graph G of D has altogether s_+ vertices and n_+ edges in its positive blocks and s_- vertices and n_- edges in its negative blocks (here we

⁵Note added in revision: a proof is now available as a preprint [9]. It proceeds as described in this paragraph, by settling Conjecture 1.6.

⁶In order to be more precise, let us note that the sutured manifold in question is a handlebody, so that the set of Spin^c structures supporting the sutured Floer homology lies in an affine \mathbb{Z} -lattice. This set is isomorphic to the set of lattice points in the hypertree polytope. Furthermore, at each Spin^c structure in the support, the homology group is isomorphic to \mathbb{Z} .

⁷We plan to clarify the meaning of this claim in a future joint paper with Dylan Thurston. Here let us only note that S_G and $\Pi(G^*)$ both have dimension $|R| - 1$.

have $n_+ + n_- = n$ but vertices where blocks are attached are counted once for each block to which they belong, whence $s = s_+ + s_- - k - l + 1$) so that the writhe of D is $w(D) = n_+ - n_-$, then the coefficient of z^{n-s+1} in the HOMFLY polynomial $P_D(v, z)$ (which is the highest power of z occurring in P) is

$$(-1)^{n-s+l} \cdot v^{w(D)-s_++s_-+k-l} \cdot \prod_{i=1}^k p(v^2) \cdot \prod_{j=1}^l p'(v^{-2}).$$

The paper is organized as follows. We start with material on arborescences and arrange them in a binary tree in Section 2. In Section 3, we use the binary tree to compute the top of the HOMFLY polynomial of a special alternating link. The proof of an important proposition will be delayed until Section 8. In Section 4, we recall some of Postnikov's results and prove Theorem 1.1. In Sections 5 and 6, respectively, we use the results of Section 3 to establish Theorems 1.3 and 1.4. In Section 7, we present some new evidence for the Conjecture made in [7] (i.e., for Conjecture 1.6 in the planar case)⁸.

Acknowledgements. We are grateful to Alexander Postnikov from whom we learned about the h -vector. The paper also benefited greatly from conversations with Dylan Thurston.

The first author is supported by a Japan Society for the Promotion of Science Grant-in-Aid for Young Scientists (B), no. 25800037, and the second author by a Grant-in-Aid for Scientific Research (C), no. 26400079.

2. Arborescences

Most arguments in the present paper are centered around a binary tree. We will give three mutually isomorphic descriptions of it. In this section, the binary tree will appear as the 'tree of arborescences.' In the next, we will describe it as the 'tree of subgraphs' and then as the 'HOMFLY computation tree.' But first, we introduce an object which will become a distinguished leaf of the tree of arborescences.

2.1. The clocked arborescence. All graphs that appear in this paper are finite. Multiple edges and loop edges are allowed. (Of course, the latter do not occur in bipartite graphs.) A *subgraph* of a graph will always have the same vertex set as the original, i.e., a subgraph will just be a subset of the edges of the graph. A *spanning tree* is a connected and cycle-free subgraph.

By definition, the edges of any plane graph G and the edges of its planar dual G^* are in a one-to-one correspondence. This gives rise to a bijection between subgraphs, where a set of edges of G is paired with the complementary set of the corresponding edges of G^* . Elements of such a pair will be called *dual subgraphs*.

⁸Note added in revision: Even though Conjecture 1.6 has been proven [9], we hope that the connections pointed out in Section 7 will be of interest to the reader.

If G and G^* are both connected, then it is well known that a subgraph (of G or of G^*) is a spanning tree if and only if its dual is one.

Definition 2.1. Let J be a directed graph (possibly with loop edges and multiple edges) and let us fix a vertex r_0 , called the *root*, in J . An *arborescence* rooted at r_0 is a subgraph of J so that

- its connected components not containing r_0 are isolated points and
- its connected component containing r_0 , called the *root component*, is a tree in which there is a (unique) directed path from r_0 to any other vertex.

A *spanning arborescence* is an arborescence without isolated points.

We remark that arborescences may never contain loop edges.

Let now G be a plane bipartite graph so that G^* is directed as explained in the introduction: if E and V are the color classes of G , then as we traverse each edge of G^* , we see an element of E to our right and an element of V to our left. We may write a directed edge as an ordered pair (startpoint, endpoint), although we should always keep it in mind that multiple edges may exist with the same initial and terminal points. Let us also recall that the vertex set of G^* is identified with the set R of regions of G (i.e., the set of connected components of $S^2 \setminus G$).

Lemma 2.2. *Let G be a plane bipartite graph. There exists a directed path from any vertex of G^* to any other vertex.*

Proof. Assume the contrary, i.e., that there exists a vertex $r_0 \in G^*$ so that the set $R' \subset R$ of vertices that are accessible from r_0 with directed paths is not R . Then the union of the corresponding (to elements of R') regions of G is not the entire sphere S^2 and hence it has non-empty boundary. (By region we mean the closure of a connected component of the complement of the embedding.) That boundary is a collection of cycles in G and since G is bipartite, each boundary component consists of at least two edges. It is easy to see that half of the edges along each boundary component are such that the corresponding edge of G^* points from an element of R' to an element of $R \setminus R'$. But that is a contradiction because by definition, the terminal points of these edges of G^* should be in R' . \square

Fix a root r_0 in G^* and an edge κ of G so that the dual edge κ^* points to r_0 . We will use a “greedy,” or depth-first, algorithm to construct a spanning arborescence of G^* determined by these data. In the process, we will select edges one-by-one so that at each stage we have an arborescence of G^* .

Let $\varepsilon_{0,1}$ be the first non-loop edge of G^* , directed away from r_0 , that we find as we turn around r_0 , starting from κ^* , in the positive (counterclockwise) direction. Let r_1 be the terminal point of $\varepsilon_{0,1}$. Now turn counterclockwise around r_1 , starting from $\varepsilon_{0,1}$, until the first edge $\varepsilon_{1,1}$ is found so that together with $\varepsilon_{0,1}$ they form an arborescence (i.e., $\varepsilon_{1,1}$ is neither a loop nor does it point to r_0). Now move to the terminal point r_2 of $\varepsilon_{1,1}$ and turn around it counterclockwise, starting from $\varepsilon_{1,1}$, until the next edge is found so that together with the first two, they form an arborescence, and so on.

If at any point in the process we select the edge $\varepsilon_{i,j} = (r_i, r_k)$ but complete a full turn around r_k without finding a suitable next edge, then we move back to r_i and continue turning counterclockwise around it from $\varepsilon_{i,j}$ until an edge $\varepsilon_{i,j+1}$ is found which forms an arborescence with the previously chosen ones. If this does not exist either, then we move back to the starting point r_l of the unique edge $\varepsilon_{l,m} = (r_l, r_i)$ in our arborescence which ends at r_i and continue searching for a suitable edge $\varepsilon_{l,m+1}$ by turning counterclockwise around r_l , starting from $\varepsilon_{l,m}$, and so on.

The edge-selecting algorithm terminates when a full turn has been completed around all vertices of G^* that we visited in the process (including r_0). We claim that the final arborescence A is spanning. Indeed, if there was an isolated point r in A then find a directed path in G^* from r_0 to r (cf. Lemma 2.2). Tracing this path backward from r , the first edge that is in A is preceded by an edge that could be added to A to form a larger arborescence (if A and the path are disjoint, then the same can be said about the first edge along the path). Moreover, when we were turning around the startpoint of this edge, we would have selected it into our arborescence, which is a contradiction.

Definition 2.3. Let G be a connected plane bipartite graph with dual graph G^* . The spanning arborescence of G^* constructed above will be called the *clocked arborescence* (relative to the vertex r_0 of G^* and the edge κ of G).

2.2. The tree of arborescences. We start with a technical digression. We stress that (say, smooth) embeddings of the graphs G and G^* into the sphere S^2 have been fixed. Let C denote the set of intersection points between edges of G and G^* . With a slight abuse of notation, we will also speak of C as the edge set of either G or G^* . Likewise, we use the symbol R both for the set of regions of G and for the set of vertices of G^* . The regions of G^* have their boundary oriented either clockwise or counterclockwise. These two sets of regions are identified with the color classes E and V , respectively, of G .

Now, let us fix smooth arcs connecting each vertex $r \in R$ to the vertices of G that lie along the boundary of the region r (but otherwise avoiding G and G^*). Together with G and G^* , these arcs form a triangulation T_R of S^2 . More precisely, the set of 0-cells is $E \cup V \cup R \cup C$ and the 1-cells are the arcs above along with the half-edges of G and G^* emanating from elements of C . Let us also fix a barycentric subdivision \mathcal{B} of T_R .

Definition 2.4. The *regular neighborhood* N_A of an arborescence A in G^* is the union of those (closed) 2-cells of \mathcal{B} that have a common point with the root component of A .

In combinatorial topology, a regular neighborhood is usually defined using a second barycentric subdivision. However for our purposes, Definition 2.4 will suffice and so we will use it to avoid unnecessary complication.

Next, we shall describe an algorithm that enumerates all arborescences of G^* rooted at r_0 and arranges them in a binary tree \mathcal{A} . The construction will depend

on the same edge κ of G as in the previous subsection. The nodes of the binary tree will actually be pairs of the form (arborescence, set of skipped edges) so that the skipped edges are edges of G^* , not in the arborescence, each of which has its startpoint in the root component. The nodes of \mathcal{A} have either no descendant or exactly two, which we will refer to as the right and left descendants. Here we imagine the descendants of each node to be ‘below’ it (with the root of the tree at the ‘top’) and slightly to the right or left. For practical reasons, in Figures 3 and 4 the root appears at the top left instead and the tree is distorted accordingly.

Our first arborescence, the root of the binary tree, is the one with no edges and no skipped edges. The right descendant of the root has the unique edge $\varepsilon_{0,1} = (r_0, r_1)$ that appeared in the construction of the clocked arborescence. It has no skipped edges. The left descendant of the root still has no edges, but it has the skipped edge $\varepsilon_{0,1}$. The following is the general description of our process.

Definition 2.5. Suppose that (A, S) has been chosen as a node of the tree \mathcal{A} , and let N_A denote the regular neighborhood of A . Let k be the intersection point of the boundary ∂N_A and the edge κ^* (which can never be in A , hence k exists). Now, let us move counterclockwise around ∂N_A starting from k until we reach the first edge δ of G^* that is not in S and which is such that $A \cup \{\delta\}$ is an arborescence. We will refer to δ as the *augmenting edge* of (A, S) . If such a δ does not exist, then the node (A, S) will have no descendants in \mathcal{A} . Otherwise, let the *right descendant* of (A, S) be $(A \cup \{\delta\}, S)$ and let its *left descendant* be $(A, S \cup \{\delta\})$.

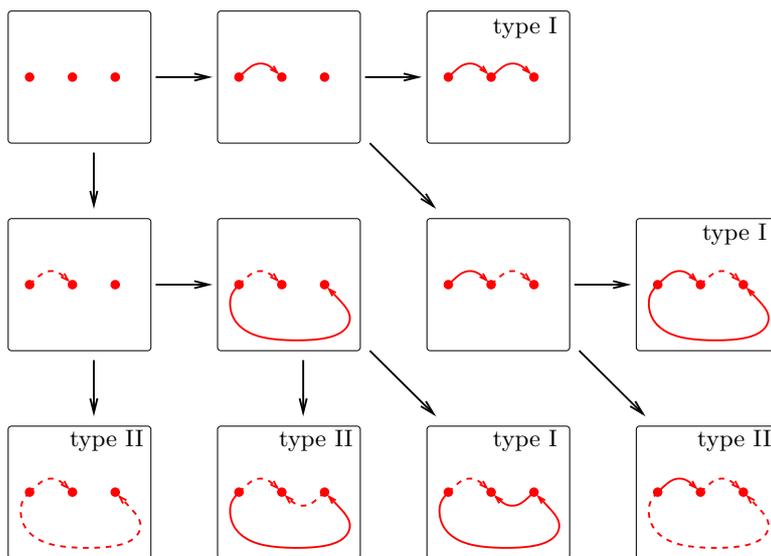


Figure 3. The tree of arborescences \mathcal{A} for the directed graph in Figure 2. For each node (A, S) , solid edges represent A and dotted edges represent S .

Example 2.6. Figure 3 shows the tree of arborescences for the complete bipartite

graph $G = K_{3,2}$. The embedding of G (and G^*), the root r_0 , and the edge κ are shown in Figure 2.

Lemma 2.7. *Along any path in \mathcal{A} that starts from the root, both the arborescences and the sets of skipped edges form an increasing sequence.*

Proof. Obvious from the construction. \square

It is easy to see that the rightmost branch of \mathcal{A} , i.e., the path when we always pass to the right descendant, leads to the clocked arborescence (with no skipped edges). As to other potential terminal nodes, we make the following observation.

Lemma 2.8. *The node (A, S) of \mathcal{A} has no descendants if and only if either*

- I. *A is a spanning arborescence or*
- II. *the set R' of vertices in the root component of A is a proper subset of R but all edges of G^* from an element of R' to an element of $R \setminus R'$ belong to S .*

With regard to the above, we will speak of *type I* and *type II leaves* of \mathcal{A} .

Proof. A spanning arborescence is also a spanning tree of G^* and hence it contains exactly $|R| - 1$ edges. Other arborescences have fewer edges. Consequently, spanning arborescences can not be extended as arborescences and so cannot have descendants in \mathcal{A} . If (A, S) fits the second description then it has no augmenting edge and hence no descendants, either.

As to the converse, let now (A, S) be so that A is not spanning and let α be an edge from the root component to an isolated point of A . Such an edge necessarily intersects the boundary ∂N_A of the regular neighborhood of A and hence it will be detected by the process described in Definition 2.5. Thus if (A, S) has no descendants in \mathcal{A} then any possible α belongs to S . \square

Lemma 2.9. *All spanning arborescences of G^* appear at a unique node (a type I leaf) of \mathcal{A} .*

Proof. Let A be a spanning arborescence of G . We will construct a path in \mathcal{A} that starts from the root and ends at (A, S) for an appropriate S . Assume that the path has already been constructed until the node (A', S') and let δ be the augmenting edge of (A', S') . If δ is an edge in A , then continue the path to the right descendant $(A' \cup \{\delta\}, S')$. Otherwise, move to the left descendant $(A', S' \cup \{\delta\})$.

We claim that the last node (\tilde{A}, \tilde{S}) along our path cannot be a leaf of type II. Since $\tilde{A} \subset A$ by construction, if $\tilde{A} \neq A$, then there has to be an edge of A going from a vertex of the root component of \tilde{A} to an isolated point of \tilde{A} . As our process never skips edges of A , this edge cannot be in \tilde{S} either, but that contradicts the definition of a type II leaf. Hence (\tilde{A}, \tilde{S}) is a type I leaf which implies $\tilde{A} = A$.

Regarding uniqueness, assume that there exists another path P in \mathcal{A} from the root to a leaf involving A . Let (A', S') be the node where the two paths (P and the one constructed above) separate. Since (A', S') cannot be a leaf, it has an

augmenting edge δ . Now the next node (A'', S'') along P after (A', S') is such that either

- $\delta \notin A$ but $\delta \in A''$ (if P takes a step to the right at (A', S')) or
- $\delta \in A$ but $\delta \in S''$ (if P takes a step to the left).

As both scenarios prevent the endpoint of P from involving A , we have a contradiction and the proof is complete. \square

If the arborescence A is not spanning then it may appear at multiple nodes of \mathcal{A} , as the example of the root and its left descendant already shows.

3. The computation tree

In this section we describe two more incarnations of the binary tree \mathcal{A} and we spell out their relation to the top of the HOMFLY polynomial P_{LG} .

It is a fairly straightforward matter to transform the tree \mathcal{A} of arborescences in G^* into the isomorphic tree \mathcal{G} of (decorated) subgraphs of G . Both trees depend on the same choice of edge κ . If (A, S) is a node in \mathcal{A} , then we replace it with the dual subgraph A^* of G . The edges of G that correspond to elements of S are in A^* and we will refer to them as the *dotted edges*. We may keep the same symbol S for the set of dotted edges. The tree structures of \mathcal{A} and \mathcal{G} are the same, that is, Definition 2.5 serves to describe \mathcal{G} as well. However it is useful to translate that description to subgraph terms.

Definition 3.1. Let G be a connected plane bipartite graph. The nodes of the *tree of subgraphs* \mathcal{G} are pairs (B, S) where B is a subgraph of G , that is a subset of the graph's edges, and $S \subset B$ is a further subset. The root of \mathcal{G} is (G, \emptyset) .

If (B, S) is already a node of \mathcal{G} then we walk along κ from its endpoint in V to the one in E and continue walking on edges of B so that on our left we always see the same region of B . We look for the first instance when an edge $\varepsilon \in B \setminus S$ is traversed from its endpoint in E to its endpoint in V and ε is such that $B \setminus \{\varepsilon\}$ is a connected subgraph of G (for example, it has no isolated points).

If such an ε is found then we let the *right descendant* of (B, S) be $(B \setminus \{\varepsilon\}, S)$ and let the *left descendant* be $(B, S \cup \{\varepsilon\})$. Else (B, S) has no descendants in \mathcal{G} .

We may refer to the region of G marked with r_0 as the ‘initial outside region.’ At each stage of the process described above, we look for a non-dotted edge ε of G to ‘puncture’ (i.e., remove) so that the outside region grows larger but its interior remains simply connected. In the right descendant of the subgraph, ε is removed. In the left descendant, the subgraph is the same but ε becomes dotted, so that it can not be punctured any more along the current branch of \mathcal{G} . A trivial induction proof shows that all nodes (B, S) of \mathcal{G} satisfy $\kappa \in B \setminus S$ and that κ , as well as all elements of S , are adjacent to the outside region.

A node of \mathcal{G} can become a leaf in two ways. Either the only region of the subgraph is the outside region, i.e., the subgraph is a tree – these are the type I

leaves. Or else, the (closed) outside region is not the entire sphere S^2 but along each boundary component, dotted and non-dotted edges alternate. In this case, that is in the case of a type II leaf, we say that the subgraph has an *alternating contour*. Note that no condition is imposed on edges that are incident on both sides with the outside region. Figure 4 shows three subgraphs (in the bottom row) with alternating contours.

Finally, in order to turn \mathcal{G} into the computation tree \mathcal{T} , we replace each subgraph with an oriented link diagram using the median construction. Here we use negative half-twists for non-dotted edges and positive ones for dotted edges. We orient the link by letting its arcs circle elements of E clockwise and elements of V counterclockwise. This way, crossings coming from non-dotted edges become positive while those coming from dotted edges are negative, cf. Figures 2 and 4. The link diagram at the root of \mathcal{T} is that of L_G . In terms of these diagrams, passing to a right descendant means *smoothing* a crossing and passing to a left descendant is equivalent to *changing* a crossing (from positive to negative).

Smoothing and changing of crossings play a crucial role in the definition of the HOMFLY polynomial $P(v, z)$, as we explained on page 5 of the Introduction. Let us quote the following well known result.

Theorem 3.2 (Morton [11]). *If an oriented link diagram contains n crossings and s Seifert circles, then in any term of the corresponding HOMFLY polynomial, the exponent of z is at most $n - s + 1$.*

The estimate is well known to be sharp for homogeneous link diagrams. This claim also follows from our main result on special alternating links (as stated in Theorems 1.3 and 1.4, although the statement in Theorem 3.5 suffices as well) and the following fact.

Theorem 3.3 (Murasugi–Przytyczki [12]). *Let D_1 and D_2 be oriented link diagrams so that they have n_1 and n_2 crossings, as well as s_1 and s_2 Seifert circles, respectively. Let us form a star product $D_1 \star D_2$ of D_1 and D_2 . Then $D_1 \star D_2$ has $n_1 + n_2$ crossings and $s_1 + s_2 - 1$ Seifert circles, and the coefficient of $z^{n_1+n_2-(s_1+s_2-1)+1}$ in $P_{D_1 \star D_2}(v, z)$ is the product of the coefficients of $z^{n_1-s_1+1}$ in $P_{D_1}(v, z)$ and that of $z^{n_2-s_2+1}$ in $P_{D_2}(v, z)$.*

Definition 3.4. For a link L that can be presented with a homogeneous diagram with n crossings and s Seifert circles, the *top of the HOMFLY polynomial* is the polynomial $T(v)$ in v that is the coefficient of z^{n-s+1} in $P_L(v, z)$.

Here $n - s$, and hence $T(v)$, does not depend on the homogeneous diagram used. When we combine the next theorem with Theorem 3.3, it follows that for any homogeneous link, all coefficients in the top of the HOMFLY polynomial have the same sign. In particular, they do not cancel when we pass to the Alexander polynomial $\Delta(t) = P(1, t^{1/2} - t^{-1/2})$, rather their sum becomes the leading coefficient.

Theorem 3.5. *Let G be a connected plane bipartite graph with s vertices and n edges, and let \mathcal{T} denote one of its computation trees that we constructed above.*

Then the top of the HOMFLY polynomial $P_{L_G}(v, z)$ is the sum to which each type I leaf of \mathcal{T} with k negative crossings (which appear as skipped edges in \mathcal{A} and as dotted edges in \mathcal{G}) contributes the monomial $v^{n-s+1+2k}$.

Proof. From the skein relation (2), it is obvious that if we label each edge of \mathcal{T} connecting a node to its right descendant with vz and each edge leading to a left descendant with v^2 , then \mathcal{T} can be used to compute the HOMFLY polynomial associated to its root L_G as the sum of the following terms: For each leaf, take the HOMFLY polynomial of the corresponding link and multiply it with the product of the edge labels along the unique path between the leaf and the root.

Type I leaves, where the corresponding subgraph is a spanning tree of G , are diagrams of the unknot. Therefore, since $P_{\bigcirc} = 1$, type I leaves contribute a single monomial, namely the product of the appropriate edge labels. Each spanning tree contains $s-1$ edges. So in order to reach a type I leaf, one needs to remove $n-s+1$ edges, that is, on the way from the root to the leaf, one has to take a step to the right exactly $n-s+1$ times. If the link diagram at the leaf contains k negative crossings, that means that we took a step to the left k times. Hence the contribution of the leaf to the HOMFLY polynomial is $(vz)^{n-s+1}(v^2)^k = v^{n-s+2k+1}z^{n-s+1}$.

The Theorem will obviously follow once we make sure that the type II leaves of \mathcal{T} do not contribute to the top of P_{L_G} . This is a consequence of the following.

Proposition 3.6. *Let the oriented link diagram D contain n crossings and s Seifert circles. Assume that there exists a region r in the complement $S^2 \setminus D$ with the following property: The arcs of D bounding r , of which there are at least two, are alternately oriented clockwise and counterclockwise so that at each crossing along ∂r , the counterclockwise arc passes under the clockwise one. We also assume that near each crossing along ∂r , only one quadrant formed by D belongs to r . (See Figure 16.) Then, in any term of the corresponding HOMFLY polynomial $P_D(v, z)$, the exponent of z is at most $n-s-1$.*

Indeed, along any path from the root of \mathcal{T} to a node in the tree, the exponent of z in the product of the corresponding edge labels is the number of steps taken to the right. That number agrees with the amount $n-n'$ by which the number of crossings decreased along the path. The number s of Seifert circles is constant throughout \mathcal{T} . Hence in order for the node to contribute terms containing z^{n-s+1} to P_{L_G} , the HOMFLY polynomial associated to the node has to contain terms with $z^{n'-s+1}$ in them. The estimate in Proposition 3.6 rules that out for type II leaves.

Note that the main assumption in the Proposition is satisfied by link diagrams that arise after (in the description of \mathcal{G} at the beginning of this section) an alternating contour in the subgraph has been achieved. There is however an extra assumption in the Proposition, namely that the outside region at this stage does not touch itself over an edge of G . General type II leaves of \mathcal{T} are obtained by connecting diagrams described in the Proposition in a tree-like fashion. Here by connecting, we mean joining the corresponding (embedded) Seifert graphs by paths

of edges (some of which may be dotted). In terms of link diagrams, that translates to joining by a sequence of (0 or more) bigons, which is just a complicated way of taking a connected sum. Because the HOMFLY polynomial is multiplicative under connected sums, the estimate discussed in the previous paragraph follows for general type II leaves from Proposition 3.6 by a short inductive argument (it even gets stronger as the number of components in the alternating contour increases). \square

Proposition 3.6 is just a slight improvement on Morton’s upper bound in Theorem 3.2, yet it turns out to be where most of the difficulty in this paper is concentrated. Notice that we did not make any assumption on the crossings not adjacent to r , even though when we apply Proposition 3.6 in the proof of Theorem 3.5, away from ∂r the diagram D is still alternating. Likewise, it is not assumed that D be special. We delay the proof of the Proposition until Section 8.

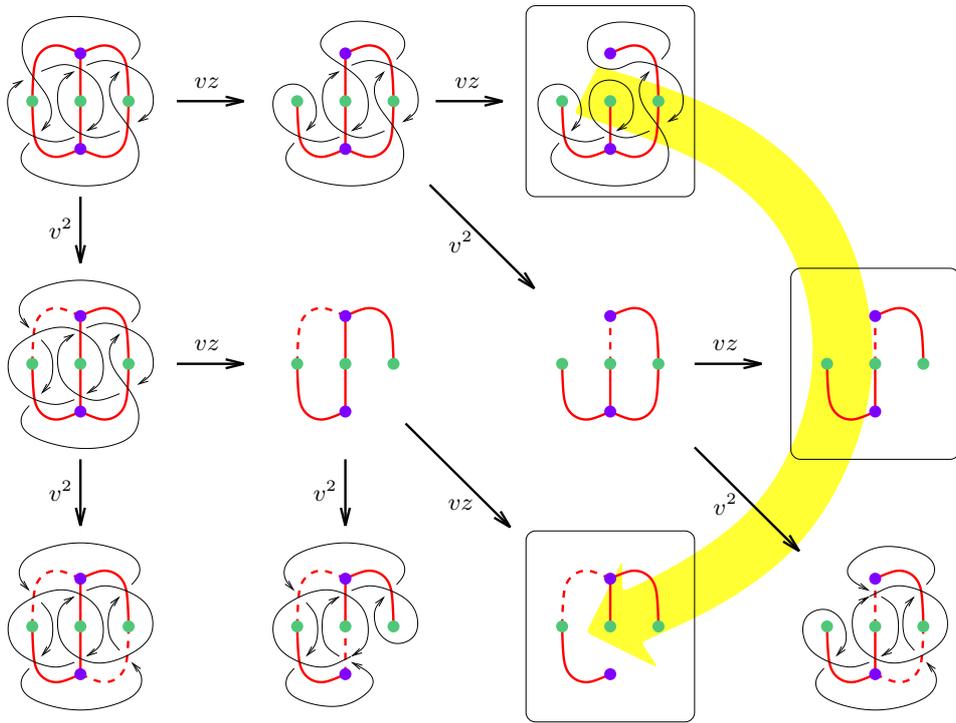


Figure 4. The tree of subgraphs \mathcal{G} , partially superimposed with the computation tree \mathcal{T} , in our running example. The semicircular arc indicates the right-to-left order of type I leaves.

Example 3.7. For the complete bipartite graph $G = K_{3,2}$ of Figure 2, the link L_G consists of three fibers of the Hopf fibration. Its diagram has $n = 6$ crossings and $s = 5$ Seifert circles, so the top of the HOMFLY polynomial is at z -exponent

$n - s + 1 = 2$. The three type II leaves of the computation tree \mathcal{T} (derived from Figure 3, shown in Figure 4) are two two-component unlinks with the HOMFLY polynomial $(v^{-1} - v)/z$ and a distant union of a Hopf link with an unknot, with $P = \begin{pmatrix} vz & \\ +vz^{-1} & -v^3z^{-1} \end{pmatrix} \cdot \frac{v^{-1} - v}{z}$. (These values are well known and easy to check. We write the HOMFLY polynomial in a slightly unconventional way to emphasize its bigraded nature.) Hence

$$\begin{aligned}
 P_{L_G}(v, z) &= \underbrace{2v^5z \cdot (v^{-1}z^{-1} - vz^{-1}) + v^4 \cdot \begin{pmatrix} 1 & -v^2 \\ +z^{-2} & -2v^2z^{-2} & +v^4z^{-2} \end{pmatrix}}_{\text{type II leaves}} \\
 &+ \underbrace{(v^2z^2 + 2v^4z^2) \cdot 1}_{\text{type I leaves}} = \begin{matrix} v^2z^2 & +2v^4z^2 & & \\ +3v^4 & -3v^6 & & \\ +v^4z^{-2} & -2v^6z^{-2} & +v^8z^{-2} & \end{matrix} \quad (5)
 \end{aligned}$$

In particular, the top of P_{L_G} is $T(v) = v^2 + 2v^4$. Notice how, in accordance with the proof of Theorem 3.5, contributions to the top came from type I leaves only. The Conway polynomial is $\nabla_{L_G}(z) = P_{L_G}(1, z) = 3z^2$ and the Alexander polynomial is $\Delta_{L_G}(t) = \nabla_{L_G}(t^{1/2} - t^{-1/2}) = 3t - 6 + 3t^{-1}$.

Remark 3.8. If one color class of G consists of valence 2 points, that is if G is obtained from a (plane) graph by placing an extra vertex at the midpoint of each edge, then (as in the previous example) the Conway polynomial $\nabla_{L_G}(z)$ is a single monomial. This follows, for example, from a result of Jaeger [5].

4. Triangulations of the root polytope

In order to relate our results on arborescences to a different kind of combinatorics, we need to review some of Postnikov's results [14]. This section also contains the proof of Theorem 1.1.

Let G be an abstract bipartite graph. That is, we do not assume an embedding of G into the plane. Let us denote the color classes of G with E and V . The graph G may have multiple edges but they do not affect the following construction.

For $e \in E$ and $v \in V$, let \mathbf{e} and \mathbf{v} , respectively, denote standard generators of $\mathbb{R}^E \oplus \mathbb{R}^V$. Let the *root polytope* of G be

$$Q_G = \text{Conv}\{\mathbf{e} + \mathbf{v} \mid ev \text{ is an edge in } G\},$$

where Conv denotes the usual convex hull.

Lemma 4.1 ([14, Lemma 12.5]). *Let G be a connected bipartite graph on s vertices. The dimension of Q_G is $s - 2$. A set of vertices of Q_G is affinely independent if and only if the corresponding edges in G form a cycle-free subgraph. In particular,*

maximal (i.e., $(s - 2)$ -dimensional) simplices in Q_G correspond to spanning trees of G . Furthermore, the volumes of such maximal simplices agree.

Definition 4.2. A collection of maximal simplices in Q_G (so that their vertices are also vertices of the root polytope) is a *triangulation* if its union is Q_G and if every two of its members intersect in a common face.

Studying these triangulations reveals many a subtle phenomenon, see [14]. First let us quote the translation, to subgraph terms, of the second condition of the definition.

Lemma 4.3 ([14, Lemma 12.6]). *Let Γ_1 and Γ_2 be spanning trees in G . The following two statements are equivalent.*

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

The proofs of the previous two lemmas are short and elementary. From the last assertion in Lemma 4.1, it follows that each triangulation of Q_G consists of the same number of simplices. Postnikov also expresses that value in terms of G .

Theorem 4.4 (Postnikov [14]). *Let G be a connected bipartite graph with color classes E and V . The number of simplices in each triangulation of the root polytope Q_G is the number of possible valence distributions, taken at elements of E , of spanning trees of G .*

There is an obvious sense in which (V, E) is a hypergraph [8] and in that context it is fairly natural to rename (essentially) the valence distributions above as follows.

Definition 4.5. Let G be a connected bipartite graph with color classes E and V . A function $\mathbf{f}: E \rightarrow \mathbb{N}$ is called a *hypertree* (in the hypergraph (V, E)) if G has a spanning tree with valence $\mathbf{f}(e) + 1$ at each $e \in E$.

With this, Theorem 4.4 says that the number of simplices needed to triangulate Q_G is the number of hypertrees in (V, E) . Of course, the same can be claimed regarding the ‘abstract dual’ (a.k.a. transpose) hypergraph (E, V) as well.

Proof of Theorem 1.1. Even though all the necessary ingredients are included in a previous paper [8], we spell out a proof for completeness and in order to elaborate on some details. As usual, let us denote the color classes of G by E and V .

Let us consider a pair A_1, A_2 of arborescences in G^* (both rooted at r_0). Assume that their dual trees Γ_1, Γ_2 violate the condition in Lemma 4.3, that is, there exists a cycle Φ in G of length at least 4, composed of edges alternately from Γ_1 and Γ_2 . The root r_0 does not lie along Φ because neither A_1 nor A_2 has edges ending at r_0 . Let Φ bound the disks U and U' in S^2 . Since elements of E and V alternate along Φ , it easily follows that all edges of A_1 that cross Φ do so from one side to the other, say from U to U' . The same is true for edges of A_2 crossing Φ , but those travel from U' to U . But then, the fact that r_0 cannot be in U and in U' simultaneously prevents one of A_1 and A_2 from being an arborescence: If,

say, $r_0 \in U$, then A_2 is not an arborescence because the startpoints of its edges crossing Φ (which exist because A_2 is connected) cannot be reached by an oriented path from r_0 .

Now that we have seen that all pairs of simplices resulting from our arborescences satisfy the compatibility condition in Definition 4.2, we just have to make sure that there are enough of them so that their union is the entire root polytope. By Theorem 4.4, it suffices to show that any hypertree in (V, E) is realized by a spanning tree of G that is dual to a spanning arborescence of G^* . But that is part of the statement of Theorem 10.1 in [8]. \square

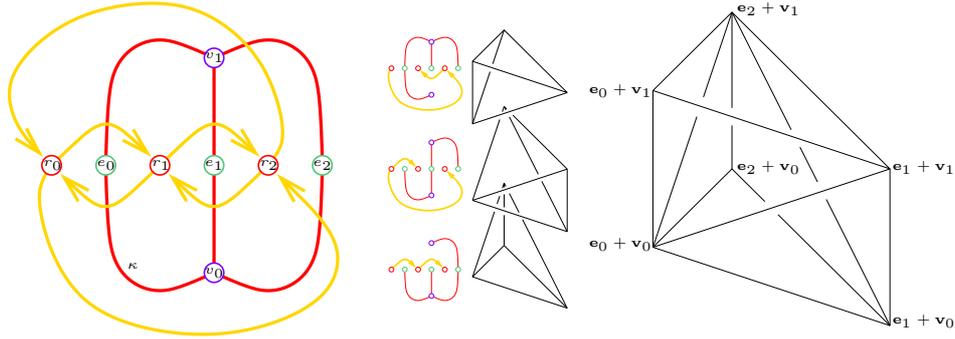


Figure 5. Left: The graph $G = K_{3,2}$ with color classes $E = \{e_0, e_1, e_2\}$ and $V = \{v_0, v_1\}$, and its dual G^* . Middle: The spanning arborescences in G^* (relative to the root r_0), their dual spanning trees in G , and the corresponding maximal simplices in Q_G . Right: The three simplices triangulate Q_G .

Example 4.6. The root polytope of the complete bipartite graph $G = K_{3,2}$ is the product of an interval and a triangle. (In general, Q_G is obtained from the product $\Delta_E \times \Delta_V$ of the $(|E|-1)$ -dimensional unit simplex Δ_E and the $(|V|-1)$ -dimensional unit simplex Δ_V by truncating vertices corresponding to non-edges of G .) In Figure 5, we show the triangulation of Q_G corresponding to the arborescences found in Example 2.6.

5. The HOMFLY polynomial and the root polytope

In this section we prove Theorem 1.3. As explained in the Introduction, the type I leaves of the binary tree \mathcal{A} have a natural order from right to left (see Figure 4 for an example). The first (smallest) element in the order is the clocked arborescence of subsection 2.1. The definition of shelling order was given on page 4 of the Introduction.

Theorem 5.1. *Let G be a connected plane bipartite graph. Let us choose a root r_0 and adjacent edge κ to triangulate the root polytope Q_G as in Theorem 1.1. Then,*

the order on the set of maximal simplices induced by the right-to-left order of the corresponding spanning arborescences is a shelling order. In particular, the triangulations of Q_G described in Theorem 1.1 are shellable.

Proof. As always, let us fix the root r_0 and the edge κ that are used to construct the tree \mathcal{A} . Recall that the vertex set of Q_G is identified with the set C , which in turn can be viewed both as the edge set of G and as the edge set of G^* .

Let A and B be spanning arborescences so that $A < B$, i.e., the unique leaf of \mathcal{A} that involves A (cf. Lemma 2.9) is to the right of the leaf involving B . We have to find a third spanning arborescence A' (allowing for $A' = A$) so that the corresponding simplices $\sigma_A, \sigma_{A'}, \sigma_B \subset Q_G$ (with vertex sets identified with $C \setminus A$, $C \setminus A'$, and $C \setminus B$, respectively) satisfy

- (i) $\sigma_B \cap \sigma_{A'}$ is a codimension one face, i.e., A' and B differ in exactly one edge
- (ii) $\sigma_B \cap \sigma_A \subset \sigma_{A'}$, that is, $A' \subset A \cup B$
- (iii) $\sigma_{A'}$ precedes σ_B , i.e., $A' < B$ in the right-to-left order.

Let us find the node (A_0, S_0) of \mathcal{A} that is the last common node along the paths connecting the root to A and B , respectively. Let δ be the augmenting edge of (A_0, S_0) . Then by Lemma 2.7, δ is an element of A and it is a skipped edge for B . We construct A' by adding δ to B and removing the edge of B with the same terminal point. By [8, Lemma 9.8], this procedure is well defined and results in a spanning arborescence A' . (Furthermore, A' is the unique spanning arborescence that contains δ and all but one edge of B .) Let us check that A' satisfies our requirements.

The condition (i) is obviously true by construction. The only edge of A' that is not an edge of B is δ ; since that is an edge of A , (ii) holds as well. Finally, we claim that the leaf of \mathcal{A} involving A' (cf. Lemma 2.9) is either a descendant of the right descendant $N = (A_0 \cup \{\delta\}, S_0)$ of (A_0, S_0) , or else, the path from the root to A' separates from the path to N so that at some node, the former goes to the right and the latter to the left. This implies (iii) immediately.

Indeed, if the path from the root to A' separated from the path to N when taking a step to the left, then an edge of $A_0 \cup \{\delta\}$ (namely the augmenting edge at the parting of the paths) would be a skipped edge for A' (cf. Lemma 2.7). But since $A_0 \cup \{\delta\} \subset A'$ (note that $A_0 \subset B$ and the unique edge in $B \setminus A'$ cannot be in A_0 because it shares terminal points with the augmenting edge δ of A_0), this is impossible and the proof is complete. \square

Theorem 5.2. *When we compute the h -vector of the triangulation in Theorem 1.1 using the shelling of Theorem 5.1, the contribution c_i (cf. (1)) of each simplex σ_i is equal to the number of skipped edges for the corresponding type I leaf of \mathcal{A} .*

Proof. In the proof of Theorem 5.1, a spanning arborescence A' was constructed for any spanning arborescence B and skipped edge $\delta \in S_B$. (Here S_B is the set of skipped edges for the unique leaf of \mathcal{A} involving B , cf. Lemma 2.9. Indeed, A' depended on A only through the choice of δ .) Since $A' \cap S_B = \{\delta\}$ by construction,

it is clear that different choices of δ yield different arborescences A' . By (i) and (iii) of the proof, we then see that for any spanning arborescence B , the number c_B of maximal simplices that precede the corresponding simplex σ_B in the shelling order and share a common facet with it, is at least $|S_B|$.

To prove the converse inequality, fix B and let A be a spanning arborescence that precedes B in the right-to-left order so that A only differs from B in one edge. As in the proof of Theorem 5.1, let (A_0, S_0) be the last common node of \mathcal{A} along the paths from the root of \mathcal{A} to A and B , respectively. Then, since $A < B$, it is clear that the path toward A passes through the right descendant of (A_0, S_0) and the path toward B passes through the left descendant. Therefore, if δ is the augmenting edge of (A_0, S_0) , then $\delta \in A$ but δ is a skipped edge for B . As A and B do not otherwise differ, it is clear that A coincides with the arborescence A' (for the δ we have just chosen) of the previous paragraph. \square

Proof of Theorem 1.3. Using the description (1) of the h -vector and the fact that the dimension of Q_G is $d = |E| + |V| - 2 = s - 2$, we have

$$\begin{aligned} \text{top of } P_{LG} &= v^{n-s+1} \sum_{\text{type I leaves}} v^{2k} = v^{n-s+1} \sum_{\text{maximal simplices}} v^{2c_i} \\ &= v^{n+s-1} \sum_{\text{maximal simplices}} v^{-2(s-1)+2c_i} = v^{n+s-1} h(v^{-2}), \end{aligned}$$

where the first equation is the statement of Theorem 3.5 and the second follows from Theorem 5.2. \square

Example 5.3. In Figure 4, we indicated the right-to-left order of the type I leaves of the computation tree for the graph $K_{3,2}$. In the middle panel of Figure 5, we find the three corresponding simplices arranged from bottom (smallest) to top. That is a shelling order for the triangulation so that $c_1 = 0$ and $c_2 = c_3 = 1$. Hence the h -vector is $h(x) = x^4 + 2x^3$. This, when compared to (5), confirms Theorem 1.3 in this case.

6. The HOMFLY polynomial and parking functions

The goal of this section is to prove Theorem 1.4. The following definition is due to Postnikov and Shapiro [15]. The only modification we made was to replace (relative) out-degree with in-degree, which is equivalent to an overall reversal of orientation in the directed graph.

Definition 6.1. Let $J = (R, C)$ be a directed graph with root $r_0 \in R$. For a non-empty subset $R' \subset R \setminus \{r_0\}$ and $r \in R'$, define the *relative in-degree* $\deg_{R'}(r)$ of r as the number of edges in C with endpoint r and startpoint outside of R' .

A function $\pi: R \setminus \{r_0\} \rightarrow \mathbb{N}$ is called a *parking function* of J with respect to r_0 if any non-empty subset $R' \subset R \setminus \{r_0\}$ contains a vertex r so that $\pi(r) < \deg_{R'}(r)$.

Let us denote the set of parking functions with Π and introduce the polynomial

$$p(u) = \sum_{\pi \in \Pi} u^{\left(\sum_{r \in R \setminus \{r_0\}} \pi(r) \right)},$$

which we call the *parking function enumerator*.

Example 6.2. We determine parking functions for the directed graph G^* of our running example (see the left panel of Figure 5 for notation). If r_0 plays the role of root, then the values $\pi(r_1), \pi(r_2)$ are subject to three conditions. The single element of $\{r_i\}$ has relative in-degree 2 ($i = 1, 2$), which implies $0 \leq \pi(r_i) \leq 1$. Both elements of $\{r_1, r_2\}$ have relative in-degree 1 and hence for any parking function π , one of $\pi(r_1), \pi(r_2)$ has to be 0.

There are three solutions for $(\pi(r_1), \pi(r_2))$: $(0, 0)$, $(1, 0)$, and $(0, 1)$. Therefore the parking function enumerator for G^* is $p(u) = 1 + 2u$ and, by comparing this to (5), we see that Theorem 1.4 holds in this case.

For any directed graph J and root r_0 , the set of parking functions is ‘closed downward’: if π is a parking function and $0 \leq \rho(r) \leq \pi(r)$ for all $r \in R \setminus \{r_0\}$, then ρ is a parking function too. The identically 0 function is a parking function (i.e., Π is non-empty) if and only if any subset of $R \setminus \{r_0\}$ has edges of J reaching it from the outside. Our directed graphs $J = G^*$ have this property by Lemma 2.2. Let us now investigate whether the next simplest candidates are parking functions.

Lemma 6.3. *If G is a connected plane bipartite graph and r_0 is an arbitrary vertex of G^* , then the following two statements are equivalent.*

- (a) *For all non-root vertices $r \in R \setminus \{r_0\}$ of G^* , the indicator function $\mathbf{i}_{\{r\}}$ is a parking function with respect to r_0 .*
- (b) *The graph G does not have multiple edges.*

Proof. (a) \implies (b). Let us assume that G has a pair of parallel edges ε_1 and ε_2 . They bound two bigons on S^2 , one of which does not contain r_0 . Let us denote the set of regions of G (i.e., vertices of G^*) in this bigon with R' . (Note that R' does not have to consist of a single point.) The set R' has a single incoming edge and we may assume it is the dual ε_1^* of ε_1 with the terminal point $r \in R'$. (Then ε_2^* is the single outgoing edge of R' .) Now it is easy to see that the function $\mathbf{i}_{\{r\}}$ and the set R' fail to satisfy the condition in Definition 6.1, hence the former is not a parking function.

(b) \implies (a). If $\mathbf{i}_{\{r\}}$ is not a parking function, then there exists a set $R' \subset R \setminus \{r_0\}$ with either no incoming edges (which would contradict Lemma 2.2) or with a single incoming edge ε^* whose terminal point is r . Since each vertex of G^* has the same in-degree as out-degree (in fact, for each region q of G , elements of E and V alternate along ∂q and hence incoming and outgoing edges alternate around the corresponding vertex of G^*), it follows that R' also has a single outgoing edge δ^* .

Let us now turn clockwise around r , starting from ε^* , until we find the next edge of G^* adjacent to r . It is necessarily such that its initial point is r . At

the terminal point of this edge, let us again turn clockwise until we find the next edge and so on. This way, a directed path is constructed which we follow until its next edge, which is necessarily δ^* at that stage, leaves R' . On the left side of each edge of the path, as well as on the left side of ε^* and δ^* , we see the same region of G^* , marked by a vertex v from the color class V of G . (This observation also shows why our path has to be finite: because the initial point of ε^* is not in R' , it is not possible to keep turning around v without leaving R' .) By turning counterclockwise instead of clockwise, we may construct an analogous path from r to the initial point of δ^* and this shows that on their right side, ε^* and δ^* are adjacent to the same vertex e from the color class E of G . This means that the dual edges ε and δ of G both connect e and v , which concludes our argument. \square

It is well known that parking functions of J with respect to r_0 are in a one-to-one correspondence with spanning arborescences of J rooted at r_0 [15]. Several bijections have been given between the two sets (see [1], [4], and references therein). In the next theorem we describe another identification that is special to the case when $J = G^*$ for a connected plane bipartite graph G . Recall that the tree \mathcal{A} constructed in Section 2 is such that for any choice of auxiliary data r_0, κ , and for any spanning arborescence A of G^* rooted at r_0 , there is a unique set S of skipped edges so that (A, S) becomes a type I leaf of \mathcal{A} .

Theorem 6.4. *For a connected plane bipartite graph G , root $r_0 \in G^*$, and adjacent edge $\kappa \in G$, let (A, S) be a type I leaf of the tree of arborescences \mathcal{A} . For every vertex $r \neq r_0$ of G^* , let $\pi(r)$ be the number of edges in S that point to r . Then π is a parking function with respect to r_0 .*

Furthermore, every parking function (for the graph G^ with root r_0) arises from a unique spanning arborescence in the manner described above.*

Proof. Let π be derived from A as in the Theorem. Fix a non-empty set $R' \subset R \setminus \{r_0\}$. Regarding the unique path in \mathcal{A} from the root (\emptyset, \emptyset) to (A, S) , let (A', S') be the first node along the path so that the root component of A' has a vertex, say r , from R' . We have $A' \subset A$ and $S' \subset S$ by Lemma 2.7. All edges of S that end at r actually belong to S' because once r is in the root component, no edges that end there can enlarge the arborescence and therefore they will never serve as augmenting edges. It is also clear that all edges in S' ending at r , as well as the edge of A' that ends at r , have their startpoint in $R \setminus R'$. Hence $\pi(r) < \deg_{R'}(r)$, which proves our first claim.

Let us now fix an arbitrary parking function $\pi: R \setminus \{r_0\} \rightarrow \mathbb{N}$. We will construct the corresponding spanning arborescence by finding the path in \mathcal{A} that leads to it from the root. The root (\emptyset, \emptyset) is obviously such that the number of skipped edges ending at any vertex $r \in R \setminus \{r_0\}$ is at most $\pi(r)$. Suppose that we have already built a path in \mathcal{A} ending at the node (A, S) that also has the property that

$$\text{each vertex } r \in R \setminus \{r_0\} \text{ has at most } \pi(r) \text{ elements of } S \text{ pointing to it.} \quad (6)$$

Let the augmenting edge of (A, S) be $\delta = (q, r)$. If $\pi(r)$ is strictly more than the number of edges in S pointing to r , then we pass to the left descendant $(A, S \cup \{\delta\})$ of (A, S) . If $\pi(r)$ equals the number of elements of S ending at r , then we go to the right descendant $(A \cup \{\delta\}, S)$. In either case, the new node of \mathcal{A} has property (6), so we may continue along our path until we arrive at a leaf (A_π, S_π) in \mathcal{A} .

We claim that (A_π, S_π) cannot be a type II leaf. Indeed, assume for contradiction that the set R'' of isolated points of A_π is non-empty. Since π is a parking function, there exists a vertex $r \in R''$ so that the number of edges starting in the root component of A_π and ending at r is more than $\pi(r)$. But because by Lemma 2.8 all of those edges belong to S_π , this contradicts (6).

By the previous paragraph, A_π is a spanning arborescence. It is clear from the construction that its associated parking function is π : for each $r \in R \setminus \{r_0\}$, we did select an edge pointing to r into A_π and we did so when exactly $\pi(r)$ skipped edges lead to r . After that, no more edges with terminal point r get skipped.

Finally, we argue that spanning arborescences $A' \neq A_\pi$ may not induce π . Let the paths in \mathcal{A} that lead from the root to (A_π, S_π) and to (A', S') , respectively (for the appropriate set S'), part ways at the node (A, S) . Let (A, S) have the augmenting edge $\delta = (q, r)$. If our earlier choice (leading to (A_π, S_π)) was the left descendant of (A, S) then the right descendant, along with all its subsequent descendants including (A', S') , is such that the number of skipped edges pointing to r is less than $\pi(r)$ (they “reach r too soon”). On the other hand, if earlier we chose the right descendant, then along the branch of \mathcal{A} corresponding to the left descendant, we have more than $\pi(r)$ skipped edges ending at r (i.e., those arborescences “reach r too late”). This completes the proof. \square

Proof of Theorem 1.4. We will describe parking functions in terms of spanning arborescences as in Theorem 6.4. Since every skipped edge for the spanning arborescence A has its unique terminal point in $R \setminus \{r_0\}$, the number of skipped edges for A is the same as the sum of the values (i.e., the index) of the corresponding parking function. Thus by Theorem 3.5 and the definition of the parking function enumerator (cf. (3), (4)), we have

$$\text{top of } P_{L_G} = v^{n-s+1} \sum_{\text{type I leaves}} v^{2k} = v^{n-s+1} \sum_{\text{parking functions}} v^{2i(\pi)} = v^{n-s+1} p(v^2),$$

which completes the proof. \square

7. HOMFLY and interior polynomials

It has long been conjectured⁹ by the first author that for a connected plane bipartite graph G of color classes E and V , the top of the HOMFLY polynomial P_{L_G} is equivalent to the interior polynomial I of the hypergraph $\mathcal{H} = (V, E)$. This is

⁹Note added in revision: a proof is now available as a preprint [9].

what motivated the development of the interior polynomial [8] in the first place. If the conjecture was proved then, among other things, we would obtain a way of deriving certain HOMFLY coefficients from Floer homology. (See [7] and the Introduction to this paper for an explanation.) We expect the conjecture to be resolved via the h -vector, as explained in the Introduction. However in this section, we point out several ways in which our results on parking functions also support the conjecture.

First, we note that in the case when the directed graph J is derived from a connected undirected graph K by replacing each edge with a pair of oppositely oriented edges between the same two vertices, the parking function enumerator p_J is related to the Tutte polynomial T_K via the formula

$$p_J(u) = u^{b_1(K)} T_K(1, 1/u),$$

where $b_1(K)$ is the first Betti number of K (viewed as a one-dimensional complex). To see this, one has to note the (quite direct) connection [15] between parking functions and critical configurations of the abelian sandpile model (a.k.a. chip-firing game), and then apply a formula of Merino López [10].

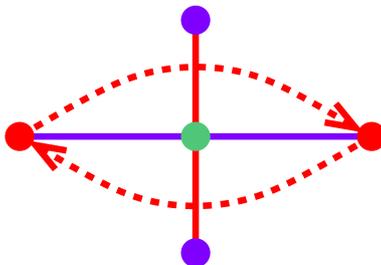


Figure 6. The two edges of G adjacent to the valence 2 point $e \in E$ are shown vertically in red. Their union can be viewed as an edge of the graph K^* . The two horizontal segments together form an edge of K . The dotted edges parallel to it belong to G^* .

Let us now consider a plane bipartite graph G so that one of its color classes, say E , contains only degree 2 vertices. In other words, G is obtained from some plane graph K^* by subdividing its edges. (We may formally write $K^* = (V, E)$.) Now the planar dual K of K^* induces a directed graph J as above and it turns out that $J = G^*$; see Figure 6. Finally, the duality formula of the Tutte polynomial implies

$$p_{G^*}(u) = u^{b_1(K)} T_K(1, 1/u) = u^{|V|-1} T_{K^*}(1/u, 1) = I_{(V,E)}(u),$$

cf. [8, equation (9)], where I is the interior polynomial of the (hyper)graph $K^* = (V, E)$. Thus, Theorem 1.4 reconfirms the equivalence of the top of the HOMFLY polynomial and the interior polynomial in this case. Of course, for this kind of graph G the equivalence is also a consequence of Jaeger’s formula [5].

Returning to the case of general connected plane bipartite graphs, if n is the number of edges in G and $s = |E| + |V|$ is the number of its vertices, then it is not hard to show that the coefficient of $(vz)^{n-s+1}$ in P_{L_G} and the constant term in I agree, namely they are both equal to 1. The argument below will re-establish this fact, and it will provide information on the coefficient of $v^{n-s+3}z^{n-s+1}$ as well.

First we show that the presence of multiple edges in G does not affect either the interior polynomial or the top of P_{L_G} . The former is obvious because edge multiplicities do not influence the structure of \mathcal{H} at all.

Lemma 7.1. *Increasing the multiplicity of an existing edge in the connected plane bipartite graph G changes the top $T(v)$ of the HOMFLY polynomial of L_G by a multiplicative factor of v for each new edge. In particular, the sequence of HOMFLY coefficients along the top remains the same.*

Proof. Assume that G contains n edges and s vertices. Let ε be an edge of G and let us construct G' by adding a parallel copy ε' of ε to G , so that G' has $n + 1$ edges but still only s vertices. (Note that ε and ε' do not necessarily bound an empty bigon of S^2 .) We apply the skein relation (2) to $L_{G'}$ and the crossing c that corresponds to ε' . Smoothing c results in L_G and the coefficient vz ensures that v times the top of P_{L_G} is part of the top of $P_{L_{G'}}$. Changing c gives a diagram that admits a link isotopy called a flype. If we carry out the flype, the number of crossings decreases to $n - 1$ but the number of Seifert circles remains the same. Since the corresponding coefficient is v^2 , Theorem 3.2 implies that we get no new contributions to the top of $P_{L_{G'}}$. \square

The bipartite graph $\text{Bip } \mathcal{H}$ of [8], associated to \mathcal{H} , is G with edge multiplicities reduced to 1. Hence if G has no multiple edges, which we may now assume without loss of generality, then the coefficient of the linear term in $I_{\mathcal{H}}$ is the first Betti number of G [8, Theorem 6.3]. The constant term of $I_{\mathcal{H}}$ is easily seen to be 1. (Unfortunately, beyond the first two, individual coefficients of I do not have similar easy descriptions¹⁰.) From Lemma 6.3 and Theorem 1.4, we see that the second coefficient along the top of P_{L_G} (that is, the coefficient of $v^{n-s+3}z^{n-s+1}$) is again the first Betti number of G . From either Theorem 1.3 or 1.4, it is easy to see that the first coefficient is 1. Hence we may conclude that the first two coefficients in the top of the HOMFLY polynomial P_{L_G} and in the interior polynomial $I_{\mathcal{H}}$ agree.

8. An improvement on Morton's inequality

This section adds the last remaining piece to establish our results. Namely, our present goal is to prove that type II leaves in the computation tree of Section 3 do not affect the top of the HOMFLY polynomial. We have seen how this boils

¹⁰Note added in revision: In [9, Proposition 5.5], a formula is found for the third (quadratic) coefficient. It is easy to equate that number to the number of parking functions with index 2.

down to Proposition 3.6, an important technical result that slightly strengthens Morton’s inequality for a specific kind of link diagram.

First, we establish several lemmas on oriented curves immersed in the plane. We assume these immersions to be generic, i.e., to have no self-tangencies or triple points. Since that is exactly the class of curves that we get if we forget the crossing information in an oriented link diagram, we will refer to our immersed curves as *link projections*. It is useful for us to study them because the numbers of crossings and of Seifert circles in a link diagram (and thus the upper bound that we are seeking in Proposition 3.6) only depend on the corresponding link projection. We may apply Reidemeister moves in this context as well.

Definition 8.1. We introduce the following terms for certain isotopies of link projections.¹¹

- A *Reidemeister I-a move* removes a kink of a link projection (see Figure 7). Here, either orientation is allowed.

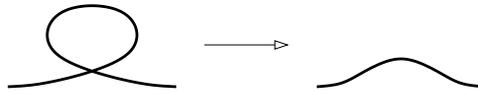


Figure 7. Reidemeister I-a move

- We call the move shown in Figure 8 (Figure 9, respectively) a *cyclic Reidemeister II-a move* (*noncyclic Reidemeister II-a move*, respectively).

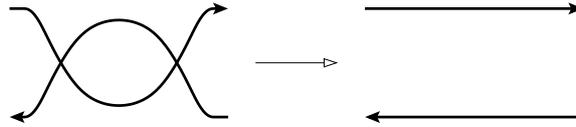


Figure 8. Cyclic Reidemeister II-a move

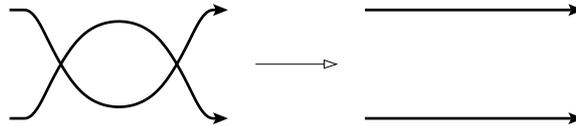


Figure 9. Noncyclic Reidemeister II-a move

- The isotopy shown in Figure 10 is called a *global noncyclic Reidemeister II-a move*. Here, the shaded region may contain arcs from the link projection. A noncyclic Reidemeister II-a move is a special case of a global noncyclic Reidemeister II-a move.

¹¹A note on terminology: ‘b’ moves would be the inverses of the ‘a’ moves below, but those will not play a role in our treatment.

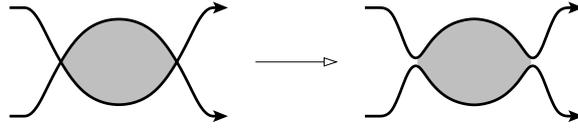


Figure 10. Global noncyclic Reidemeister II-a move

- We call the move shown in Figure 11 a *noncyclic Reidemeister III move*.

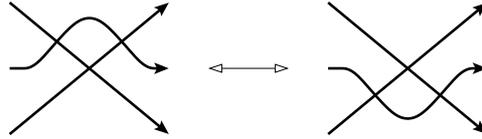


Figure 11. Noncyclic Reidemeister III move

For a link projection D , let $n(D)$ and $s(D)$ denote the number of crossings and Seifert circles, respectively. Let us analyze how the value $n(D) - s(D)$, which is essentially our desired upper bound, changes under the moves of Definition 8.1.

Definition 8.2. If an isotopy (either of a link projection or of a link diagram) reduces $n - s$, then we call it a *good move*. If the isotopy preserves $n - s$, then we call it a *fair move*.

Lemma 8.3. A Reidemeister I-a move is a fair move, a cyclic Reidemeister II-a move is a good or fair move, a (local or global) noncyclic Reidemeister II-a move is a good move, and a noncyclic Reidemeister III move is a fair move.

Proof. From the left panel in Figure 12 (where the dashed curves indicate Seifert circles) one sees immediately that a Reidemeister I-a move decreases both the number of Seifert circles and the number of crossings by one. So a Reidemeister I-a move is a fair move.



Figure 12. Seifert circles before a Reidemeister I-a and a cyclic II-a move.

Let us consider the case of a cyclic Reidemeister II-a move. Before the move, either three or two Seifert circles are adjacent to the two crossings involved in the move. (In Figure 12, right panel, the Seifert circle leaving at the top-right point may or may not re-emerge at the top-left point.) In the first case, the move

decreases the number of Seifert circles by two, and in the second case, the number of Seifert circles remains the same. Since the number of crossings is reduced by two, a cyclic Reidemeister II-a move is either a good or a fair move.

Under a (global) noncyclic Reidemeister II-a move, the configuration of Seifert circles does not change. As the number of crossings gets reduced by two, a (global) noncyclic Reidemeister II-a move is a good move.

Since neither the number of Seifert circles nor the number of crossings changes under a noncyclic Reidemeister III move, it is a fair move. \square

If one of the good moves above is applied to a link diagram, then the value of $n - s$ actually drops by 2. Hence if a link diagram admits a good Reidemeister move (or in fact, any good move), then (by the invariance of the HOMFLY polynomial) Theorem 3.2 implies Proposition 3.6 for that case. It is also easy to see the following.

Lemma 8.4. *If we are able to apply a fair or good move to a link diagram and the estimate of Proposition 3.6 holds after it, then it also holds before the move.*

However for now, we are still working in the category of link projections. As our next intermediate step, we will establish a way of handling certain self-intersections. By a *domain* of a link projection, let us mean a connected set that is a union of several regions. Note that domains are closed. By *emptying a domain* of a link projection (or diagram), we mean an isotopy that leaves the boundary arcs of the domain fixed so that at the end, the interior of the domain is disjoint from the projection/diagram. (I.e., after the isotopy the domain becomes a region.)

Lemma 8.5. *Any monogon in a link projection can be emptied by a finite sequence of fair and good moves.*

Proof. We will mostly rely on Reidemeister I-a moves, cyclic Reidemeister II-a moves, (global) noncyclic Reidemeister II-a moves, and noncyclic Reidemeister III moves, cf. Lemma 8.3. In addition, we may sometimes move a connected component of a link projection from one region (of the rest of the projection) to another, which is obviously a fair move.

By starting with an innermost one, we may assume that the monogon contains no other monogons. In other words, we assume that no arc inside the monogon intersects itself.

We proceed by induction on the number of crossings inside and on the boundary of the monogon. If this number is 1, then the monogon only contains disjoint simple closed curves and those can always be moved out of it by fair moves. Our goal is to show that if our number of crossings is at least 2, then a sequence of the specified moves can be performed to reduce it.

If a global noncyclic Reidemeister II-a move is possible, then we can achieve our goal immediately. That is the case if the monogon contains a simple closed curve within the link projection that is intersected by another arc. Hence from now on, we may assume that the monogon does not contain closed curves at all, only (directed, non-self-intersecting) arcs that connect two of its boundary points.

Furthermore, we may assume that any bigon formed by our arcs (including the monogon's boundary) is cyclic.

Let γ_0 be the arc in the monogon whose exit point occurs last along the boundary. (We may assume without loss of generality that the monogon is clockwise oriented as in Figure 13.) Let B_0 be the bigon determined by γ_0 and the monogon. Let us also put $\alpha_0 = \partial B_0 \setminus \gamma_0$ for the monogon arc that bounds B_0 .

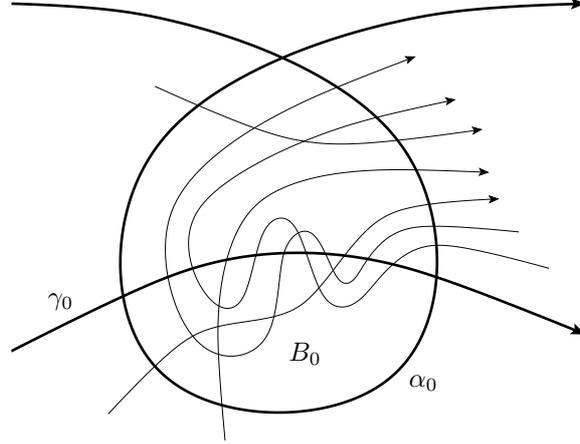


Figure 13. Monogon with arcs that form cyclic bigons with it.

For the remaining part of the proof, most bigons inside our monogon will be considered with their two boundary arcs designated as *upper* and *lower*. For the bigon B_0 of the previous paragraph, α_0 is the lower arc and γ_0 is the upper one. By our choice of γ_0 , we see that B_0 has the property that

$$\text{no arc exits the bigon through its lower boundary arc.} \quad (7)$$

We will use a recursive procedure to define a nested sequence of bigons B_0, B_1, \dots , all of which will satisfy property (7). All bigons in the sequence will of course be cyclic and have their boundary oriented the same way (i.e., clockwise).

Keeping in mind our assumption that all bigons are cyclic, it is easy to see that if a bigon has property (7), then there can only be two kinds of arcs crossing it. The two cases are depicted in Figure 14, and will be referred to as *cutting* and *biting* arcs, respectively.



Figure 14. Cutting (left) and biting (right) arcs in a bigon.

Let now B_i be a bigon with upper arc γ_i and lower arc α_i so that property (7) holds. Assume that B_i has a biting arc. Let α_{i+1} denote the biting arc whose exit point from B_i is first along γ_i . Let B'_i be the bigon formed by γ_i and α_{i+1} , see Figure 15. If no arc (inside B_i) crosses α_{i+1} downward, then B'_i , bounded by lower arc α_{i+1} and upper arc $\gamma_{i+1} = \gamma_i$ (the latter appropriately shortened) has property (7) and we denote it by $B_{i+1} = B'_i$. Otherwise, let p_i denote the last point along α_{i+1} where an arc γ_{i+1} exits B'_i .

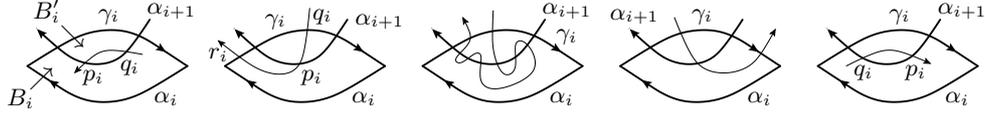


Figure 15. Hypothetical positions of the arc γ_{i+1} . Only the last option can be reconciled with our assumptions.

We claim that the portion of γ_{i+1} in B'_i may only be arranged as in the rightmost panel of Figure 15. The reasons for this are as follows. Let q_i denote the point where γ_{i+1} enters B'_i for the last time before leaving at p_i .

- The point q_i may not be on α_{i+1} and before p_i because then α_{i+1} and γ_{i+1} would form a noncyclic bigon.
- Assuming that q_i is on γ_i , we consider the point r_i where γ_{i+1} leaves B_i for the first time after passing through p_i . By property (7), r_i is on γ_i .
- The point r_i may not occur on γ_i before $\partial B'_i$ because of the way that α_{i+1} was chosen.
- The point r_i may not be on $\partial B'_i$ and (with respect to the orientation of γ_i) before q_i because in order to get there, γ_{i+1} would have to form a noncyclic bigon with α_{i+1} . (See Figure 15, middle panel. Recall that γ_{i+1} does not intersect itself. Hence it would have to intersect the portion of α_{i+1} that starts at p_i and ends at γ_i . Consider the first such intersection point along γ_{i+1} after p_i . This and p_i are the vertices of the desired bigon.)
- The point r_i may not be after q_i along γ_i because then γ_{i+1} and γ_i would form a noncyclic bigon.

In this case, then, we let the bigon B_{i+1} be bounded by the lower arc α_{i+1} and the upper arc γ_{i+1} . Because of the way p_i was chosen, B_{i+1} satisfies property (7).

We continue constructing the sequence B_0, B_1, \dots until we arrive at a bigon B_N (with lower boundary arc α_N and upper boundary arc γ_N) that has no biting arcs. It can of course have cutting arcs and, unfortunately, those may still form cyclic bigons inside B_N . Those bigons do not necessarily satisfy property (7), but nonetheless the following argument is valid.

If B_N is empty then we eliminate it by a (cyclic) Reidemeister II-a move and by doing so we reduce the number of crossings in the original monogon by 2. Otherwise, consider the cutting arc κ in B_N that has the last entry point along

α_N . To the left of κ ('left' being defined in terms of the orientation of κ) there is a triangular part T of B_N into which arcs can enter only through κ . If κ bounds some (of course cyclic) bigons to its left, then consider those bigons among them that are minimal with respect to containment. Now from these minimal bigons, take the one that is 'farthest along' κ (i.e., the arc that forms the bigon with κ has the last entry point into T along κ). Let us declare (the appropriate portion of) κ as the lower boundary arc for this bigon. Then the bigon satisfies condition (7) and hence it can be the starting element in a new nested sequence of bigons, just like B_0 was the starting element of the first sequence (the only difference is that the new bigons have counterclockwise oriented boundary). So we go back to that part of the argument and iterate.

If κ does not bound bigons on its left side, then all arcs crossing it have to do so from right to left and from there they have to proceed through the triangle T to γ_N . If no arcs cross κ then we 'pull κ out' of B_N by applying a noncyclic Reidemeister III move to T and then repeat the procedure from the beginning of the previous paragraph¹². If some arcs do cross κ then consider the one with the first crossing point and call it κ_1 . On the portion of κ_1 between κ and γ_N , we can use the same reasoning as above to conclude that either κ_1 is not crossed by any other arc, or it is only crossed from right to left, or that κ_1 bounds a bigon satisfying property (7). In the latter two cases we may iterate the ideas we presented so far.

In particular in the middle case, when κ_1 is crossed but only from right to left, we consider the arc κ_2 with the first crossing point. We argue that κ_2 exits T through γ_N without further intersecting κ_1 and apply the rest of the same reasoning that we used for κ and κ_1 , too. If this case keeps coming up, i.e., if we never loop back to the case of a (smaller) bigon with property (7), then eventually there will be an arc κ_M with no intersection points on it (between κ_{M-1} and γ_N , that is). To simplify notation in the next part of the argument, we may as well assume that $M = 1$, that is that κ_1 is void of intersection points with other arcs (which was the first of the three cases at the end of the previous paragraph).

We return to κ and continue along it beyond the startpoint of κ_1 to the next crossing point (which, if it exists, is again from right to left). We follow this new arc to γ_N and apply the same separation of cases to it as we did with κ_1 . (A crossing from left to right leads to a bigon with property (7) – recall that κ_1 has no intersection points on it. The first crossing from right to left gives us a new arc to investigate. If that does not exist, then look for the next intersection point along κ .) Eventually, after exhausting all possible bigons and intersection points, what we find is an empty triangle with one of its sides along γ_N (recall that B_N has no biting arcs). It is noncyclic because its other two sides are cutting arcs. To explain the last step in the proof of the Lemma, it suffices to assume that the triangle bounded by κ , κ_1 , and γ_N is already empty.

Indeed if that is the case, then we can apply a noncyclic Reidemeister III move

¹²An inductive argument on the number of cutting arcs is suppressed here. It is fairly obvious and we felt that the proof was complicated enough as is.

to the triangle. That does not change the number of cutting arcs but it removes the intersection of κ and κ_1 from B_N , so that if we return to B_N now and run our ‘simplification-searching algorithm’ on it again, we have one less crossing (that is to say, a ‘quantifiable simpler situation,’ which can be made precise with another trivial inductive loop) to deal with. This completes the proof of the Lemma. \square

So far in this paper we treated skein computation trees from the point of view of the Seifert graph. There is however an older approach, based on the notion of a descending diagram [2]. We will borrow some ideas from that context to prove an equivalent version of Proposition 3.6, which is the main result of this section.

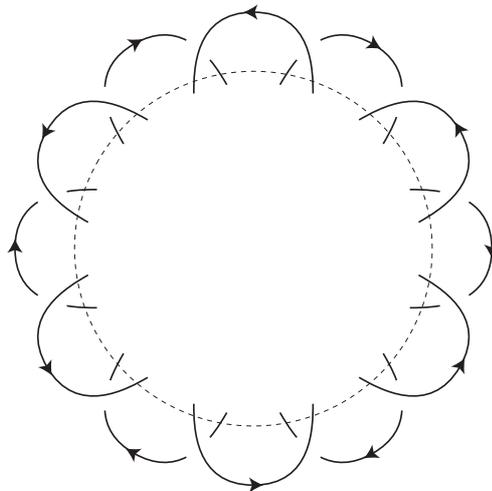


Figure 16. An alternating contour with $k = 6$ outer over-arcs and outer under-arcs.

Theorem 8.6. *Let D be a link diagram containing the part shown in Figure 16, and L the associated link. That is, outside of the dashed circle, there is no piece of D other than the $2k$ arcs shown, where $k \geq 1$. Then Morton’s inequality (Theorem 3.2) is not sharp; in fact, we have*

$$\maxdeg_z(P_L(v, z)) \leq n(D) - s(D) - 1. \quad (8)$$

Proof. We will refer to the link diagrams described in the Theorem (and in Proposition 3.6) as having an *alternating contour*. A diagram with an alternating contour is manifestly not alternating, but we are confident that this will not lead to confusion. This notion is related to plane graphs with an alternating contour (cf. Section 3), but it is best to treat the two concepts as separate.

Let us call the k arcs passing over the other arcs appearing in Figure 16 *outer over-arcs*, while the k arcs passing under the outer over-arcs are the *outer under-arcs*. Let us denote the dashed circle shown in the diagram, along which the $4k$ endpoints of all outer arcs lie, with C .

We are going to use induction on the number N of crossings that the diagram D has inside C . When that number is 0, we will first argue that outer over-arcs and outer under-arcs may not lie along the same component of L . Let us number the $4k$ points of $C \cap D$ counterclockwise around C . A quick examination of Figure 16 shows that the modulo 4 remainder class of each point tells exactly if it is a start- or an endpoint and whether of an over-arc or an under-arc. For any arc of D across the interior of C that is obtained by continuing an outer over-arc, its endpoint minus its startpoint has to equal 1 modulo 4: the difference has to be 1 or 3 so that the other arcs across C may complete a crossingless matching of the $4k$ points, but 3 is ruled out if we consider the prescribed orientations. This means that the arc leaves C at another over-arc and the claim follows.

Now the previous paragraph implies that when $N = 0$, the link L is an unlink of at least 2 components: The components containing the outer over-arcs form an unlink (with a crossingless diagram, no less), the same is true for the components through the under-arcs, and the two unlinks are separated by a copy of S^2 . In addition, L may have components that project onto disjoint simple closed curves in the interior of C . Let the number of the latter be m . Since the HOMFLY polynomial of a c -component unlink is $(v^{-1}z^{-1} - vz^{-1})^{c-1}$, we have $\max \deg_z(P) \leq -m - 1$ in our case. Hence it suffices to show that $-m - 1 \leq n(D) - s(D) - 1 = 2k - s(D) - 1$. But as every Seifert circle of D (in the $N = 0$ case) is either a trivial connected component or it has to pass through the midpoint of at least one outer arc, this follows easily.

As to the inductive step, let us start with a general observation. Let (L_+, L_-, L_0) be a skein triple and suppose that L_+ (or L_-) and L_0 satisfy the inequality (8):

$$\begin{aligned} \max \deg_z(P_{L_\pm}(v, z)) &\leq n(D_\pm) - s(D_\pm) - 1; \\ \max \deg_z(P_{L_0}(v, z)) &\leq n(D_0) - s(D_0) - 1. \end{aligned}$$

Since $s(D_+) = s(D_-) = s(D_0)$ and $n(D_+) = n(D_-) = n(D_0) + 1$, we obtain the inequalities

$$\begin{aligned} \max \deg_z(P_{L_\pm}(v, z)) &\leq n(D_\mp) - s(D_\mp) - 1, \\ \max \deg_z(P_{L_0}(v, z)) &\leq n(D_\mp) - s(D_\mp) - 2. \end{aligned}$$

On the other hand, as $P_{L_\mp}(v, z) = v^{\mp 2}P_{L_\pm}(v, z) \mp zv^{\mp 1}P_{L_0}(v, z)$ from (2), we have

$$\max \deg_z(P_{L_\mp}(v, z)) \leq \max\{\max \deg_z(P_{L_\pm}(v, z)), \max \deg_z(P_{L_0}(v, z)) + 1\}.$$

Therefore we see that

$$\max \deg_z(P_{L_\mp}(v, z)) \leq n(D_\mp) - s(D_\mp) - 1.$$

This means that if L_0 and one member of the pair L_+, L_- satisfy (8), then so does the other member.

Let us now fix k and assume that we have a diagram D as in the Theorem, with N crossings inside C , as well as that (8) holds whenever the number of crossings inside C is less than N . Because in a skein triple, L_0 always has one less crossing than L_+ or L_- , the observation above means that it suffices to show the following:

For every diagram D as in the Theorem, with $2k$ outer arcs, (9)
 it is possible to change some of the N crossings inside C so
 that (8) holds for the link thus obtained.

We wish to apply isotopies and to rely on Lemma 8.4 to prove (9). If a Reidemeister move is possible for the corresponding link projection, then it becomes possible for the diagram as well after changing at most one crossing. None of the Reidemeister moves listed in Lemma 8.3 and used in the proof of Lemma 8.5 increases the number of crossings.

A special note is in order on global noncyclic Reidemeister II-a moves. When we apply such a move to a link diagram, first we change crossings along the boundary of the bigon (including at most one of the vanishing crossings). This is done while the total number of crossings inside C is N or less, so that we can rely on the inductive hypothesis. Then, ‘during’ the isotopy, the number of crossings may temporarily exceed N but that is fine since no crossing change is necessary at those stages.

Thus, with the help of Lemma 8.5 and the fact that a Reidemeister I-a move is fair, (9) reduces to

For every diagram D as in the Theorem with $2k$ outer arcs and at (10)
 most N crossings inside C , so that arcs inside C do not self-intersect,
 it is possible to change some of the crossings inside C so that (8) holds
 for the resulting link.

We keep insisting on making changes inside C only so that our main induction can proceed. We first establish (10) in two special cases.

A We assume that there exists an outer over-arc in D whose endpoints are joined by another arc α inside the dashed circle. See Figure 17. By changing some crossings along α , let us ‘float to the top’ the unknot component arising from our assumption, and then let us separate it from the rest of the diagram as in the right panel of Figure 17. We denote the resulting diagram with \tilde{D} .

We will show that $n(D) - s(D) \geq n(\tilde{D}) - s(\tilde{D}) + 2$, i.e., that (10) for D follows from applying Morton’s inequality to \tilde{D} . Let N_α be the number of crossings of D along α (in particular, inside C). Let D' be the diagram that results from smoothing all N' crossings of D away from α and inside C , as well as the $2k$ crossings outside of C . Let s' be the number of components in D' . Then we have

$$n(D) - s(D) \geq 2k + N_\alpha + N' - (s' + N_\alpha) = 2k + N' - s',$$

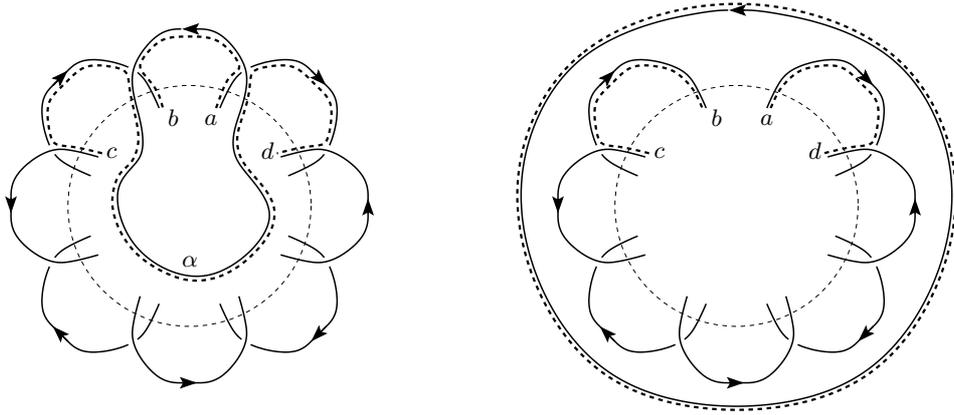


Figure 17. Left: The diagram D in case A. The dashed arcs indicate pieces of components of D' . Right: The diagram \tilde{D} , with several pieces of its Seifert circles.

because every time we smooth one of the remaining N_α crossings of D' , the number of components changes by ± 1 . On the other hand we have

$$n(\tilde{D}) - s(\tilde{D}) \leq 2k - 2 + N' - s',$$

from which our claim follows. To see why there are at least s' Seifert circles in \tilde{D} , notice that those components of D' that do not pass through the points a , b , c , or d of Figure 17 are found in \tilde{D} as well. There are 1 or 2 components of D' that do pass through those points and \tilde{D} has at least 2 components (the unknot we pulled out, and the one through, say, a) to account for them.

B We assume that as we continue each outer over-arc of D through the interior of C , we emerge at an outer under-arc. Let α_0 be one of the arcs that we have just described. The terminal point of α_0 is adjacent along C to the initial point b_1 of an outer over-arc. If we follow D backward from b_1 , then we hit either C or α_0 first. In the latter case, a noncyclic bigon is formed so that one of its corners is outside of C but its other corner, as well as each arc of D inside the bigon, is inside C . Therefore it is possible to change crossings inside C so that a global noncyclic Reidemeister II-a move becomes possible for D . Hence in this case, we are done by Lemma 8.3 and Theorem 3.2.

In the case when the arc through C that ends at b_1 is disjoint from α_0 , let us call the arc α_1 and note that by our assumption in case B, the initial point a_1 of α_1 has to be the terminal point of an outer under-arc. Now a_1 is adjacent along C to the terminal point a_2 of an outer over-arc. Notice that a_2 and α_0 are on opposite sides of α_1 . From here we iterate our argument: If the arc of D that starts at a_2 hits α_1 before C , then a noncyclic bigon and hence a good move can be found. Otherwise, follow the arc to its terminal point b_2 on C ,

which is necessarily on an outer under-arc, denote the arc a_2b_2 with α_2 , take the point b_3 , adjacent along C to b_2 , where an outer over-arc starts, note that α_2 separates α_1 and b_3 , and continue the iteration.

Our argument above produces a sequence $\alpha_0, \alpha_1, \dots$ of parallel, disjoint chords in C that are arranged monotonously. Since such a sequence cannot be infinite, it is guaranteed that after finitely many steps we find a global noncyclic Reidemeister II-a move, which completes the proof of case B.

In the rest of the proof, we are going to verify (10) by a secondary induction on k (keeping N fixed). When $k = 1$, the diagram D falls under one of the cases A and B above.

Let us now assume that (10) holds whenever the number of outer over-arcs is less than k and let D be a link diagram with an alternating contour and k outer over-arcs. Having established cases A and B, we may assume that there exists an outer over-arc in D which, when continued across C , meets a different outer over-arc. Let this arc across C be called α .

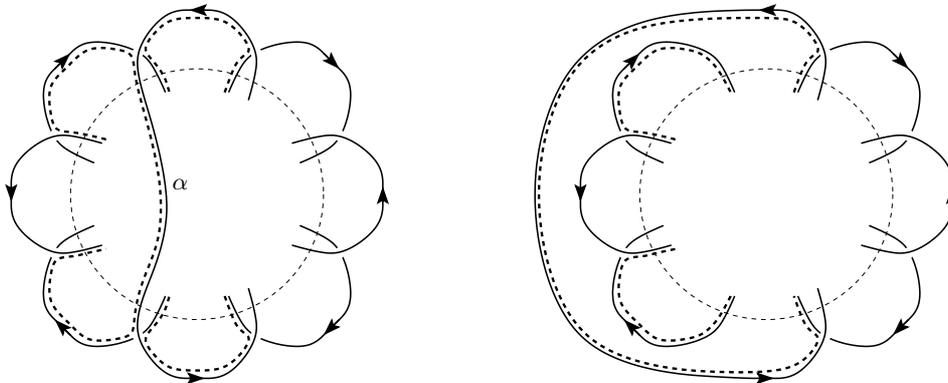


Figure 18. Left: Diagram D with an alternating contour. Dashed arcs indicate components of D' . Right: the diagram \tilde{D} .

Given a diagram D as in the left panel of Figure 18, let s' be the number of components in the diagram D' that results from smoothing all crossings in D except for those along α . Let N_α be the number of crossings on α and let N' be the number of crossings of D inside C that are not on α . Then there is a total of $2k + N_\alpha + N'$ crossings in D and the number of Seifert circles is at most $s' + N_\alpha$ by the same reason as before. Hence we have

$$2k + N' - s' \leq n(D) - s(D).$$

After changing crossings if necessary, let us pull α out as in the right panel of Figure 18 and call the resulting diagram \tilde{D} . Then the number of crossings in \tilde{D} is $2k - 2 + N'$ and the number of Seifert circles is at least $s' - 2$. (It is possible

that the arcs of D' indicated in the left panel of Figure 18 belong to three different components, whereas the arcs of \tilde{D} indicated on the right might belong to just one component. But it is easy to see that nothing worse than that can happen.) Therefore we have

$$n(\tilde{D}) - s(\tilde{D}) \leq 2k + N' - s',$$

meaning that the isotopy from D to \tilde{D} was at least a fair move. But since \tilde{D} is also a diagram with an alternating contour but with less than k outer over-arcs, we are done by Lemma 8.4 and the inductive hypothesis.

This completes the secondary induction (on k), thus we have proved (10) and hence the inductive step in our main induction (on N) is now also established. \square

References

- [1] D. Chebikin and P. Pylyavskyy, A family of bijections between G -parking functions and spanning trees. *J. Combin. Theory Ser. A* **110** (2005), no. 1, 31–41.
- [2] P. R. Cromwell, Homogeneous links. *J. London Math. Soc.* **39** (1989), no. 3, 535–552.
- [3] R. Gopakumar and C. Vafa, On the gauge theory/geometry correspondence. *Adv. Theor. Math. Phys.* **3** (1999), no. 5, 1415–1443.
- [4] A. E. Holroyd, L. Levine, K. Mészáros, Y. Peres, J. Propp, and D. B. Wilson, Chip-Firing and Rotor-Routing on Directed Graphs. In *In and Out of Equilibrium 2* (eds. V. Sidoravicius and M. E. Vares), Progress in Probability vol. 60, Birkhäuser, 2008.
- [5] F. Jaeger, Tutte polynomials and link polynomials. *Proc. Amer. Math. Soc.* **103** (1988), no. 2, 647–654.
- [6] V. F. R. Jones, Hecke algebra representations of braid groups and link polynomials. *Ann. Math.* **126** (1987), 335–388.
- [7] A. Juhász, T. Kálmán, and J. Rasmussen, Sutured Floer homology and hypergraphs. *Math. Res. Lett.* **19** (2012), no. 6, 1309–1328.
- [8] T. Kálmán, A version of Tutte’s polynomial for hypergraphs. *Adv. Math.* **244** (2013), no. 10, 823–873.
- [9] T. Kálmán and A. Postnikov, Root polytopes, Tutte polynomials, and a duality theorem for bipartite graphs. arXiv:1602.04449.
- [10] C. Merino, Chip firing and the Tutte polynomial. *Ann. Comb.* **1** (1997), no. 3, 253–259.
- [11] H. R. Morton, Seifert circles and knot polynomials. *Math. Proc. Camb. Phil. Soc.* **99** (1986), 107–109.
- [12] K. Murasugi and J. Przytycki, The skein polynomial of a planar star product of two links. *Math. Proc. Camb. Phil. Soc.* **106** (1989), 273–276.
- [13] H. Ooguri and C. Vafa, Knot invariants and topological strings. *Nucl. Phys. B* **577** (2000), no. 3, 419–438.
- [14] A. Postnikov, Permutohedra, Associahedra, and Beyond. *Int. Math. Res. Not.* 2009, no. 6, 1026–1106.

- [15] A. Postnikov and B. Shapiro, Trees, parking functions, syzygies, and deformations of monomial ideals. *Trans. Amer. Math. Soc.* **356** (2004), no. 8, 3109–3142.
- [16] E. Swartz, From polytopes to enumeration. <http://www.math.cornell.edu/~ebs/math455.pdf>.

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

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

Hitoshi Murakami, Graduate School of Information Sciences, Tohoku University, Aramaki aza Aoba 6-3-09, Aoba-ku, Sendai 980-8579, Japan

E-mail: starshea@tky3.3web.ne.jp