FROM DIMERS TO WEBS

CHRIS FRASER, THOMAS LAM, AND IAN LE

ABSTRACT. We formulate a higher-rank version of the *boundary measurement map* for weighted planar bipartite *networks* in the disk. It sends a network to a linear combination of SL_r-*webs*, and is built upon the r-fold dimer model on the network. When r equals 1, our map is a reformulation of Postnikov's boundary measurement used to coordinatize positroid strata. When r equals 2 or 3, it is a reformulation of the SL_2- and SL_3-*web immanants* defined by the second author. The basic result is that the higher rank map factors through Postnikov's map. As an application, we deduce generators and relations for the space of SL_r-webs, reproving a result of Cautis-Kamnitzer-Morrison. We establish compatibility between our map and restriction to positroid strata, and thus between webs and total positivity.

1. INTRODUCTION

The Grassmannian $Gr(k, n)$ of k-planes in $\mathbb{C}^n$ is an algebraic variety which is well-loved in combinatorics. This paper links two combinatorial tools – briefly, *dimer configurations* and *webs* – which have been used to study $Gr(k, n)$ and its (homogeneous) coordinate ring. Both approaches have a similar flavor – each involves certain planar diagrams in a disk, and each relies heavily on local diagrammatic moves/relations amongst such diagrams. We show that this resemblance is not coincidental, and that these approaches are dual in a sense that we make precise. Moreover, statements on each side can be translated to the other to give meaningful consequences.

The first approach starts with a choice of *network N*, meaning a planar bipartite graph in the disk whose edges are weighted by nonzero complex numbers. Such a network comes with two parameters: n (the number of boundary vertices) and k (the *excedance*, cf. (2.1)). The key operation is Postnikov's *boundary measurement map*

$$(1.1) \qquad\qquad N \mapsto (\Delta_I(N))_{I \in \binom{[n]}{k}} \in Gr(k, n)$$

sending a network N to its $\binom{n}{k}$ *boundary measurements*. Each boundary measurement $\Delta_I(N)$ is a complex number obtained by summing over *dimer configurations of N* whose *boundary data* is I. The striking feature is that for any network, the $\Delta_I(N)$ satisfy the well-known *Plücker relations*, so the image of the boundary measurement is a point in the Grassmannian.

The second approach considers a distinguished class of functions on $Gr(r, n)$, indexed by planar diagrams known as SL_r-*webs*. Let us consider the space $\mathcal{W}(r, n) = Hom_{SL_r}((\mathbb{C}^r)^{\otimes n}, \mathbb{C})$, i.e. the vector space of SL_r-invariant multilinear function of n vectors. The homogeneous coordinate ring $\mathbb{C}[Gr(r, n)]$ can be viewed as an algebra of SL_r-invariants, and $\mathcal{W}(r, n)$ is a certain subspace of $\mathbb{C}[Gr(r, n)]$. We will call elements of $\mathcal{W}(r, n)$ *tensor invariants*.

An SL_r-*web diagram* is a planar bipartite graph in the disk, with its edges labeled by positive integers so that the labels around each internal vertex sum to r. (See the

Key words and phrases. dimer, web, boundary measurement, positroid, Grassmannian.

left hand side of (5.4) for an example when $r = 4$. Edges with label 1 are suppressed.) We denote by $\mathcal{FS}(r)$ the vector space of formal sums of SL_r-web diagrams (it is related to the *free spider category* [1]). Each web diagram determines an element[1] of $\mathcal{W}(r, n)$. Intuitively, web diagrams determine tensor invariants as follows: an edge in a web diagram with label a corresponds to a copy of the exterior power $\bigwedge^a(\mathbb{C}^r)$, and each interior vertex v corresponds to an SL_r-invariant map between the exterior powers adjacent to v. The web diagram indicates how to compose these building blocks to create more complicated tensor invariants. The map $\mathcal{FS}(r) \to \mathcal{W}(r, n)$ is surjective; i.e. SL_r-invariants coming from webs span the space of tensor invariants.

Webs can be used to study more general tensor invariant spaces, and other homogeneous pieces of $\mathbb{C}[\mathrm{Gr}(k, n)]$, but in this introduction, we will focus on a particular situation where our results are easiest to state. Thus, we restrict attention to networks N whose number of boundary vertices n is a multiple of its excedance k, i.e. $n = kr$. For these N, we construct an r-fold boundary measurement map

$$(1.2) \qquad\qquad N \mapsto \mathrm{Web}_r(N) \in \mathcal{W}(r, n),$$

sending a network to a particular formal sum of webs. Its key feature (our Theorem 4.4) is that it factors through the boundary measurement map (1.1). That is – if N and N' are two networks with the same boundary measurements (and hence which give the same point in the affine cone over the Grassmannian $\widetilde{\mathrm{Gr}}(k, n)$), then $\mathrm{Web}_r(N) = \mathrm{Web}_r(N')$ as tensor invariants, even though these will typically look different as elements of $\mathcal{FS}(r)$. When r is 2 or 3, this reduces to the construction of Temperley-Lieb immanants and web immanants via the double-dimer and triple-dimer models studied by the second author [10]. We also show (Theorem 8.1) that (1.2) induces a natural isomorphism $\mathcal{W}(k, n) \cong \mathcal{W}(r, n)^*$. Moreover, both spaces $\mathcal{W}(k, n)$ and $\mathcal{W}(r, n)$ are irreducible S_n-modules, where S_n acts on the tensor invariant spaces by permuting the vectors. As S_n-modules, we have that $\mathcal{W}(k, n) \cong \mathcal{W}(r, n)^* \otimes \epsilon$ where ϵ is the sign representation. Thus, we get a canonically defined S_n-equivariant pairing between SL_k- and SL_r-invariant spaces. For small values of k and r, we draw the resulting "web duality pictures" in an Appendix.

We will lay out the contents of this paper while mentioning the applications of our main construction. Section 2 reviews the boundary measurement map for networks via dimer configurations [10], as well as the local moves on networks which preserve the boundary measurements. We also review in this section how boundary measurements, networks, and local moves, can be used to study the *positroid stratification* of the Grassmannian. Section 3 gives the basics of tensor invariants and web diagrams. In Section 4 we make our main definition, i.e. the r-fold boundary measurement map (1.2), and introduce the closely related *immanant map*. Section 5 gives a self-contained proof that the r-fold boundary measurement factors through Postnikov's boundary measurement map. The key tool is a lemma about sign-coherence for webs that we conjecture is related to total positivity.

Section 6 discusses one of our main applications. One of the deepest results on the boundary measurement map is that if networks N and N' have the same boundary measurements, then N' is connected to N by a sequence of local moves [16]. Likewise, a

---

[1]in fact, the diagrams defined in this introduction only determine a tensor invariant up to a sign. The sign is determined by a procedure called a *tagging* of the web.

guiding problem in the history of web combinatorics was to find a complete set of diagrammatic moves describing the kernel of the map $\mathcal{FS}(r) \to \mathcal{W}(r,n)$. This problem was solved for $r = 3$ by Kuperberg [9], studied further by Kim [6] and Morrison [13], and settled for all r by Cautis, Kamnitzer, and Morrison in [1]. We show that the completeness of the relations in [1] follows from Theorem 4.4 and the connectedness result for networks. We remark that the results in [1] hold in greater generality than ours – those authors work in a quantum deformation of our setting. We are hopeful that an understanding of this will lead to a quantum deformation of the dimer model.

In Section 7, we make a connection between webs and *positroid varieties*. Here, we extend the results of Sections 4 and 5 to the context of positroids. The *totally nonnegative Grassmannian* $\mathrm{Gr}(k,n)_{\geq 0}$ is the set of real points in $\mathrm{Gr}(k,n)$ with nonnegative Plucker coordinates [16]. The matroid $\mathcal{M}$ of a totally nonnegative point $x \in \mathrm{Gr}(k,n)_{\geq 0}$ is known as a *positroid*. The Zariski-closures of the positroid strata are known as *positroid varieties* [7]. The homogeneous coordinate ring $\mathbb{C}[\Pi]$ of a positroid variety $\Pi = \Pi_{\mathcal{M}}$ is a quotient of $\mathbb{C}[\mathrm{Gr}(k,n)]$. We show that the homogeneous pieces of $\mathbb{C}[\Pi]$ are dual to certain naturally defined subspaces of tensor invariants (cf. (7.4)), and suggest a way to compute the dimensions of these pieces using webs. We believe that our duality reflects deep relations between positroids and webs.

We denote by $[n]$ the set of positive integers $\{1, \ldots, n\}$, and by $\binom{S}{k}$ the collection of k-element subsets of a set S.

2. Networks, boundary measurements, and positroid strata

2.1. **The dimer model in the disk.** We denote by $\mathrm{Gr}(k,n)$ the Grassmannian of k-dimensional subspaces in a fixed n-dimensional complex vector space. It has a projective embedding $\mathrm{Gr}(k,n) \hookrightarrow \mathbb{P}^{\binom{n}{k}-1}$, and we denote by $\widetilde{\mathrm{Gr}}(k,n) \subset \mathbb{C}^{\binom{n}{k}}$ the affine cone over $\mathrm{Gr}(k,n)$ with respect to this embedding. Points in the affine cone are collections of $\binom{n}{k}$ coordinates satisfying the well-known *Plücker relations*. The homogeneous coordinate ring $\mathbb{C}[\mathrm{Gr}(k,n)]$ (or equivalently, the coordinate ring of the affine cone $\widetilde{\mathrm{Gr}}(k,n)$) is generated by *Plücker coordinates* $(\Delta_I)_{I \in \binom{[n]}{k}}$. If we represent a point in $x \in \widetilde{\mathrm{Gr}}(k,n)$ by a $k \times n$ matrix of maximal rank, then $\Delta_I(x)$ is the $k \times k$ minor of this matrix with columns I.

By a *planar bipartite graph in the disk* we mean a graph G embedded in a closed disk, with its vertices colored in two colors (black and white) such that edges join vertices of opposite color. Furthermore, label the vertices on the boundary of the disk $1, \ldots, n$ in

counterclockwise order, and require that that ith *boundary vertex* is incident to at most one *boundary edge* b_i.

Furthermore, for the sake of simplicity, throughout this paper we will require that each of the boundary vertices of G is black.

By a *network* N we will mean a planar bipartite graph in the disk whose edges have been weighted by nonzero complex numbers. In examples, an edge drawn without an edge weight is implicitly assumed to have edge weight equal to 1.

A *dimer configuration* on N (or *almost perfect matching* of N), is a collection π of edges of N such that each interior vertex is used exactly once in π (and each boundary vertex is used one or zero times). The *boundary subset* $\partial(\pi) \subset [n]$ is the set of boundary vertices that are used in π. The cardinality $k = |\partial(\pi)|$ depends only on the underlying bipartite graph G, not on the choice of π or the edge weights on N. Explicitly:

(2.1) $k = |\partial(\pi)| =$ no. of interior white vertices in G minus no. of interior black vertices.

We call the number k in (2.1) the *excedance* of N.

Remark 2.1. The requirement that boundary vertices of a planar bipartite graph G are always black simplifies our statements, but can be removed. Any graph G', possibly with white vertices on the boundary, can be turned into a graph G whose boundary vertices are all black, by adding one edge at each white boundary vertex (if i is a white boundary vertex, we drag it into the interior of the disk, and connect it to a newly created black boundary vertex by a newly created edge). Dimer configurations π on G with boundary $\partial(\pi) = I$ are in bijection with dimer configurations π' of G' such that I is the union of the black vertices used in π' with the white vertices not used in π', cf. [10] for further details.

For a network N and a dimer cover π on N, the *weight* $\mathrm{wt}(\pi)$ is the product of the weights of the edges used in π.

The *boundary measurement* $\Delta_I(N)$ is a weight generating function for dimer configurations with boundary I:

(2.2) $$\Delta_I(N) = \sum_{\pi: \partial(\pi)=I} \mathrm{wt}(\pi).$$

Proposition 2.2 (Kuo [8], Postnikov-Speyer-Williams [17], Lam [11])**.** *Let N be a network of excedance k, with n boundary vertices, and with at least one almost perfect matching. Then the boundary measurements $(\Delta_I(N))_{I \in \binom{[n]}{k}} \in \mathbb{C}^{\binom{n}{k}}$ determine a point $\tilde{X}(N)$ in the affine cone $\widetilde{\mathrm{Gr}}(k,n)$, and thus a point $X(N) \in \mathrm{Gr}(k,n)$.*

That is, the boundary measurements satisfy the Plücker relations. If N has no almost perfect matchings, then all of its Plücker coordinates are zero, so $X(N)$ is not a well-defined point in $\mathrm{Gr}(k,n)$. We make the standing assumption that all networks N considered in this paper have an almost perfect matching.

2.2. **Local moves.** It is easy to verify that the following local moves can be applied to a network N to yield a new network N' satisfying $X(N') = X(N)$. Thus the Plücker coordinates for $\tilde{X}(N)$ and $\tilde{X}(N')$ differ by a common scalar α and we write $\tilde{X}(N) = \alpha \tilde{X}(N')$.

(G) Gauge equivalence: If edges $e_1, e_2, \ldots, e_d$ are the edges incident to an interior vertex v, we can rescale all of their edge weights by the same constant $\alpha \in \mathbb{C}^*$. The resulting network N' satisfies $\tilde{X}(N) = \alpha\tilde{X}(N')$.

(M1) Spider move, square move, or urban renewal: assuming the leaf edges of the spider have been gauge fixed to 1, the transformation is

$$(2.3) \qquad a' = \frac{a}{ac + bd} \qquad b' = \frac{b}{ac + bd} \qquad c' = \frac{c}{ac + bd} \qquad d' = \frac{d}{ac + bd}$$
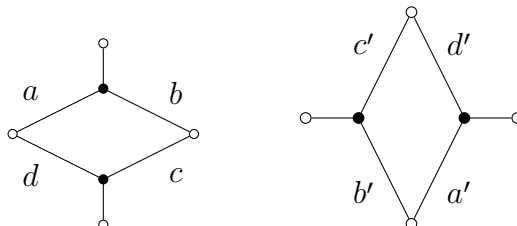
and satisfies $\tilde{X}(N) = (ac + bd)\tilde{X}(N')$.

(M2) two-valent vertex removal. If v has degree two, we can gauge fix both incident edges (v, u) and (v, u') to have weight 1, then contract both edges (that is, we remove both edges, and identify u with u'). Note that if v is a two-valent vertex adjacent to boundary vertex i, with edges (v, i) and (v, u), then this move can only be applied when u has degree at most two.

(R1) Multiple edges with same endpoints is the same as one edge with sum of weights.

(R2) Leaf removal: Suppose v is a leaf, and (v, u) the unique edge incident to it. Then we can remove both v and u, and all edges incident to u. If (u, v) is a boundary edge b_i, then the leaf cannot be removed.

(R3) Dipoles (two degree one vertices joined by an edge) can be removed.

On the other hand, the following is one of the deepest results on the combinatorics of networks:

Theorem 2.3. *[Postnikov [16]] If N and N' satisfy $X(N) = X(N')$, then they are connected to each other by a finite sequence of these moves. If furthermore both of these networks have the minimal number of faces in their move-equivalence class, then they are connected by a finite sequence using (M1) and (M2) only.*

2.3. **Positroids.** Considering boundary measurement maps for various bipartite graphs G leads to a special stratification of the Grassmannian by *positroid varieties*. The positroid varieties are distinguished in many senses – they are exactly the varieties that can be obtained by projecting Richardson varieties from the flag variety [7], they are exactly the compatibly split Frobenius subvarieties of $\mathrm{Gr}(k, n)$ with respect to the standard splitting [7], the (open versions of) positroid varieties are exactly the symplectic leaves with respect to a Poisson structure on $\mathrm{Gr}(k, n)$ [4]. We only review what we will need here and refer the reader to [7; 11].

The *totally nonnegative Grassmannian* $\mathrm{Gr}(k, n)_{\geq 0}$ consists of points in the Grassmannian $\mathrm{Gr}(k, n)$ that can be given by $k \times n$ matrices with real entries, all of whose Plücker coordinates are nonnegative.

For any point $x \in \mathrm{Gr}(k, n)$, the *matroid* of x is the realizable matroid $\mathcal{M}(x)$ formed by the subsets $I \in \binom{[n]}{k}$ such that $\Delta_I(x) \neq 0$. The *matroid variety* associated to a matroid $\mathcal{M}$ is the closure in $\mathrm{Gr}(k, n)$ of $\{x \in \mathrm{Gr}(k, n) : \mathcal{M}(x) = \mathcal{M}\}$.

A matroid is a *positroid* if it is the matroid of a totally nonnegative point $x \in \mathrm{Gr}(k, n)_{\geq 0}$. Its corresponding matroid variety $\Pi = \Pi_{\mathcal{M}}$ is called a *positroid variety*. The *open positroid variety* $\overset{\circ}{\Pi}_{\mathcal{M}}$ is defined to be the subset of Π not belonging to a lower-dimensional positroid variety; it is a Zariski open subset of Π. The intersection $\overset{\circ}{\Pi} \cap \mathrm{Gr}(k, n)_{\geq 0}$ is called a *positroid cell* and is homeomorphic to $\mathbb{R}_{>0}^{\dim \Pi}$.

Unlike general matroid varieties which can have arbitrarily bad singularities, positroid varieties are normal and Cohen-Macaulay. Likewise, while the problem of indexing matroid varieties (that is, recognizing representable matroids) is essentially hopeless, positroids are indexed by a well-behaved family of combinatorial structures. Here is one way of indexing positroid strata:

Theorem 2.4 (Postnikov, Knutson-Lam-Speyer [7])**.** *Let G be a planar bipartite graph in the disk. Then as N varies over all possible edge-weightings of G by nonzero complex numbers, the boundary measurements $X(N)$ sweep out a Zariski dense subset of a single positroid variety $\Pi = \Pi(G)$. Furthermore, every positroid variety Π arises in this way from some planar bipartite graph G. If we restrict attention to networks N with positive real edge weights, then the boundary measurements sweep out the entire positroid cell $\mathrm{Gr}(k, n)_{\geq 0} \cap \overset{\circ}{\Pi}$. In particular, if a given boundary measurement Δ_I is not identically zero on G, then $\Delta_I(N)$ is positive for each choice of network N with $\mathbb{R}_{>0}$ edge weights.*

We denote by $\mathcal{M}(G)$ the positroid such that $\Pi_{\mathcal{M}} = \Pi(G)$. We say that G (or N) *represents the top cell* if its positroid $\mathcal{M}(G)$ is the uniform matroid, or equivalently, if its positroid variety $\Pi(G)$ is equal to the Grassmannian $\mathrm{Gr}(k, n)$.

3. Tensor invariants and webs

Let U be an r-dimensional vector space. We denote by $\bigwedge^a(U)$ the ath exterior power of U. Throughout this paper we assume we have chosen a basis $E_1 \ldots, E_r$ for U, thus giving an isomorphism $\bigwedge^r(U) \cong \mathbb{C}$ under which the volume form $E_1 \wedge \cdots \wedge E_r \mapsto 1$.

Let $V_1, \ldots, V_n$ be a sequence of irreducible representations of $\mathrm{SL}(U)$. A *tensor invariant* is an element of the space

$$\mathrm{Hom}_{\mathrm{SL}(U)}\left(\bigotimes_{i=1}^{n} V_i, \mathbb{C}\right).$$

In this paper, we will be interested in tensor invariants of fundamental representations. Let $\lambda = (\lambda_1, \ldots, \lambda_n)$ be a sequence of integers satisfying

(3.1) $\qquad 0 \leq \lambda_i \leq r$ for $i = 1, \ldots, n$, and $\lambda_1 + \lambda_2 + \cdots + \lambda_n = kr$ for some k.

We are interested in the space

(3.2) $$\mathcal{W}_\lambda(U) = \mathrm{Hom}_{\mathrm{SL}(U)}\left(\bigotimes_{i=1}^{n} \bigwedge^{\lambda_i}(U), \mathbb{C}\right).$$

The dimension of the space (3.2) is the coefficient of the Schur polynomial $s_{((k^r))}(x_1, \ldots, x_r)$ in the product of r elementary symmetric polynomials $e_{a_1}(x_1, \ldots, x_r) \cdots e_{a_r}(x_1, \ldots, x_r)$. By

iterating the dual Pieri Rule, this is the number of $r \times k$ tableaux, of content λ, whose entries are strictly increasing along rows and weakly increasing down columns.

The simplest example, which the reader is encouraged to have in mind throughout, is when $n = kr$ and $\lambda = (1, \ldots, 1)$. In this case we prefer to denote $\mathcal{W}_{(1,\ldots,1)}(U)$ by $\mathcal{W}(r, n)$. The space $\mathcal{W}(r, n)$ is the vector space of SL_r-invariant multilinear functions on n vectors. Its dimension is the number of standard Young tableaux with r rows and k columns.

Webs are a particular way of encoding tensor invariants by using planar diagrams. We will define these tensor invariants in two steps, by introducing *untagged webs* followed by *tagged webs.*

Definition 3.1. An *untagged web* W is a planar graph in the disk, with each directed edge $e = (u, v)$ labeled by a multiplicity $m(e) = m(u, v) \in [r-1]$, so that reversing directions replaces a by $r - a$, that is $m(u, v) + m(v, u) = r$ for all directed edges (u, v). Furthermore, at each interior vertex, there exists a way to direct the edges so that the sum of the incoming multiplicities equals the sum of the outgoing multiplicities. The degree $\lambda(W)$ is the sequence of multiplicities $(m(b_1), \ldots, m(b_n))$ obtained by directing all boundary edges away from the boundary.

We will see that an untagged web W determines a tensor invariant in $\mathcal{W}_\lambda(U)$, up to a sign. The sign is determined by choosing a *tagging* of the web, which we now define.

Definition 3.2 (Tagged web). A *tagged* SL_r*-web* $\hat{W}$ is a planar graph in the disk subject to the following. First, each edge or half-edge of $\hat{W}$ is directed and labeled by a multiplicity in one of the three following ways

(3.3)   pair tag   source tag   .
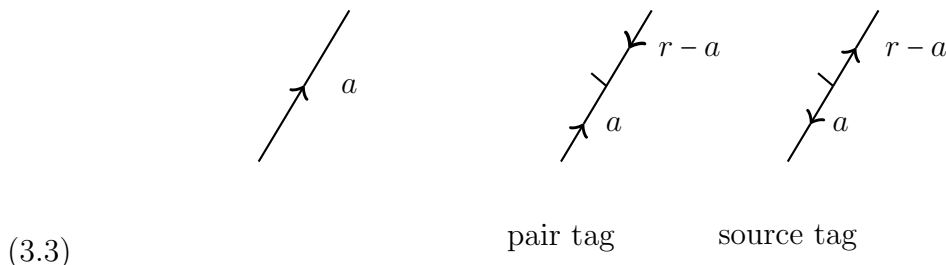
The "tiny edges" decorating the second and third edge types are called *tags* and they provide us with a preferred choice of side for the given edge. Second, we require that each interior vertex of $\hat{W}$ is modeled on one of the following two pictures:

(3.4)   wedge   shuffle   .
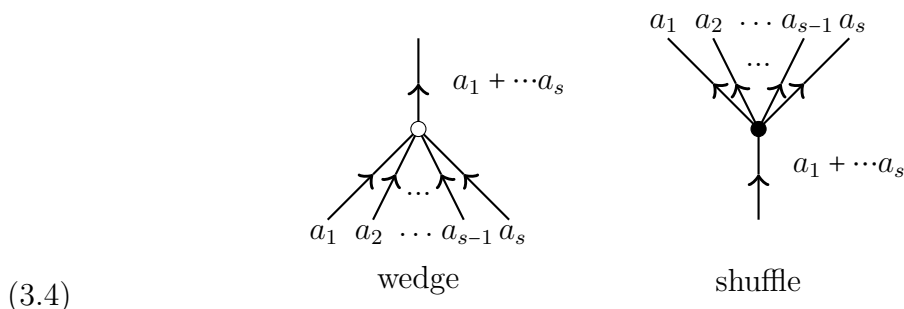
Thirdly and finally, we require that each boundary edge in $\hat{W}$ is a source. The *degree* $\lambda(\hat{W}) = (\lambda_1, \ldots, \lambda_n)$ of $\hat{W}$ is the sequence of edge multiplicities of the boundary edges.

We adopt the convention throughout that $\hat{W}$ stands for a tagged web as in Definition 3.2, and will reserve the symbol W for untagged webs, or for the weblike subgraphs that we discuss later. We will refer to the data of tagging by the letter $\mathcal{O}$, writing $\hat{W} = \hat{W}(W, \mathcal{O})$.

The definition of a tagged web could in principle allow for "wedge *and* shuffle" vertices, in which there are multiple edges flowing in both the inward and outward direction in (3.4). However, these will not come up in practice in our paper, so we do not work with this definition.

It will be important for us to note that by tagging sufficiently many edges in $\hat{W}$, one can always ensure that there are no oriented cycles in $\hat{W}$. We will always assume that $\hat{W}$ has been tagged in this way, because this simplifies the association of a tensor invariant to each web.

Denote by $\mathcal{FS}_\lambda(r)$ the space of finite formal $\mathbb{C}$-linear combinations of tagged SL_r-web diagrams of degree λ. Now we explain how an SL_r-web $\hat{W}$ of degree λ determines a tensor invariant $\hat{W} \in \mathcal{W}_\lambda(U)$. This induces a linear map $\mathcal{FS}_\lambda(r) \to \mathcal{W}_\lambda(U)$ that is known to be surjective.

An edge with multiplicity $m(e)$ in $\hat{W}$ encodes a copy of $\bigwedge^{m(e)}(U)$. Each of the interior vertices (3.4) encodes an $\mathrm{SL}(U)$-invariant map between the indicated tensor powers of fundamental representations. The first vertex encodes the exterior product map

$$(3.5) \qquad \overset{a_1}{\bigwedge}(U) \otimes \cdots \otimes \overset{a_s}{\bigwedge}(U) \to \overset{a_1 + \cdots + a_s}{\bigwedge}(U)$$

given by $x_1 \otimes x_2 \otimes \cdots \otimes x_s \mapsto x_1 \wedge x_2 \wedge \cdots \wedge x_s$. The second vertex encodes the map

$$(3.6) \qquad \overset{a_1 + \cdots + a_s}{\bigwedge}(U) \to \overset{a_s}{\bigwedge}(U) \otimes \cdots \otimes \overset{a_1}{\bigwedge}(U)$$

given by sending the wedge product $x_1 \wedge \cdots \wedge x_b \in \bigwedge^b(U)$ to the signed sum of shuffles

$$(3.7) \quad \sum \pm (x_{i_1} \wedge x_{i_2} \wedge \cdots \wedge x_{i_{a_s}}) \otimes \cdots \otimes (x_{i_{b-a_1+1}} \wedge x_{i_{b-a_1+2}} \wedge \cdots \wedge x_{i_b}) \in \overset{a_s}{\bigwedge}(U) \otimes \cdots \otimes \overset{a_1}{\bigwedge}(U)$$

where the summation is over permutations $(i_1, i_2, \ldots, i_b)$ of $(1, 2, \ldots, b)$ such that indices are increasing in each block: $i_1 < i_2 < \cdots < i_{a_s}$, $\ldots$, $i_{b-a_1+1} < \cdots < i_b$. The sign $\pm$ is the sign of the permutation $(i_1, i_2, \ldots, i_b)$, multiplied by the "global sign" of the permutation $(b - a_s + 1, \ldots, b, \ldots, 1, \ldots, a_1)$ Note that in both cases, the cyclic order of the edges in (3.4) is crucial for specifying signs.

The tagged edges in (3.3) should be thought of as degenerate cases of the maps (3.5) and (3.6), where the tag stands encodes a copy of $\bigwedge^r(U)$, which we canonically identify with $\mathbb{C}$ using the volume form. Thus the pair tag $\bigwedge^a \otimes \bigwedge^{r-a} \to \bigwedge^r \cong \mathbb{C}$ produces a number obtained by pairing the two incoming tensors, and source tag "creates" two tensors using the shuffle $\bigwedge^r \to \bigwedge^a \otimes \bigwedge^{r-a}$. Note that the side of the edge that the tag occurs on matters, because, for example, the maps $\bigwedge^a \otimes \bigwedge^{r-a} \to \bigwedge^r \cong \mathbb{C}$ and $\bigwedge^{r-a} \otimes \bigwedge^a \to \bigwedge^r \cong \mathbb{C}$ differ by a sign.

To evaluate $\hat{W}$ on a simple tensor $x_1 \otimes \cdots \otimes x_n \in \bigotimes_{i=1}^n \bigwedge^{\lambda_i}(U)$, we imagine placing each tensor x_i at boundary vertex i. We repeatedly compose the four basic morphisms – wedge, shuffle, pair, and source – as indicated by the arrows in $\hat{W}$ to obtain the value $\hat{W}(x_1 \otimes \cdots \otimes x_n)$.

Remark 3.3 (Taggings and perfect orientations). In practice, the following recipe will work to specify a tagging for the webs we consider in this paper. Suppose that we have an untagged web W whose underlying graph G is a bipartite graph. Orient all the edges in W from white to black, and further suppose that the weights of the directed edges around each vertex sum to r. These conditions characterize the r-weblike subgraphs that appear in this paper. In this situation, we can tag the web by choosing a *perfect orientation* of G.

A *perfect orientation* of a planar bipartite graph G in the disk is a choice of direction on each edge of G such that every white vertex has outdegree 1 and every interior black vertex has indegree 1. To get a perfect orientation of G, just choose an almost perfect matching of G: given an almost perfect matching π, we obtain a perfect orientation by directing each edge e from white to black if $e \in \pi$ (resp. black to white if $e \notin \pi$). The r-weblike subgraphs defined below always have almost perfect matchings (in fact, they are a union of r almost perfect matchings), so they can be perfectly oriented. Furthermore, if G has excedance k, then any perfect orientation for G will have exactly k boundary sinks (and exactly $n - k$ boundary sources). By adding a pair tag at each of the k boundary sinks, we get a tagged web. The only reason we don't use this construction in general is that there is no guarantee that the resulting web does not contain oriented cycles. To obtain such a web, we may need to add more tags along interior edges, introducing sources and pairings.

Example 3.4. Consider the SL_4-web $\hat{W}$ in (3.8). The edge mutiplicities are red, and multiplicities equal to one are omitted. The underlying bipartite graph for $\hat{W}$ has excedance 2. In the second figure, we show how to evaluate $\hat{W}$ on the tensor product of basis vectors $E = E_1 \otimes E_2 \otimes E_3 \otimes E_4 \otimes E_3 \otimes E_2 \otimes E_1 \otimes E_4 \in U^{\otimes 8}$. We abbreviate $E_{123} = E_1 \wedge E_2 \wedge E_3$ etc.

(3.8)
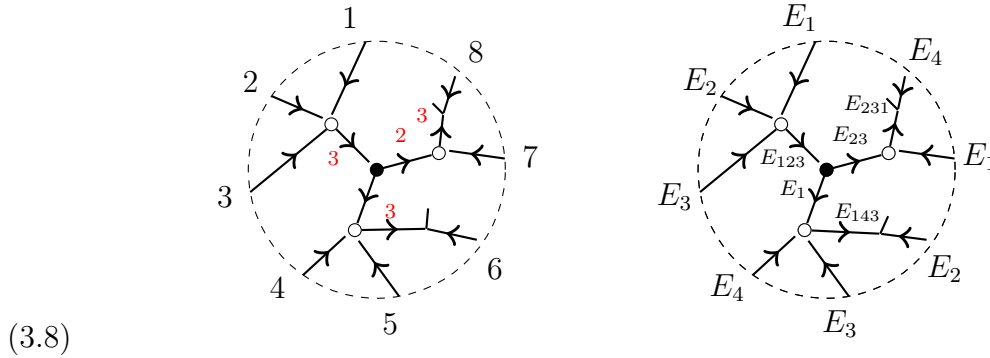
To begin, place the ith vector in E at boundary vertex i. Vectors flow along directed arrows until they reach a black vertex, which in this case happens when $E_1 \wedge E_2 \wedge E_3$ arrives at the interior black vertex. The map (3.6) splits this tensor up as a signed sum

$$(3.9) \qquad E_1 \otimes (E_2 \wedge E_3) - E_2 \otimes (E_1 \wedge E_3) + E_3 \otimes (E_1 \wedge E_2).$$

The signs come from (3.7). The evaluation is a sum of contributions from three terms, but only the first term pairs nontrivially with the vectors at 6 and 8. The pairings for this term are $(E_1 \wedge E_4 \wedge E_3) \wedge E_2 = -1$ and $(E_2 \wedge E_3 \wedge E_1) \wedge E_4 = 1$. The evaluation is $\hat{W}|_E = 1 \cdot -1 \cdot 1 = -1$, obtained as the product of the sign at the shuffle, at vertex 6, and at vertex 8.

Remark 3.5. Definition 3.2 is *not* the usual definition of a web [1; 9]. Typically, the definition comes with an added requirement that every interior vertex be *trivalent* (as a degenerate case of this, bivalent interior vertices are also allowed). Requiring trivalence is reasonable, because it can be shown that all possible $\mathrm{SL}(U)$-equivariant maps amongst tensor products of various $\bigwedge^a(V)$ come from compositions of maps involving three tensor factors [13]. Likewise, using (M2), any network N can be transformed to a network N' with interior vertices of valence at most three, without changing the boundary measurements. However, the trivalent restriction is not necessary for the current work.

Example 3.6 (Webs in small rank). Let us describe SL_r-webs $\hat{W}$ in the multilinear case $\lambda = (1, \ldots, 1)$, for $r = 1, 2, 3$. In each case, we have vectors $v_1, \ldots, v_n$ sitting at the boundary vertices $1, \ldots, n$. When $r = 1$, then $U \cong \mathbb{C}$, and $\hat{W}$ is a union of isolated interior edges (these do not affect the tensor invariant and can be removed) and edges based at boundary vertices. Thus $\hat{W}$ encodes a monomial in the vectors $v_1, \ldots, v_n$. When $r = 2$, an SL_2-web $\hat{W}$ consists of a disjoint union of a) tagged cycles, and b) directed paths $v_i \to v_j$ between boundary vertices. An oriented cycle contributes a multiplicative factor of ± 2 (depending on the tagging) when $\hat{W}$ is evaluated on $v_1, \ldots, v_n$. Changing the orientation on a path changes the web by a minus sign. Thus, ignoring these signs, and removing all cycles, SL_2-webs are spanned by crossingless matchings on the boundary vertices. In fact, these crossingless matchings are a basis for $\mathcal{W}(2, 2r)$.

A result in similar spirit is true when $r = 3$. In this case, the sign of a web does not depend on tagging. SL_3-webs are typically drawn as bipartite graphs without directed edges, with the convention that the edge multiplicities are given by $m(b, w) = 1$, where b is a black vertex and w is a white vertex. The SL_3 *skein relations* provide diagrammatic rules for expressing any SL_3-web in terms of a basis of *non-elliptic webs* [9], i.e. webs that are without 2-valent vertices or interior faces bounded by four or fewer sides (cf. [3] for a summary of these results).

For $r \geq 4$, the set of SL_r-webs is a distinguished spanning set, but the existence of a *web basis* satisfying enough "desirable" properties is unknown.

Finally, let us recall the relationsip between webs and the Grassmannian. Each Plücker coordinate $\Delta_I \in \mathbb{C}[\mathrm{Gr}(k, n)]$ is an SL_k-invariant function of the column vectors $v_1, \ldots, v_n \in \mathbb{C}^k$ representing a point in $\mathrm{Gr}(k, n)$.

In general, there is a $\mathbb{Z}^n$-grading of $\mathbb{C}[\mathrm{Gr}(k, n)]$ given by the degree in each column. For example, $\Delta_{123}\Delta_{256} \in \mathbb{C}[\mathrm{Gr}(3, 6)]$ has degree $(1, 2, 1, 0, 1, 1)$. The coordinate ring decomposes as

$$(3.10) \qquad \mathbb{C}[\mathrm{Gr}(k, n)] = \bigoplus_{r=1}^{\infty} \bigoplus_{\lambda_1 + \cdots + \lambda_n = rk} \mathbb{C}[\mathrm{Gr}(k, n)]_{\lambda},$$

where the inner direct sum is over λ as in (3.1). Note that $\mathbb{C}[\mathrm{Gr}(k, n)]_{\lambda}$ is given by the invariant space

$$\mathrm{Hom}_{\mathrm{SL}_k}\left(\bigotimes_{j=1}^{n} \mathrm{Sym}^{\lambda_j}(\mathbb{C}^k), \mathbb{C}\right).$$

Our main construction will give a duality between the graded piece $\mathbb{C}[\mathrm{Gr}(k, n)]_{\lambda}$ of the homogeneous coordinate ring of the Grassmannian and the SL_r-invariant space $\mathcal{W}_{\lambda}(U)$, where λ satisfies (3.1). We will discuss in Section 8 that these spaces are naturally dual

using purely representation-theoretic considerations. However, our work manifests this duality explicitly in terms of webs and dimers.

If $n = rk$, the invariant space $\mathcal{W}(k, n)$ sits inside $\mathbb{C}[\mathrm{Gr}(k, n)]$ as the multilinear piece $\lambda = (1, \ldots, 1)$. It is spanned by the r-fold products of Plücker coordinates $\Delta_{I_1} \cdots \Delta_{I_r}$ satisfying $I_1 \cup \cdots \cup I_r = [n]$. Thus our results will provide a duality between $\mathcal{W}(k, n)$ and $\mathcal{W}(r, n)$. We examine this duality pictorially in the Appendix.

4. The main construction

In this section we make the connection between dimer configurations and webs. Suppose we are given r dimer configurations $\pi_1, \ldots, \pi_r$ on N. By superimposing them we naturally obtain what we call an r-weblike subgraph:

Definition 4.1. An *r-weblike subgraph* $W \subset G$ (alternatively, $W \subset N$) is a subgraph of G, using all the vertices of G, with each edge e of W labeled by a multiplicity $m(e) \in [r]$, in such a way that the sum of the multiplicities around each interior vertex is r. The *weight* of W is the product $\mathrm{wt}(W) = \prod_e \mathrm{wt}(e)^{m(e)}$ of the edge weights raised to the indicated multiplicity. The *degree* of W is its sequence $\lambda(W) = (m(b_1), \ldots, m(b_n))$ of boundary multiplicities, where $b_1, \ldots, b_n$ are the boundary edges.

Each r-weblike subgraph W determines an untagged SL_r-web by giving each black-white edge (b, w) the multiplicity of this edge in W. Thus, one can get a tensor invariant $\hat{W}(W, \mathcal{O})$ from W by choosing a tagging $\mathcal{O}$. For any two such choices of tagging $\mathcal{O}$ and $\mathcal{O}'$, the resulting tensor invariants are equal up to a sign: $\hat{W}(W, \mathcal{O}) = \pm \hat{W}(W, \mathcal{O}')$. In Section 5, we define a sign $\mathrm{sign}(W, \mathcal{O}) \in \{\pm 1\}$ so that the tensor invariant

$$(4.1) \qquad \mathbf{W} := \mathrm{sign}(W, \mathcal{O})\hat{W}(W, \mathcal{O}).$$

does not depend on $\mathcal{O}$. In other words, there is a "correct" choice of sign for the tensor invariant represented by W. With this in mind, we can formulate the main definition of this paper.

Definition 4.2. Let N be a network of excedance k and λ satisfy (3.1). We denote by

$$(4.2) \qquad \mathrm{Web}_r(N; \lambda) := \sum_{W \subset G, \ \lambda(W) = \lambda} \mathrm{wt}(W) \, \mathbf{W} \in \mathcal{W}_\lambda(U),$$

the weighted sum over the r-weblike subgraphs of G with degree λ. It is a $\mathbb{C}$-linear combination of the various boldface $\mathbf{W} \in \mathcal{W}_\lambda(U)$, with the $\mathrm{wt}(W)$'s serving as coefficients.

We think of $\mathrm{Web}_r(N; \lambda)$ as an *r-fold boundary measurement.* When $r = 1$, the choice of λ is equivalent to the choice of $I \in \binom{[n]}{k}$, and $\mathrm{Web}_1(N; I)$ is a variant of the boundary measurement $\Delta_I(N)$ (2.2). Whereas $\Delta_I(N)$ is a number, our $\mathrm{Web}_1(N; I)$ lies in $\mathcal{W}_\lambda(U)$, which in this case is isomorphic to $\mathbb{C}$.

We denote by $\mathrm{Web}_r(N) := \sum_\lambda \mathrm{Web}_r(N; \lambda) \in \bigoplus_\lambda \mathcal{W}_\lambda(U)$.

Example 4.3. Consider the pair of networks N and N' given by

(4.3)

where a', b', c', d' are related to a, b, c, d according to the spider move (2.3)

$$(4.4) \qquad a' = \frac{a}{ac + bd} \qquad b' = \frac{b}{ac + bd} \qquad c' = \frac{c}{ac + bd} \qquad d' = \frac{d}{ac + bd}.$$

Their parameters are $k = 2$ and $n = 6$, and the letters $a, a', b, b', \ldots, e, f \in \mathbb{C}^\times$ denote edge weights. Let us consider $r = 3$ and $\lambda = (1, 1, 1, 1, 1, 1)$. The network N has three r-weblike subgraphs with degree λ. The 3-fold boundary measurement $\mathrm{Web}_3(N; \lambda)$ is a linear combination

(4.5)

with coefficients depending on $a, \ldots, f$. On the other hand, $\mathrm{Web}_3(N'; \lambda)$ is a linear combination of two webs

(4.6)

The following theorem says that the r-fold boundary measurements factors through Postnikov's boundary measurement map.

Theorem 4.4. *If N and N' are two networks satisfying $\tilde{X}(N) = \tilde{X}(N')$, then $\mathrm{Web}_r(N; \lambda) = \mathrm{Web}_r(N'; \lambda) \in \mathcal{W}_\lambda(U)$.*

That is, even though the right hand sides of (4.2) are different elements of $\mathcal{FS}_\lambda(r)$, they are equal as tensor invariants in $\mathcal{W}_\lambda(U)$. We prove Theorem 4.4 in Section 5.

Example 4.5. Continuing with Example 4.3, recall that a', b', c', d' are related to a, b, c, d by (2.3). The networks N and N' are related by (M1), and $\tilde{X}(N) = (ac + bd)\tilde{X}(N')$. It follows that $\mathrm{Web}_3(N; \lambda) = (ac + bd)^3 \mathrm{Web}_3(N'; \lambda)$. By equating the coefficient of $abcdef$

in (4.5) with the coefficient of $\frac{abcdef}{(ac+bd)^3}$ in (4.6), we deduce the *square move* for SL_3-webs:

(4.7)

4.1. **Immanants.** Now consider any functional $\varphi \in \mathcal{W}_\lambda(U)^*$ on the SL_r tensor invariant space, where λ satisfies (3.1). From Theorem 4.4, the number

$$(4.8) \qquad\qquad \varphi(\mathrm{Web}_r(N;\lambda))$$

is independent of the choice of network N representing the point $\tilde{X}(N) \in \widetilde{\mathrm{Gr}}(k,n)$.

Now let G be a bipartite graph that represents the top cell. As we vary over networks N with underlying graph G, the map $\tilde{X}(N) \mapsto \varphi(N)$ defines a function on the subset of $\widetilde{\mathrm{Gr}}(k,n)$ swept out by $\tilde{X}(N)$'s.

Proposition 4.6. *The function* $\tilde{X}(N) \mapsto \varphi(\mathrm{Web}_r(N;\lambda))$ *extends to a (uniquely defined) element of* $\mathbb{C}[\mathrm{Gr}(k,n)]_\lambda$*; this function does not depend on the choice of* G*.*

We will prove Proposition 4.6 in Section 5. It leads naturally to the following definition:

Definition 4.7 (Immanant map). The linear map $\mathrm{Imm} : \mathcal{W}_\lambda(U)^* \to \mathbb{C}[\mathrm{Gr}(k,n)]_\lambda$ defined by $\mathrm{Imm}(\varphi)(\tilde{X}(N)) = \varphi(\mathrm{Web}_r(N;\lambda))$ is called the *immanant map.*

Another of our main theorems is the following.

Theorem 4.8. *The immanant map* $\mathrm{Imm} : \mathcal{W}_\lambda(U)^* \to \mathbb{C}[\mathrm{Gr}(k,n)]_\lambda$ *is an isomorphism.*

Thus, we obtain a canonical pairing of $\mathcal{W}_\lambda(U)$ with $\mathbb{C}[\mathrm{Gr}(k,n)]_\lambda$, described more explicitly in (5.16) and (5.17). In the multilinear case $\lambda = (1,\ldots,1)$, this pairing is the unique (up to scalars) S_n-equivariant pairing of the web spaces $\mathcal{W}(r,n)$ and $\mathcal{W}(k,n)$, cf. Section 8.

Remark 4.9. When r is 2 or 3, the second author previously defined SL_2- and SL_3-*web immanants* [10]. These are obtained via (4.8) when the functional $\varphi \in \mathcal{W}(r,n)^*$ is the dual functional to an element of the web basis. In the case that $r \geq 4$, since we are without a notion of a web basis, we believe the r-fold boundary measurement $\mathrm{Web}_r(N)$ is the more fundamental object.

5. PROOF OF THE MAIN THEOREMS

Definition 5.1 (Consistent labeling). Let W be an r-weblike subgraph of a planar bipartite graph G with n boundary vertices, and with boundary edges $b_1,\ldots,b_n$. A *consistent labeling of* W is a labeling of each edge e in W by a subset $S(e) \subset [r]$ so that $|S(e)| = m(e)$, and such that the union of the sets around each interior vertex is $[r]$. Equivalently, one can require that the sets $S(e)$ are disjoint around each vertex.

Let ℓ be a consistent labeling. From the following pair of equalities of multisets

$$\bigcup_{\text{all edges}} S(e) = \bigcup_{\text{white vertices}} [r]$$

$$\bigcup_{\text{all non-boundary edges } e} S(e) = \bigcup_{\text{interior black vertices}} [r],$$

it follows that the multiset $S(b_1) \cup \cdots \cup S(b_n)$ satisfies

$$(5.1) \qquad\qquad S(b_1) \cup \cdots \cup S(b_n) = \{1^k, 2^k, \ldots, r^k\}.$$

We refer to $\mathcal{S} = (S(b_1), \ldots, S(b_n))$ as a list of *boundary label subsets*.

We can also associate to any labeling ℓ a list of *boundary location subsets* $\mathcal{I} = (I_1, \ldots, I_r) \subset [n]$, defined by $I_i = \{j \in [n] : i \in S(b_j)\}$. From the preceding argument, we know that each $|I_i|$ has size k. Furthermore,

$$(5.2) \qquad\qquad I_1 \cup \cdots \cup I_r = \{1^{\lambda_1}, 2^{\lambda_2}, \ldots, n^{\lambda_n}\}$$

as multisets, where λ is the degree of W.

The data of boundary labels $\mathcal{S} = (S_1, \ldots, S_n) \subset [r]$ satisfying (5.1), and the data of boundary locations $\mathcal{I} = (I_1, \ldots, I_r) \subset [n]$, satisfying (5.2), are dual to each other (in the sense of combinatorial design theory), i.e. one can be recovered from the other.

Recall we have a vector space U with basis $E_1, \ldots, E_r$ satisfying $E_1 \wedge E_2 \wedge \cdots \wedge E_r = 1$. To any subset $S \subset [r]$ with $|S| = a$ we can associate the tensor $E_S \in \bigwedge^a(U)$ by taking the wedge of the basis vectors labeled by S, in ascending order. If $\mathcal{S} = (S_1, \ldots, S_n)$ is a list of boundary label subsets, we obtain in this way a tensor $E_{\mathcal{S}} \in \bigotimes_{j=1}^n \bigwedge^{\lambda_j}(U)$ by

$$(5.3) \qquad\qquad E_{\mathcal{S}} = E_{S_1} \otimes E_{S_2} \otimes \cdots \otimes \cdots \otimes E_{S_n}.$$

It will be convenient to replace the multiset $\{1^k, 2^k, \ldots, r^k\}$ by an honest set $\mathcal{A} = \{1_1, 1_2, \ldots, 1_k, 2_1, 2_2, \ldots, 2_k, \ldots, r_1, \ldots, r_k\}$ which we refer to as the *alphabet*. We introduce the obvious total order $1_1 < 1_2 < \cdots < 1_k < 2_1 < \cdots < r_k$ on $\mathcal{A}$. We let $\mathbb{C}\langle \mathcal{A} \rangle$ be the free vector space with basis $\mathcal{A}$, and with volume form $1_1 \wedge 1_2 \wedge \cdots \wedge r_k = 1 \in \bigwedge^{rk}(\mathbb{C}\langle \mathcal{A} \rangle)$. We will think of the numbers $1, \ldots, r$ as *colors*: for $i_s, j_t \in \mathcal{A}$ we say that i_s and j_t have the same *color* if $i = j$.

If we read the indices of the basis vectors in (5.3) from left to right, we get a word $w'(\mathcal{S})$ in which each of the numbers $1, \ldots, r$ appears exactly k times. We let $\text{sign}(\mathcal{S}) = 1$ or -1 depending on whether the number of inversions of $w'(\mathcal{S})$ is even or odd. Equivalently, we can think of the entries of $w'(\mathcal{S})$ as elements of $\mathcal{A}$ by making the subscripts increase by one from left to right. We let $w(\mathcal{S}) \in \mathcal{A}$ denote the word obtained from $w'(\mathcal{S})$ in this way. Then $\text{sign}(\mathcal{S})$ is the value of the wedge product of $w(\mathcal{S}) \in \bigwedge^{rk}(\mathbb{C}\langle \mathcal{A} \rangle)$.

Example 5.2. Let G be the bipartite graph underlying Example 3.4. Here is a particular 4-weblike subgraph $W \subset G$, and a consistent labeling ℓ of W:

(5.4)
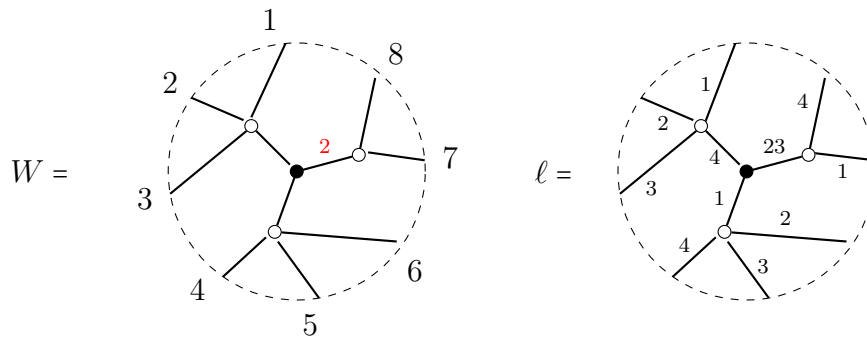
The labeling ℓ of W has boundary location subsets $\mathcal{I} = (\{1,7\}, \{2,6\}, \{3,5\}, \{4,8\})$ and boundary label subsets $\mathcal{S} = (\{1\}, \{2\}, \{3\}, \{4\}, \{3\}, \{2\}, \{1\}, \{4\})$. The tensor $E_{\mathcal{S}}$ (5.3) is the 8-fold tensor E from Example 3.4. The sign $\text{sign}(\mathcal{S})$ is given by $1_1 \wedge 2_1 \wedge 3_1 \wedge 4_1 \wedge 3_2 \wedge 2_2 \wedge 1_2 \wedge 4_2 = -1$. The SL_4-web $\hat{W}$ from Example 3.4 is a tagging of W, obtained by directing the edges of G according to (3.8) (cf. Remark 3.3). The labeling ℓ of W determines the flow of tensors along the edges of $\hat{W}$ pictured in (3.8), specifically we replace S in ℓ by E_S for edges directed black to white and by $E_{[4] \setminus S}$ for edges directed white to black.

As Example 5.2 suggests, consistent labelings of W with boundary $\mathcal{S}$ are closely related to evaluating $\hat{W}$ on the tensor $E_{\mathcal{S}}$.

Definition 5.3. Denote by $a(\mathcal{S}; W)$ the number of consistent labelings of W with fixed boundary label subsets $\mathcal{S}$.

The following seemingly innocuous lemma is the key technical result underpinning our main theorems. The proof of the lemma is somewhat intricate, and has minimal bearing on the rest of the paper, so it may be skipped on a first reading.

Lemma 5.4. *Let W be an r-weblike subgraph of a planar bipartite graph G, and $\hat{W} = \hat{W}(W, \mathcal{O})$ a choice of tagging for W. Then there is a sign $\text{sign}(W, \mathcal{O}) \in \{\pm 1\}$, such that for any list of boundary label subsets $\mathcal{S}$, we have*

$$(5.5) \qquad \hat{W}\big|_{E_{\mathcal{S}}} = \text{sign}(\mathcal{S}) \text{sign}(W, \mathcal{O}) \, a(\mathcal{S}; W).$$

It is fairly easy to see (and we elaborate on this below) that consistent labelings with boundary $\mathcal{S}$ give rise to terms in the evaluation of $\hat{W}$ on $E_{\mathcal{S}}$. Therefore, the key assertion in this lemma is that each of these terms contributes to the evaluation with the same sign. Let us also emphasize that the sign $\text{sign}(W, \mathcal{O})$ is independent of the boundary input $E_{\mathcal{S}}$.

Thus for any r-weblike subgraph W, the tensor invariant $\mathbf{W} = \text{sign}(W, \mathcal{O})\hat{W}(W, \mathcal{O}) \in \mathcal{W}_\lambda(U)$ of (4.1) is characterized by the equation

$$(5.6) \qquad \mathbf{W}\big|_{E_{\mathcal{S}}} = \text{sign}(\mathcal{S}) \, a(\mathcal{S}; W).$$

for every list of boundary label subsets $\mathcal{S}$.

Proof. Let us consider a consistent labeling ℓ with boundary label subsets $\mathcal{S}$. The choice of tagging on $\hat{W}$ prescribes how to evaluate $\hat{W}$ on $E_{\mathcal{S}}$. Our first goal is to explain why the consistent labeling ℓ gives a term in the evaluation of $\hat{W}$ on $E_{\mathcal{S}}$. This evaluation involves many terms because of the vertices where we perform the shuffle operation.

For a given labeling ℓ, we can place tensors on the edges of $\hat{W}$ as governed by ℓ. If an edge (or half-edge if the edge is tagged) is labeled by the set S in ℓ, let us place the tensor E_S along that edge if it is directed black to white and place the complementary tensor $E_{[r]\setminus S}$ along the edge if it is directed white to black. A tensor along a given edge in a given term represents some partial stage of the evaluation $\hat{W}$ on $E_{\mathcal{S}}$. Because ℓ is consistent, the incoming flow of basis vectors equals the outgoing flow at each interior vertex, and no two basis vectors flow through the same interior vertex. Thus, once we have placed the appropriate tensor on each edge of $\hat{W}$ as indicated by ℓ, we get a sign for this term which contributes to the final evaluation.

Thus, each consistent labeling corresponds to a term in $\hat{W}\big|_{E_{\mathcal{S}}}$, and the nonzero terms in the evaluation of $\hat{W}$ on $E_{\mathcal{S}}$ are counted by $a(\mathcal{S}; W)$. The nontrivial assertion is that once $\mathcal{S}$ is fixed, every term contributes with the same sign. The content of the proof lies in understanding these signs. The signs come from two places: the shuffles at interior vertices and the pairings at each tag.

We begin our sign analysis by associating to a consistent labeling ℓ a flow of vectors along the directed edges of $\hat{W}$. Specifically, if $i \in S \subset [r]$ and E_S is at a given edge (or half-edge) of this term, then we say that the basis vector E_i flows along that edge. The evaluation begins with tensors $E_{S_1}, \ldots, E_{S_n}$ at the boundary vertices. Each of these E_{S_i} is itself a wedge product of basis vectors. By keeping track of how such a boundary basis vector flows along the web, we get a path from the boundary to a tag.

In this way, the labeling ℓ provides us with rk paths starting at boundary sources and ending at tags. These paths are naturally indexed by the multiset $\{1^k, \ldots, r^k\}$, and we index them by elements of the alphabet $\mathcal{A}$ by adding subscripts in counterclockwise order along the boundary.

For the sake of simplicity, we will assume that there are no interior sources in $\hat{W}$, so that the term is completely described by the union of these rk paths (if there are interior sources, the argument below can be repeated with slight modifications).

Let us take these rk paths from the boundary to the tags and isotope them slightly so that they are disjoint at the boundary, and are noncrossing except in transverse crossings in small neighborhoods of interior shuffle vertices. Thus when a of these paths flow along an edge of multiplicity a, the a paths have a natural order from left to right. For each tag there will be exactly r paths ending at the tag.

Definition. A *flow* on a tagged web $\hat{W}$ is a collection of rk paths starting at the boundary and ending at tags, where the number of paths flowing along each edge is the multiplicity of the edge, and where the paths are labeled by the alphabet $\mathcal{A}$. We will usually isotope these paths so that we get a parallel paths along each edge of multiplicity a and so that the crossings of these paths are concentrated in a neighborhood of shuffle vertices.

To summarize, a consistent labeling ℓ provides us with a flow, which we will also call ℓ. The consistent labeling from Example 5.2, which corresponds to the term in the evaluation

computed in Example 3.4, has the following flow:

(5.7)
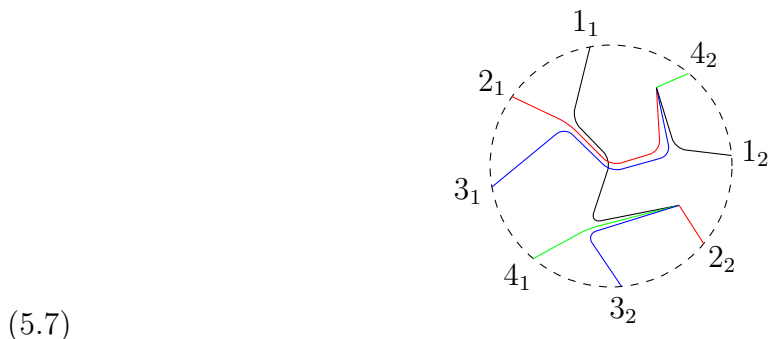
We will now consider flows which do not necessarily come from consistent labelings. In a general flow, it is not required that the paths along each edge are of different colors (likewise for the paths flowing into a tag). This requirement holds precisely for flows coming from consistent labelings. As an intermediate step, we will associate a sign $\epsilon(\ell) = \epsilon(\ell, \mathcal{O})$ to *any* flow ℓ on $\hat{W}$, not necessarily coming from a consistent labeling. The sign $\epsilon(\ell)$ is a topological invariant of the rk paths from the boundary to the tags. It knows nothing about structure of the set $\mathcal{A}$ other than that it consists of rk elements (and thus it does not know about the colors of these rk elements).

Let $t_1, \ldots, t_k$ be a list of the tags in $\hat{W}$. For each tag t_i, we let $C_i \in \binom{\mathcal{A}}{r}$ be the set of labels of the r paths ending at t_i. We refer to C_i as a *tag subset*. Within each tag subset, the terms are ordered (from left to right) according to how they flow into the tag, where we recall that if x and y are complementary tensors flowing into a tag and x is left of y (when the tag points north) then the tensors are ordered $x \wedge y$. The ordered set C_i naturally gives rise to an element of $\wedge^r \mathbb{C}\langle \mathcal{A} \rangle$, which by abuse of notation we also call C_i. Then the sign $\epsilon(\ell)$ is defined by

$$\epsilon(\ell) \coloneqq C_1 \wedge C_2 \wedge \cdots \wedge C_k \in \bigwedge^{rk} \mathbb{C}\langle \mathcal{A} \rangle.$$

For any list of boundary subsets $\mathcal{S}$, there is a *canonical* way of flowing according to $\mathcal{O}$, namely the one in which strands cross maximally at each interior shuffle vertex. According to (3.7), there is no sign at interior shuffle vertices associated with such a flow. We let $\ell_{\mathrm{canon}}(\mathcal{S})$ denote the canonical flow with boundary $\mathcal{S}$. This canonical flow is a device that gives us a reference point for analyzing signs.

Let $\mathcal{S}_0$ be list of boundary label subsets whose word is lexicographically smallest, i.e. $w(\mathcal{S}_0) = 1_1 1_2 1_3 \ldots r_k$. For any other choice of boundary label subsets $\mathcal{S}$, the canonical flows have signs related by

(5.8) $$\epsilon(\ell_{\mathrm{canon}}(\mathcal{S}_0)) = \epsilon(\ell_{\mathrm{canon}}(\mathcal{S})) \operatorname{sign}(\mathcal{S}).$$

Now suppose we have a consistent labeling ℓ with boundary subset $\mathcal{S}$. We can get from $\ell_{\mathrm{canon}}(\mathcal{S})$ to ℓ by performing a sequence of swaps at shuffle vertices. Clearly, each such swap changes the sign ϵ by a factor of -1. We have that
(5.9)
$$\epsilon(\ell) = \epsilon(\ell_{\mathrm{canon}}(\mathcal{S})) (-1)^{\# \text{ of shuffle swaps } \ell_{\mathrm{canon}}(\mathcal{S}) \to \ell} = \epsilon(\ell_{\mathrm{canon}}(\mathcal{S})) (\text{sign of shuffle swaps}).$$

Now we compute $\epsilon(\ell)$ by computing the number of transpositions required to transform $\epsilon(\ell)$ into $1_1 \wedge 1_2 \wedge \cdots \wedge r_k$. First, within each tag subset, we can rearrange the terms so that they are increasing in $\mathcal{A}$. We will refer to this as the *inversions within tags*. In addition to alphabetizing within each tag, we must perform a transposition for each inversion involving two elements of different tag subsets. We will refer to this latter number of inversions as the *inversions between tags*. Thus, we have that

(5.10)
$$\epsilon(\ell) = (-1)^{\text{inversions within tags}} (-1)^{\text{inversions between tags}} = (\text{sign within tags}) \cdot (\text{sign between tags}).$$

We need one final observation: the contribution of ℓ to $\hat{W}\big|_{E_S}$ has sign equal to (sign of shuffle swaps) times (sign within tags). Putting this together with (5.8) through (5.10), the sign of the contribution of ℓ is thus

(5.11)
$$(\text{sign within tags}) (\text{sign of shuffle swaps}) = \epsilon(\ell_{\text{canon}}(\mathcal{S}_0)) (\text{sign between tags}) \operatorname{sign}(\mathcal{S}).$$

Thus the proof is finished once we show the following: the sign between tags is the same for *any* consistent labeling ℓ of $\hat{W}$, regardless of the choice of boundary label subsets $\mathcal{S}$. Recall that for a consistent labeling ℓ the rk paths are labeled by the alphabet $\mathcal{A}$ by adding subscripts in counterclockwise order.

This last claim follows from two subclaims: first, since ℓ is a consistent labeling, all of the tag subsets are copies of $[r]$. Thus, the number of inversions between tags involving *elements of a different color* does not depend on ℓ. For example, if 4 is in the tag subset C_3, then it will be in inversion with the copies of $1, 2, 3$ in the tag subsets $C_4, C_5, \ldots$. Second, we need to check that the total number of inversions *between elements of the same color* (and in different tags) does not depend on the consistent labeling ℓ. Thus we have reduced the proof of the lemma to the second subclaim, which we now state formally:

Claim. Fix a tagged web $\hat{W}$. Consider any consistent labeling ℓ of $\hat{W}$. Suppose that ℓ has boundary label subset $\mathcal{S}$. Adding subscripts in counterclockwise order, we get a flow on $\hat{W}$ with paths labeled by $\mathcal{A}$. Then the signs between tags for elements of $\mathcal{A}$ of the same color is independent of ℓ (and hence $\mathcal{S}$).

Up to this point, the sign analysis has been mostly "soft" reasoning. Proving this second subclaim requires using the planarity of $\hat{W}$ in an essential way.

Let us cut the boundary of the disk between boundary vertices n and 1, flattening the web $\hat{W}$ so that the boundary vertices are drawn on the x-axis in $\mathbb{R}^2$ in the order $1, \ldots, n$, and the web $\hat{W}$ is drawn in the upper half plane. The flow ℓ gives rk directed paths from the x-axis to the tags, colored by the elements of $[r]$. Let us make the assumption that the x-coordinates of the boundary vertices, the tags t_i, and the intersection points between the rk paths are pairwise distinct. Let us assume furthermore that the tags are ordered $t_1, \ldots, t_k$ by the x-coordinates, from left to right.

Fixing the boundary data $\mathcal{S}$, the flow $\ell_{\text{canon}}(\mathcal{S})$ is the one in which these paths cross each other maximally, and any consistent labeling ℓ with boundary $\mathcal{S}$ is obtained by resolving certain of the crossings in $\ell_{\text{canon}}(\mathcal{S})$. Furthermore, a key observation that we use later is that if the labeling is consistent, then two paths of the same color never intersect.

For each pair (p, t_j) of a path p ending at t_i, and a tag t_j (where $i \neq j$), we define the *nesting number* $n(p, t_j)$ of the pair (p, t_j) to be the total number of intersection points of

p with the vertical ray R_j going upwards from t_j to ∞. We define $n(p, t_i) = 0$ if p ends at t_i. For any labeling ℓ, we can define the *total nesting number* of ℓ as the sum of the nesting numbers for all pairs (p, t_j).

We now argue that when we resolve a crossing in a flow, the total nesting number does not change. Let p, q be two paths ending at t_i and t_j respectively. Since we have assumed that $\hat{W}$ has no oriented cycles, there are no oriented cycles in the flow. It follows that

(5.12) p (resp. q) has no self intersections, and

(5.13)   the intersection points of $p \cap q$ appear in the same order on p and q.

When we resolve a crossing of p and q, the quantity $n(p, t_k) + n(q, t_k)$ for $k \neq i, j$ clearly remains unchanged. Let us decompose the path p (resp. q) into two segments $p_b p_a$ (resp. $q_b q_a$) so that the crossing happens when p_a ends and p_b begins (likewise for q_a and q_b). We will be done if we show that $n(p, t_j) + n(q, t_i) \equiv n(q_b p_a, t_i) + n(p_b q_a, t_j) \mod 2$. We can decompose $n(p, t_j)$ as $\#p_a \cap R_j + \#p_b \cap R_j$ and so on for the other paths. Certain of these contributions cancel modulo 2, and our claim reduces to checking that

(5.14) $$\#(p_a \cup q_a) \cap (R_i \cup R_j) \equiv 0 \mod 2.$$

Indeed, each of $\#(p_a \cup q_a) \cap (R_i)$ and $\#(p_a \cup q_a) \cap (R_i)$ is even, which follows from the fact that t_j (resp. t_i) cannot lie inside a bounded region enclosed by p_a, q_a and the x-axis by (5.12) and (5.13).

Thus the total nesting number of any consistent labeling ℓ with boundary $\mathcal{S}$ is the same as the total nesting number of $\ell_{\mathrm{canon}}(\mathcal{S})$. The total nesting number of $\ell_{\mathrm{canon}}(\mathcal{S})$ is independent of $\mathcal{S}$ because the flow of the paths for ℓ_{canon} is defined topologically (without regard to color).

Moreover, because ℓ is a consistent labeling, and therefore two paths of the same color never cross, the total nesting number coincides with the number of inversions between elements of the same color modulo 2. This completes the proof of the claim and the lemma. $\square$

The following result relates tensor evaluation of webs to boundary measurements of networks; it serves as the fundamental link between webs and dimers.

Proposition 5.5. *Let N be a network with excedance k. Let $\mathcal{I} = (I_1, \ldots, I_r)$ be a list of subsets satisfying (5.2), dual to $\mathcal{S} = (S_1, \ldots, S_n)$ satisfying (5.1). Then*

(5.15) $$\Delta_{I_1}(N) \cdots \Delta_{I_r}(N) = \mathrm{sign}(\mathcal{S}) \mathrm{Web}_r(N; \lambda)\big|_{E_{\mathcal{S}}}.$$

We still find the statement of Proposition 5.5 surprising, yet the importance of this statement for us is belied by the simplicity of the proof.

Proof. We have that

$$\Delta_{I_1}(N) \cdots \Delta_{I_r}(N) = \sum_{r\text{-weblike subgraphs } W, \text{ degree } = \lambda} a(\mathcal{S}; W) \mathrm{wt}(W)$$
$$= \mathrm{sign}(\mathcal{S}) \sum_W \mathbf{W} \mathrm{wt}(W)$$
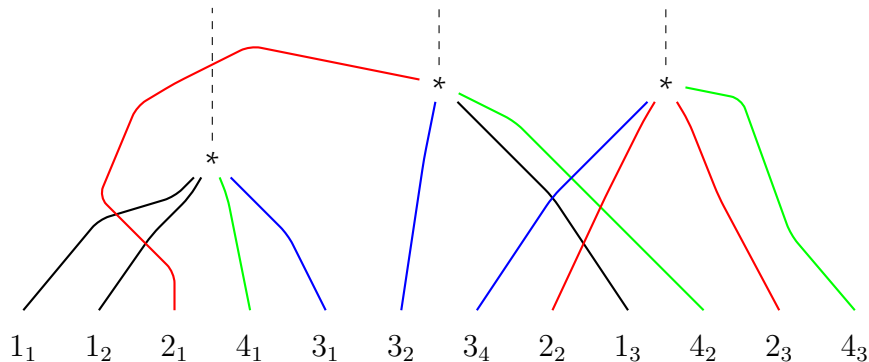$$= \mathrm{sign}(\mathcal{S}) \mathrm{Web}_r(N; \lambda)\big|_{E_{\mathcal{S}}}.$$

FIGURE 1. An example of a flow that could come from an SL_4-web on 12 vertices. The $*$'s are tags and the dashed lines are the rays. The total number of nestings for this flow is 1, witnessed by the path through 2_1 going above the first tag. The flow is not consistent because two black strands end at tag 1. It can be made consistent (for example) by resolving the crossing between strands 1_1 and 2_1 and also between 2_2 and 1_3, which does not change the total nesting number. Once these two resolutions are made, the total number of inversions between colors is 1 (witnessed by the inversion between 1_2 and 1_1), which matches the initial total number of nestings.

The first equality follows by considering superpositions of dimer configurations $(\pi_1, \ldots, \pi_r)$ with boundary subsets $(I_1, \ldots, I_r)$ and grouping them according to the r-weblike subgraph they determine. The second equality follows from Lemma 5.4, and the third equality from the definition (4.2). $\qquad\square$

Proof of Theorem 4.4. From (5.15), the evaluation of $\mathrm{Web}_r(N; \lambda)$ on a basis tensor $E_{\mathcal{S}}$ can be expressed in terms of boundary measurements, which only depend on $\tilde{X}(N)$. Evaluation at $E_{\mathcal{S}}$ for varying $\mathcal{S}$ determines an element of $\mathcal{W}_\lambda(U)$. This proves Theorem 4.4. $\qquad\square$

Proof of Proposition 4.6 and Theorem 4.8. First note that for each list of boundary label subsets $\mathcal{S}$, there is a functional $\mathrm{eval}(E_{\mathcal{S}}) \in \mathcal{W}_\lambda(U)^*$ given by evaluation at $E_{\mathcal{S}}$. Furthermore, by (5.15), the function $\tilde{X}(N) \to \mathrm{eval}(E_{\mathcal{S}})(\mathrm{Web}_r(N; \lambda))$ agrees with the function $\tilde{X}(N) \mapsto \mathrm{sign}(\mathcal{S})\Delta_{I_1}(N)\cdots\Delta_{I_r}(N)$. This latter function is a regular function on $\widetilde{\mathrm{Gr}}(k, n)$. Since the set of points of the form $\tilde{X}(N)$ is Zariski dense in $\widetilde{\mathrm{Gr}}(k, n)$, the function $\tilde{X}(N) \to \mathrm{eval}(E_{\mathcal{S}})(\mathrm{Web}_r(N; \lambda))$ extends uniquely to a regular function on all of $\widetilde{\mathrm{Gr}}(k, n)$. The resulting extension is the image of the immanant map. Now we note the functionals $\mathrm{eval}(E_{\mathcal{S}})$ span $\mathcal{W}_\lambda(U)^*$, so we conclude that every element of $\varphi \in \mathcal{W}_\lambda(U)^*$ gets sent to a a regular function on $\widetilde{\mathrm{Gr}}(k, n)$ by the immanant map. This proves Proposition 4.6.

Since r-fold products $\Delta_{I_1}\cdots\Delta_{I_r}$ of Plücker coordinates span the graded piece of $\mathbb{C}[\mathrm{Gr}(k, n)]_\lambda$, the immanant map is surjective. A dimension count verifies that the map is an isomorphism. We will see another argument for the equality of dimensions in Section 8. This proves Theorem 4.8. $\qquad\square$

Finally, let us extract some consequences. First, the pairing $\mathcal{W}_\lambda(U) \otimes \mathbb{C}[\mathrm{Gr}(k,n)]_\lambda \to \mathbb{C}$ can be given in simple terms as follows:

$$(5.16) \qquad \langle \mathbf{W}, \Delta_{I_1} \cdots \Delta_{I_r} \rangle = \mathrm{sign}(\mathcal{S}) \mathbf{W}\big|_{E_\mathcal{S}} = a(\mathcal{S}; W)$$

$$(5.17) \qquad \langle \mathrm{Web}_r(N; \lambda), f \rangle = f(\tilde{X}(N)).$$

In (5.16) we let $(I_1, \ldots, I_r)$ be a list of subsets, dual to $\mathcal{S}$, and let W be an r-weblike subgraph with tensor invariant $\mathbf{W} \in \mathcal{W}_\lambda(U)$. In (5.17) we let $f \in \mathbb{C}[\mathrm{Gr}(k,n)]_\lambda$ and N be any network of excedance k. Note in particular that (5.16) is always nonnegative. The equation (5.17) follows from checking the special case where $f = \Delta_{I_1} \cdots \Delta_{I_r}$, which was the content of Proposition 5.5.

Second, we can give another interpretation of the duality between $\mathbb{C}[\mathrm{Gr}(k,n)]_\lambda$ and $\mathcal{W}_\lambda(U)$. Let G be a bipartite graph representing the top cell. Let us pick a collection of edges of G so that varying the weights w_e of those edges (keeping the rest of the weights as 1) gives a a rational parametrization $(\mathbb{C}^\times)^{k(n-k)} \to \mathrm{Gr}(k,n)$. We now think of the weights w_e as variables, and form the sum

$$(5.18) \qquad \mathrm{Web}_r(G; \lambda) := \sum_{\substack{W \subset G \,\lambda(W) = \lambda}} \mathrm{wt}(W)\, \mathbf{W}.$$

Here, $\mathrm{wt}(W)$ are monomials in the edge weight variables w_e. Specializing the edge-weights to complex numbers gives an element of $\mathcal{W}_\lambda(U)$. Thus we get elements of $\mathcal{W}_\lambda(U)$ depending algebraically on points of $\widetilde{\mathrm{Gr}}(k,n)$. Our above analysis shows that we can then view $\mathrm{Web}_r(G; \lambda)$ as giving rise to an element

$$\mathrm{Web}_r(G; \lambda) \in \mathbb{C}[\mathrm{Gr}(k,n)]_\lambda \otimes \mathcal{W}_\lambda(U)$$

which *is* the pairing between these two spaces. We have that

$$\mathrm{Web}_r(G; \lambda)\big|_{E_\mathcal{S}} = \mathrm{sign}(\mathcal{S}) \Delta_{I_1} \cdots \Delta_{I_r}$$

for any boundary label subsets $\mathcal{S}$, and in fact, this characterizes $\mathrm{Web}_r(G; \lambda) \in \mathbb{C}[\mathrm{Gr}(k,n)]_\lambda \otimes \mathcal{W}_\lambda(U)$. In particular, we emphasize that $\mathrm{Web}_r(G; \lambda)$ is independent of G (as long as G represents the top cell).

Since $\mathrm{Web}_r(G; \lambda)$ realizes the duality between $\mathbb{C}[\mathrm{Gr}(k,n)]_\lambda$ and $\mathcal{W}_\lambda(U)$, it must be a full-rank tensor in $\mathbb{C}[\mathrm{Gr}(k,n)]_\lambda \otimes \mathcal{W}_\lambda(U)$. From this, it can be seen that as N varies, the various $\mathrm{Web}_r(N; \lambda)$ span $\mathcal{W}_\lambda(U)$. This will also follow from Theorem 7.4.

Let us also formulate the following positivity conjecture.

Conjecture 5.6. *Under the isomorphism $\mathcal{W}_\lambda(U) \cong \mathbb{C}[\mathrm{Gr}(r,n)]_\lambda$, the functions $\mathbf{W}$ take nonnegative values on the cone $\widetilde{\mathrm{Gr}}(r,n)_{\geq 0}$ over the totally nonnegative Grassmannian.*

In Section 7, we show that webs in $\mathcal{W}(r,n)$ are naturally dual to certain elements of $\mathbb{C}[\mathrm{Gr}(k,n)]$, and that this duality behaves well under restriction to positroid subvarieties $\Pi \subset \mathrm{Gr}(k,n)$. Our positivity conjecture 5.6 says that SL_r-webs are positive when thought of as functions on the Grassmannian $\mathrm{Gr}(r,n)$. To briefly explain the reasoning, we first note that the positivity property of Conjecture 5.6 is known to hold for the canonical basis by the work of Lusztig [12]. In the $r = 2$ or $r = 3$ cases, $\mathbf{W}$ expands positively as a sum of non-elliptic web basis elements. These web basis elements are known to coincide with the canonical basis for $r = 2$, and share many similar properties with the canonical basis

in the case $r = 3$. Finally, we view the sign coherence in Lemma 5.4 as a manifestation of positivity.

Example 5.7. As a simple example of Conjecture 5.6, consider the *single cycle web* $W \in \mathcal{W}(3, 9)$ (the second web in the first row of Figure 3), with its boundary vertices numbered so that vertices $1, 4, 7$ connect to the interior hexagon of W. Multiplying by the Plücker coordinate Δ_{479} and using the skein relations, one sees that $\Delta_{479} \cdot W$ is a sum of two cluster monomials. It follows that W can be expressed as a Laurent monomial with positive coefficients in Plücker coordinates for $\mathbb{C}[\mathrm{Gr}(3, 9)]$, and therefore takes positive values on $\widetilde{\mathrm{Gr}}(3, 9)_{\geq 0}$. See [2, Proposition 16]), for a more general version of this argument (that does not explicitly mention positivity) in the case $r = 3$.

6. Deducing skein relations from moves on networks

According to Theorem 4.4, the element $\mathrm{Web}_r(N)$ only depends on $\tilde{X}(N)$. Thus it is unchanged when we perform the local moves from Theorem 2.3 to N. However, the *expression* for $\mathrm{Web}_r(N)$ as a weighted sum of r-weblike subgraphs will change after each local move, since the weblike subgraphs change. In this way, we obtain relations between certain linear combinations of SL_r-webs.

As it turns out, the diagrammatic relations on webs we obtain in this way are exactly the Cautis-Kamnitzer-Morrison relations [1]. This will be the content of Theorem 6.1] and Corollary 6.2. We can interpret this fact in two ways: On the one hand, it is possible (though perhaps tedious) to use the Cautis-Kamnitzer-Morrison relations to give an alternative proof Theorem 4.4 – one checks that if N and N' are related by a local move on networks, then $\mathrm{Web}_r(N)$ and $\mathrm{Web}_r(N')$ differ by a skein relation. We believe that our approach to Theorem 4.4, based on Proposition 5.5 and Lemma 5.4, is somewhat more illuminating. On the other hand, in this paper we will use Theorem 4.4 to derive diagrammatic relations amongst webs, and use our duality setup to prove that the relations we obtain in this way are a complete set of diagrammatic relations.

Our r-weblike graphs are closely related to, but slightly different from the webs considered in [1]. These differences require some care when comparing moves on bipartite graphs and diagrammatic relations on webs. Thus we will begin with diagrammatic relations on r-weblike subgraphs before proceeding to the case of webs.

6.1. r-weblike graph relations.

We begin by illustrating how the various skein relations on r-weblike graphs can be derived from relations on networks.

Consider the space $\mathcal{FW}_\lambda(r)$ of finite formal $\mathbb{C}$-linear combinations of r-weblike graphs (i.e, labeled bipartite graphs such that the sum of labels around every interior vertex equals r). We will later compare this with the space $\mathcal{FS}_\lambda(r)$ of formal $\mathbb{C}$-linear combinations of SL_r-web diagrams.

There is a surjection $\mathcal{FW}_\lambda(r) \twoheadrightarrow \mathcal{W}_\lambda(U)$ from the evaluation equation (5.6), that is, by sending a r-weblike graph W to the tensor invariant given by its boldface version $\mathbf{W}$.

We begin by listing all the diagrammatic moves that can be performed to an r-weblike graph W without changing the value of the tensor invariant $\mathbf{W}$. Of these moves, the last move is the only truly interesting one.

- Two-valent vertex removal (6.1). There is also an analogous move to (6.1) which has the colors reversed. For simplicity, we drew all neighboring edges in (6.1) with multiplicity 1, though in general they may have any multiplicity.
- Bigon removal (6.1). Parallel edges, with multiplicities a and b, can be replaced by one edge with multiplicity $a + b$, at the cost of a multplicative factor $\binom{a+b}{b}$. As before, colors can be reversed.
- Leaf and dipole removal. Edges with multiplicity r or 0 can be removed from r-weblike graphs.
- The square move for r-weblike graphs (6.2). Notice that the diagrams in (6.2) connect to the "outside" with multiplicity j from the southwest, ℓ from the southeast, $r+v-s-\ell$ from the northeast, and $r+s-v-j$ from the northwest. If j, ℓ, v are fixed, there is a square move for each s satisfying $\max(0, v-\ell, v+j-r) \le s \le \min(j, r-\ell)$.

(6.1)

(6.2)
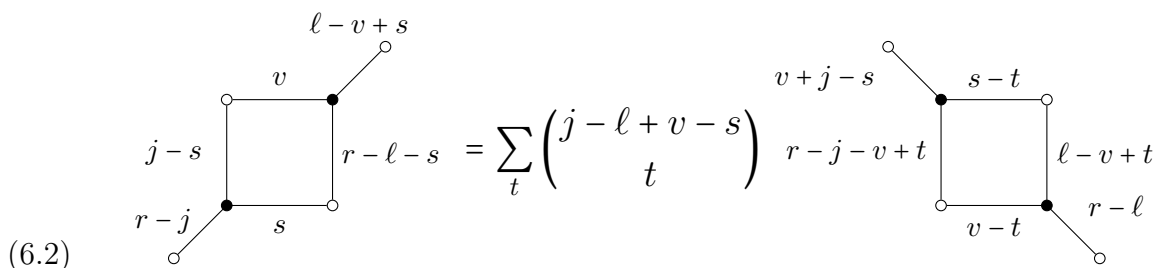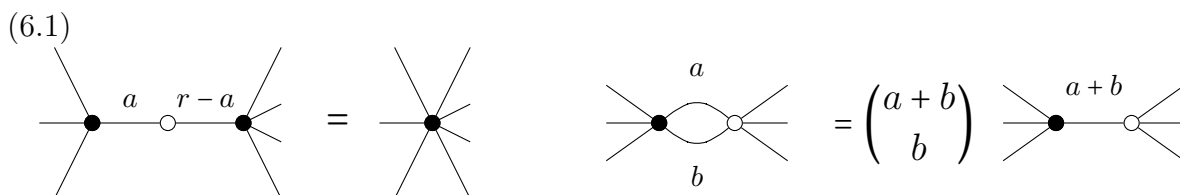
Theorem 6.1. *The relations amongst networks imposed by Theorem 4.4 generate the kernel $\mathcal{F}\mathcal{W}_\lambda(r) \twoheadrightarrow \mathcal{W}_\lambda(U)$.*
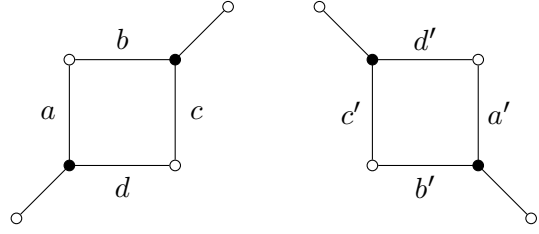
Corollary 6.2. *The four diagrammatic relations listed above generate the kernel $\mathcal{F}\mathcal{W}_\lambda(r) \twoheadrightarrow \mathcal{W}_\lambda(U)$.*

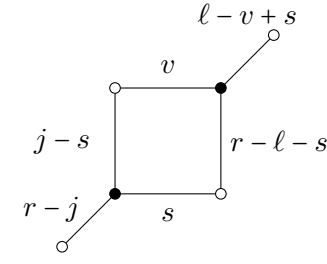Before proving Theorem 6.1, we will show how Corollary 6.2 follows from it.

Proof of Corollary 6.2. Let us briefly explain how the relations (6.1) and (6.2) can be deduced from network moves. The two-valent vertex removal (M2) for networks implies two-valent vertex removal (6.1) for r-weblike graphs. The parallel edge removal (R1) for networks implies the bigon removal move (6.1) for r-weblike graphs. Leaf and dipole removal (R2), (R3) for networks give leaf and dipole removal for r-weblike graphs.

Thus we now focus on the most interesting move on networks (M1), and show that it implies the square move for r-weblike graphs (6.2).

Suppose N and N' are are related by a square move involving edge weights $a, b, c, d \in \mathbb{C}^*$
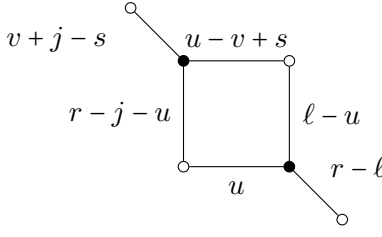

(6.3)

We need to compare the r-weblike subgraphs coming from the networks N and N'. Each such weblike subgraph can be built by fixing the edge multiplicites $m(e)$ for edges e outside the local fragment (j from the southwest, ℓ from the southeast, and so on, as described above), then choosing the edge multiplicities inside the local fragment in a way that is compatible with the outside. Thus the numbers j, ℓ and $s - v$ are fixed throughout this computation. We will let s be variable (so that s determines v).

(6.4)

The contribution to $\mathrm{wt}(W)$ from the edges in the local fragment is $a^{j-s}b^v c^{r-\ell-s}d^s$.

On the other hand, in N', for the same value of outside multiplicities, the filling inside will look like

(6.5)

where the values of $j, \ell, v - s$ agree with the fixed values, but u can vary. Notice that s and v only occur in (6.5) together via their difference $s - v$.

Now we denote by W_s the weblike subgraph indexed by s in (6.4) and denote by W'_u the one indexed by u in (6.5). We sum up the boldface $\mathbf{W}_s$ with their weights and the $\mathbf{W}_u$ with their weights. Since the points in the affine cone are related by $\tilde{X}(N) = (ac + bd)\tilde{X}(N')$, it follows that $\mathrm{Web}_r(N) = (ac + bd)^r \mathrm{Web}_r(N')$. Writing out what this means:

$$\sum_s a^{j-s}b^v c^{r-\ell-s}d^s \, \mathbf{W}_s = (ac + bd)^r \sum_u (a')^{\ell-u}(b')^u(c')^{r-j-u}(d')^{u-v+s} \, \mathbf{W}'_u$$

$$= (ac + bd)^{j-\ell+v-s} \sum_u a^{\ell-u}b^u c^{r-j-u}d^{u-v+s} \, \mathbf{W}'_u.$$

Now we fix a value of s, and we also fix the exponent t of bd in the binomial formula expansion of $(ac + bd)^{v+j-\ell-s}$. In order to match the weights on the left and right hand

side, the value of u must be $v - t$. The corresponding equality is

$$(6.6) \qquad \mathbf{W}_s = \sum_t \binom{j - \ell + v - s}{t} \mathbf{W}'_{v-t} \in \mathcal{W}_\lambda(U),$$

which is the square move (6.2).

Thus we see that the local moves on networks imply the four diagrammatic moves for r-weblike graphs that preserve the corresponding tensor invariants. $\qquad\square$

Proof of Theorem 6.1. Now, we will give an "abstract proof" that diagrammatic moves amongst r-weblike subgraphs we obtain in this way are a complete set of relations.

Let us denote by K the kernel $\mathcal{FW}_\lambda(r) \to \mathcal{W}_\lambda(U)$ and abbreviate $\mathcal{FW} = \mathcal{FW}_\lambda(r)$. Let $\mathcal{N}(k)$ denote the set of networks of excedance k whose underlying graph G represents the top cell. We will denote by $\mathbb{C}[\mathcal{N}(k)]$ the linear space of functions $f: \mathcal{N}(k) \to \mathbb{C}$ with the property that for any fixed graph G, the map f is a polynomial in the edge weights on G. The space $\mathbb{C}[\mathcal{N}(k)]$ is a very large space of functions.

We say that $f \in \mathbb{C}[\mathcal{N}(k)]$ is *consistent of order r* if $f(N) = \alpha^r f(N')$ whenever $\tilde{X}(N) = \alpha \tilde{X}(N')$. The space of consistent functions of order r is canonically identified with the space

$$\mathbb{C}[\mathrm{Gr}(k,n)]_{(r)} := \bigoplus_\lambda \mathbb{C}[\mathrm{Gr}(k,n)]_{(\lambda)}$$

where the sum is over λ such that $\lambda_1 + \cdots + \lambda_n = kr$.

There is a map $\mathrm{Imm}^\dagger : \mathcal{FW}^* \to \mathbb{C}[\mathcal{N}]$ sending a functional $\varphi \in \mathcal{FW}^*$ to the function on networks $N \mapsto \varphi(\mathrm{Web}_r(N; \lambda))$, where we now think of $\mathrm{Web}_r(N; \lambda)$ as an element of $\mathcal{FW}$. It is defined similarly to the immanant map Imm but on the space of formal sums of weblike graphs. This map obviously fits into a commutative diagram

$$(6.7) \qquad \begin{array}{c} (\mathcal{FW}/K)^* \\ \Big\uparrow \\ \mathcal{FW}^* \xrightarrow[\mathrm{Imm}^\dagger]{} \mathbb{C}[\mathcal{N}(k)] \longleftarrow \mathbb{C}[\mathrm{Gr}(k,n)]_{(\lambda)} \end{array} \qquad .$$

The diagonal arrow in this diagram is the immanant map $\mathrm{Imm} : \mathcal{W}_\lambda(U)^* \to \mathbb{C}[\mathrm{Gr}(k,n)]_\lambda$. Going diagonally and then left, we arrive at the space of consistent functions of order r on networks inside $\mathbb{C}[\mathcal{N}(k)]$ which have degree λ.

The image of the downward arrow is $K^\perp$.

The main claim which needs to be checked is that $\mathrm{Imm}^\dagger$ is an injection. Let us first see that the fact that $\mathrm{Imm}^\dagger$ is an injection allows us to finish the proof. Suppose that $\mathrm{Imm}^\dagger$ is an injection. Then it follows that $\varphi \in \mathcal{FW}^*$ defines a consistent function on networks if and only if $\varphi \in K^\perp$. To put it another way, whether or not φ lies in $K^\perp$ is detected by whether $\mathrm{Imm}^\dagger(\varphi)$ is a consistent function on networks. A function in $\mathbb{C}[\mathcal{N}(k)]$ is consistent if an only if it transforms appropriately under local moves. Therefore, the relations defining $K^\perp$ are exactly those forced on φ by the local moves on networks (cf. Proposition 2.2). In other words, we need that

$$\mathrm{Imm}^\dagger(\varphi)(N) = \alpha^r \, \mathrm{Imm}^\dagger(\varphi)(N') \qquad \text{whenever} \qquad \tilde{X}(N) = \alpha \tilde{X}(N').$$

We now prove the injectivity of $\mathrm{Imm}^\dagger$. Namely, if a functional $\varphi \in \mathcal{FW}^*$ satisfies $\varphi(\mathrm{Web}_r(N; \lambda)) = 0$ for all networks N of excedance k, then φ must be 0. Equivalently,

we must prove that the various $\{\text{Web}_r(N; \lambda)\}$ span $\mathcal{FW}$ as we vary N over all networks representing the top cell.

Now let G be a fixed bipartite graph that represents the top cell. The graph G has a finite number of r-weblike subgraphs. Let $W_1, \ldots, W_\ell$ be the list of all of the r-weblike subgraphs of G having degree λ. We claim that there exists a list $N_1, \ldots, N_\ell$ of networks with underlying graph G such that the matrix

$$(6.8) \qquad\qquad (\text{wt}_{N_i}(W_j))_{i,j=1,\ldots,\ell}$$

is invertible. First suppose that such $N_1, \ldots, N_\ell$ exist. The rows of the matrix give the terms of $\text{Web}_r(N_i; \lambda)$, so for each j there exists a linear combination of the $\{\text{Web}_r(N_i; \lambda)\}$ giving the r-weblike subgraph W_j. By varying G, we then conclude that any r-weblike subgraph W is in the span of $\{\text{Web}_r(N; \lambda)\}$, and $\{\text{Web}_r(N; \lambda)\}$ spans $\mathcal{FW}$.

The existence of $N_1, \ldots, N_\ell$ follows from a standard argument. The expression for $\text{Web}_r(N; \lambda)$ consists of a monomial in the edge weights of G multiplied by each weblike subgraph W_i. Because the W_i all differ, the monomial attached to each web has a different multidegree. Then a standard Vandermonde-type argument shows that one can specialize the edge-weights appropriately so that the matrix $(\text{wt}_{N_i}(W_j))$ is invertible. $\qquad\square$

6.2. **Cautis-Kamnitzer-Morrison relations.** Recall we denote by $\mathcal{FS}_\lambda(r)$ the space of formal sums of tagged SL_r-web diagrams of degree λ. Any tagged web gives a tensor invariant, so that we get a surjection $\mathcal{FS}(r) \twoheadrightarrow \mathcal{W}_\lambda(r)$. Cautis, Kamnitzer, and Morrison [1] gave a set of diagrammatic relations describing the kernel of this surjection. Let us briefly write down these relations (see [1, Section 2] for pictures and discussion):

- Switching a tag [1, (2.3)]. If e is a tagged edge with multiplicities a and $r - a$, then changing which side the tag is on contributes a factor $(-1)^{a(r-a)}$.
- Tag migration [1, (2.7),(2.8)]. Suppose an interior white vertex v has an incident edge e with a tag pointing in the clockwise direction around v (the tag can be either a pair or a source tag, and the vertex v can be either a shuffle or wedge vertex). Let e' be the next edge incident to v in the clockwise direction from v. Then the tag on e can be migrated to a tag on e' pointing in the counterclockwise direction, without any change of sign.
- Wedge product is associative [1, (2.6)] and shuffle is coassociative.
- Bigon removal [1, (2.4)]. A pair of directed edges $v \to u$, of multiplicities a and b, can be replaced by a single edge $v \to u$ with multiplicity $a + b$, at the cost of a factor $(-1)^{ab}\binom{a+b}{b}$.
- The square move for tagged webs:

$$(6.9)$$

Theorem 6.3. *[Cautis-Kamnitzer-Morrison] The five relations listed above generate the kernel $\mathcal{FS}(\text{SL}_r) \twoheadrightarrow \mathcal{W}_\lambda(U)$.*

We note that the results in [1] are more general than ours – they describe all diagrammatic relations amongst fundamental representations of the quantum group $U_q(\mathfrak{sl}_n)$, whereas our current proof only makes sense in the "classical" ($q = 1$) setting. Additionally, they provide certain redundant relations that are needed to describe the kernel if one wishes to work over $\mathbb{Z}[q, q^{-1}]$ instead of $\mathbb{C}(q)$. In our version, we work over $\mathbb{C}$.

Now we explain how Theorem 6.3 can be deduced from our Corollary 6.2.

There are three differences between r-weblike graphs and webs. First, and most important, r-weblike graphs do not come with a tagging. Recall that tagging an r-weblike graph W gives a web $\hat{W} = \hat{W}(W, \mathcal{O})$, where $\mathcal{O}$ refers to the data of the tagging. From any r-weblike graph W, we associate a canonical invariant $\mathbf{W} := \mathrm{sign}(W, \mathcal{O})\hat{W}$. One of the advantages of working with our tensor invariants $\mathbf{W}$ is that the signs come naturally built-in.

The second difference between r-weblike graphs and webs is purely a matter of convention and is less serious: Cautis-Kamnitzer-Morrison require that vertices in their webs are of degree three, whereas ours can be arbitrary. For this reason, they need additional associativity relations that we do not. These relations follow from two-valent vertex removal on networks.

The third difference is also less serious. We require our graphs to be bipartite, so that in order to go from a web to an r-weblike graph, we will sometimes need to contract edges. The fact that edge contractions give r-weblike graphs that are related by allowable moves again follows from two-valent vertex removal.

We now proceed with the proof. The first two relations in Theorem 6.3 have to do with tagging. These tagging relations generate a subspace of relations $\mathcal{R} \subset \mathcal{FS}_\lambda(r)$.

There is a linear map $\mathcal{FS}_\lambda(r) \twoheadrightarrow \mathcal{FW}_\lambda(r)$ defined by replacing a tagged web $\hat{W} = \hat{W}(W, \mathcal{O})$ by its "underlying r-weblike subgraph" with a sign: $\hat{W} \mapsto \mathrm{sign}(W, \mathcal{O})W \in \mathcal{FW}_\lambda(r)$.

Let us be more specific about how to obtain an r-weblike graph from $\hat{W}$. Any edges in $\hat{W}$ joining vertices of the same color should be contracted, so that we end up with a bipartite graph. If necessary, we may need to use tag switches and migrations before contracting. Once this is done, if e is an edge (or half-edge) with multiplicity a, and if e points from a white vertex to a black vertex, then e has multiplicity $r - a$ in W, while if it points away from a black vertex towards a white vertex, it has multiplicity a. Finally, forget about all the directions on the edges. The resulting labeled diagram is an r-weblike graph.

The quotient $\mathcal{FS}_\lambda(r)/\mathcal{R}$ is exactly the image of $\mathcal{FS}_\lambda(r)$ in $\mathcal{FW}_\lambda(r)$. We have that the map $\mathcal{FS}_\lambda(r) \twoheadrightarrow \mathcal{W}_\lambda(r)$ factors as

$$\mathcal{FS}_\lambda(r) \twoheadrightarrow \mathcal{FW}_\lambda(r) \twoheadrightarrow \mathcal{W}_\lambda(r)$$

There are relations on $\mathcal{FS}_\lambda(r)$ that come from the kernel of the first map. These are precisely the relations $\mathcal{R} \subset \mathcal{FS}_\lambda(r)$ coming from tagging. These relations follow from our analysis of tags in Lemma 5.4.

The remaining relations on $\mathcal{FS}_\lambda(r)$ come from "tagging" (i.e., adding tags to) relations in $\mathcal{FW}_\lambda(r)$. Thus, to find a complete set of diagrammatic relations amongst tagged webs, we need to lift the relations defining $\ker(\mathcal{FW}_\lambda(r) \twoheadrightarrow \mathcal{W}_\lambda(r))$ to $\mathcal{FS}_\lambda(r)$.

Let us do this in the most interesting case of the square move. The square move for r-weblike graphs (6.2) can be lifted to a relation in $\mathcal{FS}_\lambda(r)$ by tagging as follows
(6.10)

,

where we have chosen a tagging outside the local fragment that is compatible with the tagging inside the fragment shown above. Then (6.9) follows once one checks that all of the signs, $\text{sign}(W, \mathcal{O})$ and $\text{sign}(W'_t, \mathcal{O}'_t)$, are equal and do not depend on t. This can be proved by analyzing the canonical flows through the local fragments, as described in the proof of Lemma 5.4.

We briefly mention how to lift the remaining relations. Two-valent vertex removal for r-weblike graphs allows us to contract and expand edges to derive the associativity of wedge for tagged webs. In a similar fashion, bigon removal for r-weblike graphs lifts to the bigon removal for tagged webs.

7. Positroids and webs

In this section we fix a positroid $\mathcal{M}$ and consider a graph G such that $\mathcal{M}(G) = \mathcal{M}$ (see Section 2.3). Let $\Pi_\mathcal{M} = \Pi_G$ be the corresponding positroid variety.

7.1. **The immanant map for a positroid.** Let $\mathfrak{I}(\mathcal{M}) \subset \mathbb{C}[\text{Gr}(k, n)]$ be the (homogeneous) ideal generated by $\{\Delta_I : I \notin M\}$. By [7], $\mathfrak{I}(\mathcal{M})$ is a prime ideal, and its vanishing set in $\text{Gr}(k, n)$ is the positroid variety $\Pi_\mathcal{M}$. The homogeneous coordinate ring $\Pi_\mathcal{M}$ is thus given as the quotient

$$(7.1) \qquad\qquad \mathbb{C}[\Pi_\mathcal{M}] = \mathbb{C}[\text{Gr}(k, n)]/\mathfrak{I}(\mathcal{M}),$$

noting that by [7], $\Pi_\mathcal{M} \subseteq \text{Gr}(k, n)$ is projectively normal. The grading (3.10) descends to a grading on $\mathbb{C}[\Pi_\mathcal{M}]$. One of our goals will be to give a description of the graded pieces $\mathbb{C}[\Pi_\mathcal{M}]_\lambda$ via tensor invariants.

Recall that the immanant map (for the top cell) is an isomorphism

$$\text{Imm} : \mathcal{W}_\lambda(U)^* \to \mathbb{C}[\text{Gr}(k, n)]_\lambda.$$

We now define an immanant map $\text{Imm}_\mathcal{M} : \mathcal{W}_\lambda(U)^* \to \mathbb{C}[\Pi_\mathcal{M}]_\lambda$ for the positroid $\mathcal{M}$ by the equation

$$\text{Imm}_\mathcal{M}(\varphi)(\tilde{X}(N)) = \varphi(\text{Web}_r(N; \lambda))$$

for N any assignment of edge weights to G.

Proposition 7.1. *The function $\tilde{X}(N) \mapsto \varphi(\text{Web}_r(N; \lambda))$ extends to a (uniquely defined) element of $\mathbb{C}[\Pi_\mathcal{M}]_\lambda$.*

Proof of Proposition 7.1. The proof is the same as the proof of Proposition 4.6 after noting that the set of points of the form $\tilde{X}(N)$ is Zariski dense in $\widetilde{\Pi}_\mathcal{M}$. We get that $\text{Imm}_\mathcal{M}(\varphi)$ is a regular function on $\widetilde{\Pi}_\mathcal{M}$, and that thus the immanant map $\text{Imm}_\mathcal{M}$ is well-defined. $\square$

Let us note that the immanant map is characterized as the map that (up to sign) sends $\mathrm{eval}(E_\mathcal{S})$ to the function $\Delta_{I_1}\cdots\Delta_{I_r}$ where $\mathcal{S}$ and $\mathcal{I} = (I_1,\ldots I_r)$ are dual. Therefore $\mathrm{Imm}_\mathcal{M} = \pi_\mathcal{M} \circ \mathrm{Imm}$, where $\pi_\mathcal{M} : \mathbb{C}[\mathrm{Gr}(k,n)]_\lambda \twoheadrightarrow \mathbb{C}[\Pi_\mathcal{M}]_\lambda$ denotes the quotient map.

7.2. Partial evaluation of webs.

Let $\mathfrak{I}(\mathcal{M})_\lambda$ denote the degree λ part of the positroid ideal $\mathfrak{I}(\mathcal{M})$. We now give a description of the subspace

$$(\mathfrak{I}(\mathcal{M})_\lambda)^\perp := \{x \in W_\lambda(U) : \langle x, f \rangle = 0 \text{ for all } f \in \mathfrak{I}(\mathcal{M})_\lambda\} \subseteq W_\lambda(U)$$

where $\langle \cdot, \cdot \rangle$ denotes the pairing between $W_\lambda(U)$ and $\mathbb{C}[\mathrm{Gr}(k,n)]_\lambda$. Thus the inclusion $(\mathfrak{I}(\mathcal{M})_\lambda)^\perp \hookrightarrow W_\lambda(U)$ is dual to the surjection $\pi_\mathcal{M} : \mathbb{C}[\mathrm{Gr}(k,n)]_\lambda \twoheadrightarrow \mathbb{C}[\Pi_\mathcal{M}]_\lambda$.

We now define the notion of the partial evaluation of a web. Let $\hat{W} \in W_\lambda(U)$ be an SL_r-web, where $\lambda_1 + \cdots + \lambda_n = kr$. The copies of $\wedge^{\lambda_1}(U),\ldots,\wedge^{\lambda_n}(U)$ gives us kr locations to plug in vectors $v_1,\ldots,v_{kr}$. For example, the first λ_1 vectors are plugged into boundary vertex 1 as the wedge $v_1 \wedge \cdots \wedge v_{\lambda_1}$, and so on.

Definition 7.2. Let $I = \{i_1,\ldots,i_k\} \in \binom{[n]}{k}$. Let $U' \subset U$ be the subspace spanned by $E_1,\ldots,E_{r-1}$. Then the *partial evaluation map along I*

$$(7.2) \qquad\qquad\qquad \mathcal{W}_\lambda(U) \to \mathcal{W}_\mu(U')$$

$$(7.3) \qquad\qquad\qquad x \mapsto x\big|_{I \mapsto E_r}$$

is the linear map obtained by specializing the input vectors $v_{i_1}, v_{i_2},\ldots,v_{i_k}$ to the last basis vector E_r, and then restricting the remaining vectors to lie in the subspace U'. The resulting function is SL_{r-1}-invariant. We have $\mu_i = \lambda_i - 1$ or $\mu_i = \lambda_i$ depending on whether $i \in I$ or $i \notin I$. It satisfies $\mu_1 + \cdots + \mu_n = (r-1)k$. In the multilinear case $\lambda = (1,\ldots,1)$, partial evaluation along I is a linear map $\mathcal{W}(r,n) \mapsto \mathcal{W}(r-1, n-k)$.

We can now make the key definition of this section. For a positroid $\mathcal{M}$, we denote by

$$(7.4) \qquad\qquad \mathcal{W}_\lambda(U)(\mathcal{M}) = \{x \in \mathcal{W}_\lambda(U): x\big|_{I \mapsto E_r} = 0 \text{ for all } I \notin \mathcal{M}\}.$$

We get a collection of subspaces $\mathcal{W}_\lambda(U)(\mathcal{M})$ of the SL_r-tensor invariant space indexed by positroids $\mathcal{M}$ for the Grassmannian $\mathrm{Gr}(k,n)$.

If W is an r-weblike subgraph, its partial evaluation along I is zero if and only if W has no consistent labelings in which there are r's at the boundary edges indicated by I, cf. Example 7.6. We note that the partial evaluation of a web W will often decompose as a *sum* of several SL_{r-1}-web invariants.

Theorem 7.3. *Let $\mathcal{M}$ be a positroid with positroid variety $\Pi_\mathcal{M}$. Then we have $\mathcal{W}_\lambda(U)(\mathcal{M}) = (\mathfrak{I}(\mathcal{M})_\lambda)^\perp$. In particular, these two spaces have the same dimension.*

Proof. $\mathfrak{I}(\mathcal{M})_\lambda$ is spanned by the r-fold products of Plücker coordinates $\Delta_{I_1}\cdots\Delta_{I_r}$ satisfying (5.2) such that at least one of the I_i satisfies $I_i \notin \mathcal{M}$. By reordering indices, let us assume that $I_r \notin \mathcal{M}$.

Thus, the dual space $(\mathfrak{I}(\mathcal{M})_\lambda)^\perp$ is the subspace of tensor invariants $x \in \mathcal{W}_\lambda(U)$ that pair to zero with each of these r-fold products. From (5.15), pairing to 0 with $(I_1,\ldots,I_r)$ means that $x\big|_{E_\mathcal{S}} = 0$ where $\mathcal{S}$ is dual to $(I_1,\ldots,I_r)$. If we fix I_r, this holds for all choices of $I_1,\ldots,I_{r-1}$ if and only if the partial evaluation $x\big|_{I_r \mapsto E_r} = 0 \in \mathcal{W}_\mu(U')$. By varying I_r over all subsets not in $\mathcal{M}$, we see that the dual space is exactly (7.4) as claimed. $\qquad\square$

Theorem 7.4. *Let G be a planar bipartite graph with positroid $\mathcal{M}$. Then the subspace $\mathcal{W}_\lambda(U)(\mathcal{M})$ is spanned by either of the following sets:*

- *the elements $\mathrm{Web}_r(N; \lambda)$, as N varies over the (infinitely many) networks whose underlying graph is G;*
- *the elements $\mathbf{W}$, as W varies over the (finitely many) r-weblike subgraphs of G with degree λ.*

This characterization of $\mathbb{C}[\Pi_\mathcal{M}]_\lambda$ may be easier to work with than the characterization using promotion and cyclic Demazure crystals given in [11, Section 12].

Proof. First we prove that the various $\mathrm{Web}_r(N; \lambda)$ span the subspace $\mathcal{W}_\lambda(U)(\mathcal{M})$. From (5.17), we see immediately that $\mathrm{Web}_r(N; \lambda)$ pairs to zero with any r-fold product as in the proof of Theorem 7.3 (since $\Delta_{I_r}(N) = 0$ on $\Pi_\mathcal{M}$). Thus $\mathrm{Web}_r(N; \lambda) \in \mathcal{W}_\lambda(U)(\mathcal{M})$.

On the other hand, if $f \in \mathbb{C}[\mathrm{Gr}(k, n)]_\lambda$ pairs to 0 with every $\mathrm{Web}_r(N; \lambda)$, then by (5.16) we have $f(\tilde{X}(N)) = 0$ for every network N representing G. Since the set of points of the form $\tilde{X}(N)$ is Zariski dense in the affine cone over $\Pi_\mathcal{M}$, we deduce that f is the zero polynomial, i.e. $f \in \mathfrak{I}(\mathcal{M})_\lambda$. This shows the various $\mathrm{Web}_r(N; \lambda)$ span the dual space $(\mathfrak{I}(\mathcal{M})_\lambda)^\perp$.

Finally, the span of the various $\mathrm{Web}_r(N; \lambda)$ coincides with the span of the weblike subgraphs of G, as follows from the from the argument inverting the transition matrix (6.8) in the proof of Theorem 6.1. □

Thus for a fixed G associated to a positroid $\mathcal{M}$, the r-weblike subgraphs contained in G are forced to evaluate to 0 on certain basis tensors. This seems to be a completely new phenomenon. Also surprising is the fact that the elements $\mathbf{W}$, for r-weblike subgraphs of a bipartite graph G representing the top cell, must span the invariant space $\mathcal{W}(k, n)$.

The statement that the weblike subgraphs of G span the subspace (7.4) also implies some non-obvious compatibilities between webs, the partial evaluation map (7.2) and positroids. We indicate one such compatibility in the following examples.

Example 7.5. Let $I_{\mathrm{last}} = \{n - k + 1, \ldots, n\}$ so that $\Delta_{I_{\mathrm{last}}}$ is the Plücker coordinate occupying the last k columns. We consider the positroid $\mathcal{M} = \{\binom{[n]}{k}\} \backslash I_{\mathrm{last}}$. It is a positroid, because it can be written as an appropriate intersection of cyclically rotated Schubert matroids [14]. The ideal $\mathfrak{I}(\Pi)_\lambda$ is the set of multiples of $\Delta_{I_{\mathrm{last}}}$ inside $\mathfrak{I}(\Pi)_\lambda$, i.e. $\mathfrak{I}(\Pi)_\lambda = \Delta_{I_{\mathrm{last}}} \mathbb{C}[\mathrm{Gr}(k, n)]_\mu$, where μ is obtained from λ by decrementing along I.

If G is a bipartite graph with positroid $\mathcal{M}$, let $W_1, \ldots, W_s$ be a set of r-weblike subgraphs of G of degree λ such that $\mathbf{W}_1, \ldots, \mathbf{W}_s$ is a basis for $\mathcal{W}_\lambda(U)(\mathcal{M})$. We can extend $\mathbf{W}_1, \mathbf{W}_2, \ldots, \mathbf{W}_s$ to a basis $\mathbf{W}_1, \ldots, \mathbf{W}_s, \mathbf{W}_{s+1}, \ldots, \mathbf{W}_{s+t}$ of $\mathcal{W}_\lambda(U)$ consisting entirely of weblike graphs. Then the number t is given by

$$(7.5) \qquad t = \dim(\mathfrak{I}(\Pi)_\lambda) = \dim(\mathbb{C}[\mathrm{Gr}(k, n)]_\mu) = \dim(W_\mu(U')),$$

where the last equality follows from duality. On the other hand, since $\mathrm{span}(\mathbf{W}_{s+1}, \ldots, \mathbf{W}_{s+t}) \cap \mathcal{W}_\lambda(U)(\mathcal{M}) = \{0\}$, we deduce that the partial evaluations of $\mathbf{W}_{s+1}, \ldots, \mathbf{W}_{s+t}$ along I_{last} are linearly independent in $W_\mu(U')$. Thus, these partial evaluations are a basis for $W_\mu(U')$. As we have already said, these partial evaluations are not guaranteed to be SL_{r-1}-webs.

Example 7.6. Let us the illustrate the previous example in a particular case. We let $k = r = 3$ and $\lambda = (1, \ldots, 1)$, so we are considering the space $\mathcal{W}(3, 9)$ of SL_3-invariant
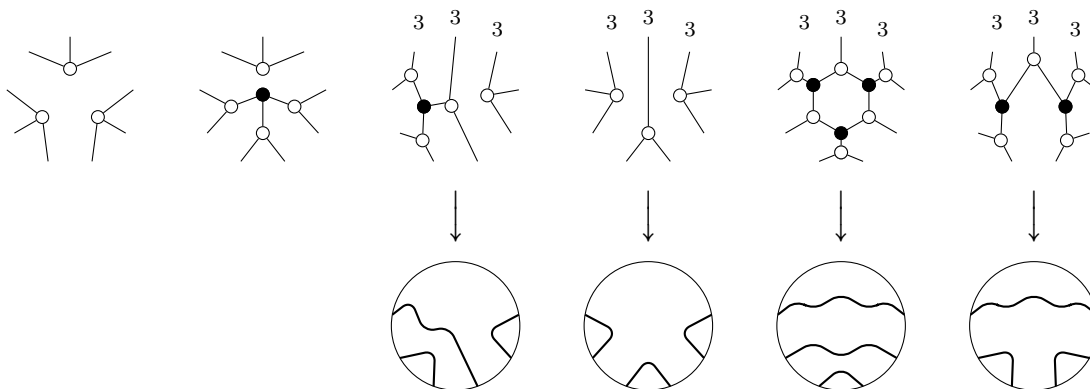
FIGURE 2. SL_2-webs from SL_3-webs via partial evaluation.

multilinear functions of 9 vectors $v_1, \ldots, v_9$. Its dimension is the number of 3×3 standard Young tableaux, which is 42. The non-elliptic basis webs (up to rotation) are listed in the top row of Figure 2. We will focus on the positroid $\mathcal{M} = \binom{[9]}{3} \backslash \{\Delta_{789}\}$ as in the preceding example. The subspace $\mathcal{W}(3,9)(\mathcal{M})$ is the kernel of the linear map $\mathcal{W}(3,9) \to \mathcal{W}(2,6)$ given by specializing $v_7 = v_8 = v_9 = E_3$. Concretely, this means placing 3's at boundary vertices $7, 8$, and 9, which will force certain interior edges to also be labeled by 3's, and then studying consistent labelings of the leftover edges once these 3's have been placed.

A web will be in the subspace $\mathcal{W}(3,9)(\mathcal{M})$ if and only if it either has a fork between vertices 7 and 8 or has a fork between vertices 8 and 9. By examining the forks in Figure 2, we see that there are exactly 37 webs that are in $\mathcal{W}(3,9)(\mathcal{M})$, and five webs that are not. We have placed 3's at the appropriate locations in four of the nonvanishing webs. The fifth nonvanishing web (not pictured) is obtained by reflecting the third web in Figure 2 along the vertical axis (with 3's at the same boundary vertices). From Theorem 7.4, it follows that $\dim(\mathbb{C}[\Pi(\mathcal{M})]_{(1,\ldots,1)}) = 37$.

To see that that $\dim(\mathbb{C}[\Pi(\mathcal{M})]_{(1,\ldots,1)}) = 37$ directly, let $W \notin \mathcal{W}(3,9)(\mathcal{M})$ be one of the five non-vanishing webs. As we said above, the placement of 3's at the boundary forces certain interior edges to be 3's. The remaining edges in W must be labeled by 1's and 2's. These remaining edges naturally decompose into a union of three noncrossing paths joining boundary vertices in pairs, as we demonstrate schematically via the downward arrows in Figure 2. The five SL_2-webs we get in this way are exactly the 5 crossingless matchings on 6 vertices, i.e. a set of basis webs for $\mathcal{W}(2,6)$. In confirmation of Theorem 7.3, we can see concretely that no nontrivial linear combination of the five SL_3-webs combines to give an element of $\mathcal{W}(3,9)(\mathcal{M})$ (and thus, $\mathcal{W}(3,9)(\mathcal{M})$ is spanned by the 37 webs that vanish under partial evaluation).

8. Duality of symmetric group representations

In this section, we deal only with the multilinear web spaces $\mathcal{W}(r,n)$ and $\mathcal{W}(k,n)$ where $n = kr$. These spaces are $\mathbb{C}[\mathrm{Gr}(r,n)]_\lambda$ and for $\mathbb{C}[\mathrm{Gr}(k,n)]_\lambda$ for $\lambda = (1, \ldots, 1)$, so we no longer use the symbol λ to denote the degree of a web. Instead λ will be used to denote a partition of n.

Both tensor invariant spaces $\mathcal{W}(k, n)$ and $\mathcal{W}(r, n)$ carry an action of the symmetric group by permutations of the vectors. Furthermore, these S_n-modules are irreducible, and are related to each other by tensoring with the sign representation ϵ.

Theorem 8.1. *The immanant map $\mathcal{W}(r, n)^* \to \mathcal{W}(k, n) \otimes \epsilon$ is an isomorphism of S_n-modules.*

Since the S_n-modules in Theorem 8.1 are irreducible, the map in Theorem 8.1 is unique up to a scalar factor. The S_n-equivariance of the immanant map does not seem obvious, since there is not a natural action of S_n on the space of networks.

Before proving Theorem 8.1 we recall the relevant notions from S_n representation theory.

Let $\lambda = (\lambda_1, \dots, \lambda_r) \vdash n$ be a Young diagram, with λ_1 the number of boxes in the first row, λ_2 the number of boxes in the second row, and so on. Let $S_\lambda = S_{\lambda_1} \times \cdots \times S_{\lambda_r} \subset S_n$ be the corresponding Young subgroup. A *tableau* for a shape $\lambda \vdash n$ is a filling of λ by the numbers $1, \dots, n$. A *tabloid* is an equivalence class of tableaux, where we identify two tableaux if their entries differ by a permutation in each row. The free vector space M_λ on the set of tabloids of shape λ is a right S_n-module. It is the induced representation $\text{Ind}\,\mathbb{1}_{S_\lambda}^{S_n}$ where $\mathbb{1}$ is the trivial representation. We use the notation $[T]$ for the tabloid determined by a tableau T.

We denote by $N_\lambda = M_\lambda \otimes \epsilon = \text{Ind}\,\epsilon_{S_\lambda}^{S_n}$. It has a basis consisting of *anti-tabloids*, which are defined similarly to tabloids but with the added requirement that swapping two entries in a given row contributes a multiplicative factor of -1. We use the notation $\{T\}$ for the antitabloid determined by T.

Both of the S_n-modules M_λ and N_λ are reducible. To obtain S_n irreducibles, we recall that the *polytabloid* determined by T is the signed sum of tabloids

$$(8.1) \qquad \text{poly}(T) = \sum_{\sigma \in \text{col}(T)} \text{sign}(\sigma)[T \cdot \sigma] \in M_\lambda,$$

where $\text{col}(T)$ is the subgroup of S_n consisting of permutations that fix the entries in each column of T. Then the *Specht Module* S_λ is the subspace of M_λ spanned by the polytabloids T, as T varies over all fillings of λ. The various Specht modules S_λ, for $\lambda \vdash n$, are exactly the irreducibles for S_n.

For a Young diagram λ, we let λ^t denote the conjugate or transpose partition (likewise T^t denotes the conjugate tableau to T). Then the Specht modules S_λ and S_{λ^t} are related by $S_\lambda = S_{\lambda^t} \otimes \epsilon$.

For the space $\mathcal{W}(k, n)$, a list of boundary location subsets $(I_1, \dots, I_r)$ is the same as an $r \times k$ tableau $T(I_1, \dots, I_r)$ whose rows are $I_1, \dots, I_r$. The boundary label subsets $\mathcal{S} = (S_1, \dots, S_n)$ are singletons such that S_j records the row of $T(I_1, \dots, I_r)$ in which j appears. The sign of $\mathcal{S}$, which we can denote by $\text{sign}(T)$, is determined by the well-known *descent number* of T modulo 2, where we recall that the number of descents of a tableaux is the number of pairs $i > j$ but i is in a lower row of T than j.

Letting λ denote the $r \times k$ rectangle, there is an isomorphism S_n-modules

$$(8.2) \qquad\qquad\qquad \mathcal{W}(k, n) \leftrightarrow S_{\lambda^t}$$

$$(8.3) \qquad\qquad\qquad \Delta_{I_1} \Delta_{I_2} \cdots \Delta_{I_r} \mapsto \text{poly}(T^t)$$

where $T = T(I_1, \ldots, I_r)$. Indeed, consider a $k \times n$ matrix of indeterminates x_{ij}. The homogeneous coordinate ring $\mathbb{C}[\mathrm{Gr}(k,n)]$ can be identified with the invariant subring $\mathbb{C}[x_{ij}]^{\mathrm{SL}_k} \subset \mathbb{C}[x_{ij}]$. For a set $J = \{j_1, \ldots, j_r\} \in \binom{[n]}{r}$, denote by x_{iJ} the monomial $x_{ij_1} x_{ij_2} \ldots x_{ij_r}$. For a tableau $T = T(I_1, \ldots, I_r)$ whose columns are $J_1, \ldots, J_k$, we can consider the monomial

$$\prod_{i=1}^{k} x_{iJ_i}.$$

Clearly this monomial only depends on $[T^t]$. The expansion of $\Delta_{I_1} \cdots \Delta_{I_r}$ as a signed sum of such monomials agrees with the expression of (8.1) of $\mathrm{poly}(T^t)$ as a signed sum of tabloids. Thus the subspace spanned by $\mathrm{poly}(T^t)$, which is the Specht module S_{λ^t}, is isomorphic to $\mathcal{W}(k,n)$ as a representation of S_n.

Proof of Theorem 8.1. The idea is to fit the immanant map into a commuting diagram of maps of S_n-modules, from which we deduce that the immanant map is S_n-equivariant. We remark that, all the maps in this diagram do not require λ to be of rectangular shape, and are unique up to scalars. The only assertion that requires λ to be of rectangular shape is the interpretation of one of these maps as the immanant map.

First, it easy to check that the pairing

$$M_\lambda \otimes N_\lambda \to \epsilon$$
$$[S] \otimes \{T\} \mapsto \delta_{S,T} \mathrm{sign}(S).$$

is S_n-equivariant, where $\delta_{S,T}$ is the Kronecker delta. In this way, we get an S_n-equivariant map $N_\lambda \to M_\lambda^* \otimes \epsilon$. Since $S_\lambda \subset M_\lambda$, we can further compose $N_\lambda \to M_\lambda^* \otimes \epsilon \to S_\lambda^* \otimes \epsilon$.

On the other hand, there is always an S_n-equivariant map $N_\lambda \to S_{\lambda^t}$ induced by $\{T\} \mapsto \mathrm{poly}(T^t)$. The map (8.3) can be thought of as a special case of this. (The space of formal sums of $\Delta_{I_1} \cdots \Delta_{I_r}$ – before imposing Plücker relations – is naturally identified with N_λ). Then the immanant map $(S_\lambda)^* \otimes \epsilon \to S_{\lambda^t}$ is the unique (up to scalars) map making the diagram

(8.4)
$$N_\lambda \longrightarrow (M_\lambda)^* \otimes \epsilon \longrightarrow (S_\lambda)^* \otimes \epsilon$$

with N_λ mapping down to S_{λ^t} and the diagonal map labeled Imm.

Let us explain why this diagonal map is the immanant map. Going down in (8.4) is the map that replaces a formal sum of products $\Delta_{I_1} \cdots \Delta_{I_r}$ with their images in $\mathbb{C}[\mathrm{Gr}(k,n)]$ (i.e., going down is imposing the Plücker relations).

Let us now examine the top row of the commutative diagram. The space M_λ can be identified with monomials in y_{ij} which are entries in an $r \times n$ matrix, and which are multilinear in each column, by the analysis above. The map $N_\lambda \to M_\lambda^* \otimes \epsilon$ takes an antitabloid $\{T\}$ to $\mathrm{sign}(T)$ times the functional that picks out the coefficient of the tabloid $[T]$. Putting this together, the antitabloid $\{T\}$ gets sent to $\mathrm{sign}(T)$ times the coefficient of $\prod_i y_{iI_i}$. When we restrict this to the Specht module, we get exactly the functional $\mathrm{sign}(\mathcal{S})\mathrm{eval}(E_{\mathcal{S}})$, where $\mathrm{eval}(E_{\mathcal{S}})$ is as in the proof of Proposition 4.6 . The commutativity now follows since the immanant map is characterized by (5.15). $\qquad \square$

This appendix is jointly written with Darlayne Addabbo, Eric Bucher, Sam Clearman, Laura Escobar, Ningning Ma, Suho Oh, and Hannah Vogel. For simplicity, we restrict attention to the multilinear case $\lambda = (1, \ldots, 1)$ and $n = kr$. The immanant map provides us with an isomorphism $\mathrm{Imm} \colon \mathcal{W}(r,n)^* \to \mathcal{W}(k,n)$. When r is equal to 2 or 3, the space $\mathcal{W}(r,n)$ has a distinguished choice of *web basis* $\mathcal{B}$ (given by crossingless matchings, and non-elliptic SL$_3$ webs, respetively). For a given basis web W, we denote by $\varphi_W \in \mathcal{W}(r,n)^*$ the dual basis element with respect to $\mathcal{B}$. Its image $\mathrm{Imm}(\varphi_W) \in \mathcal{W}(k,n)$ is called a *web immanant* in [10].

Our main observation in this appendix is the following:

Observation 8.2. Let $r = 2$ or 3, and $W \in \mathcal{W}(r,n)$ be a basis web. Then the web immanants $\mathrm{Imm}(\varphi_W) \in \mathcal{W}(k,n)$ are also web invariants (up to a sign).

That is, the unique S_n-equivariant pairing between the tensor invariant spaces $\mathcal{W}(k,n)$ and $\mathcal{W}(r,n)$ induces a duality between basis webs in these spaces, in small cases. We verified this observation in small cases by straightforward calculations which we omit here (see [10] for some examples). The resulting pairing is rotation-equivariant, and we find the resulting pictures appealing, cf. Figure 3.

Our second observation is as follows. When $r = 2$ or 3, there is a well-known bijection between standard young tableaux and non-elliptic webs, due to Khovanov and Kuperberg [5] (see also [18]). It has the property that rotation of webs is given by promotion of standard young tableaux, and this can be used to give an elementary proof [15] of the cyclic sieving phenomenon for rectangular tableaux with 2 or 3 or rows. Our next observation is that under this bijection, dual basis webs correspond to transposed tableaux:

Observation 8.3. In small cases, duality between basis webs is given by transposing standard young tableaux.

Specifically, we checked that this is true for the dualities $\mathcal{W}(3,6) \leftrightarrow \mathcal{W}(2,6)$ and $\mathcal{W}(3,9) \leftrightarrow \mathcal{W}(3,9)$. For $r \geq 4$, Westbury has given a basis of SL$_r$-webs indexed by standard young tableaux [19]. In the special case of $\mathcal{W}(4,8)$, the basis of web immanants listed in the second column of Figure 3 is *different* than Westbury's basis for $\mathcal{W}(4,8)$. The web immanant basis consists of the 4 rotation classes of the first type of web, the 8 rotation classes of the second type of web, and the 2 rotation classes of the third type of web. We remark that this third type of SL$_4$-web is fixed by rotating two units, which is not obvious, but is an instance of the SL$_4$-square move. The web immanant basis disagrees with Westbury's basis in these last two elements – the element Westbury assigns to these two tableaux has order 4 under rotation. Once the elements of Westbury's basis are replaced by the corresponding web immanants, we again have that rotation of webs is given by promotion. This suggests that there might be a slightly different choice of web basis, indexed by tableaux, that is better behaved with respect to promotion.

Let us also note that the last SL$_5$-web in Figure 3 *must* be fixed by rotating two units, because this is true of its dual SL$_2$-web. And indeed, this can be checked using the SL$_5$ diagrammatic relations – applying a square move to the top square of this web produces the same web, but rotated 4 units clockwise (so applying the square move three times, we get rotation by 2 units).

We mention that Observation 8.3 is probably implied by similar statements guaranteeing duality between the *canonical bases* for $\mathcal{W}(r,n)$ and $\mathcal{W}(k,n)$ [12], and from an agreement between the web basis and canonical basis in these instances. Nonetheless, the pictures in Figure 3 are new. When r and k are both > 3, the immanant map gives us a bilinear pairing of $\mathcal{W}(k,n)$ with $\mathcal{W}(r,n)$. It would be interesting to understand the resulting pairing between SL_k-webs and SL_r-webs combinatorially.
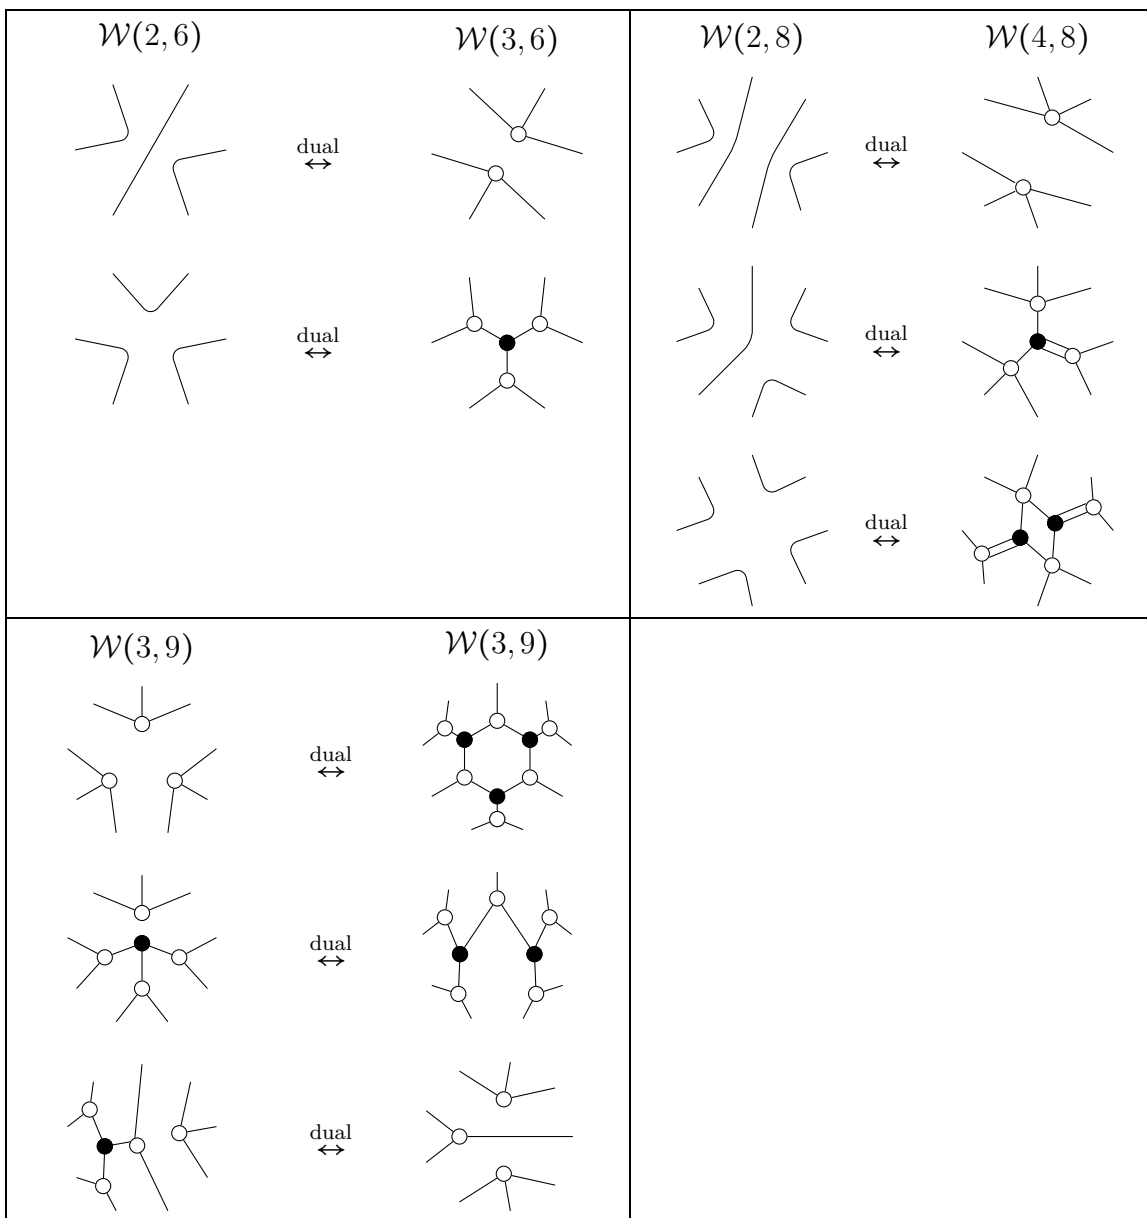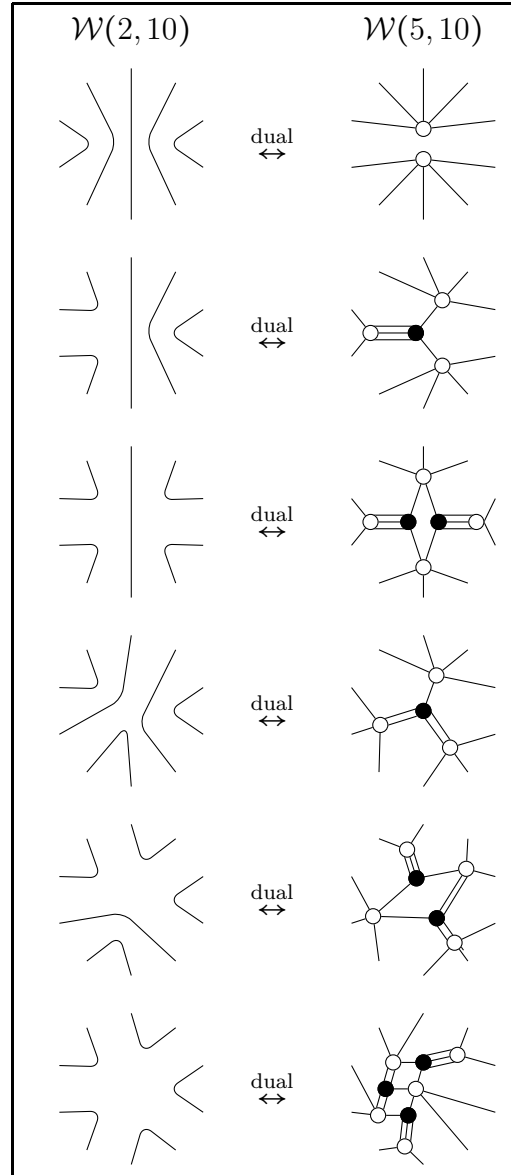

FIGURE 3. Duality between webs in small cases, drawn in terms of weblike subgraphs. The last case is on the next page. Edges of multiplicity two or three are depicted by doubled or tripled edges.

References

[1] S. Cautis, J. Kamnitzer, and S. Morrison. Webs and quantum skew Howe duality. *Math. Ann.*, 360:351–390, (2014).

[2] S. Fomin and P. Pylyavskyy. Webs on surfaces, rings of invariants, and clusters. *Proc. Natl. Acad. Sci*, 111(27):9680–9687, (2014).

[3] S. Fomin and P. Pylyavskyy. Tensor diagrams and cluster algebras. *Adv. Math.*, 300: 717–787, (2016).

[4] K. R. Goodearl and M. Yakimov. Poisson structures on affine spaces and flag varieties. ii. general case. *Trans. Amer. Math. Soc.*, 361(11):5753–5780, (2009).

[5] M. Khovanov and G. Kuperberg. Webs bases for $sl(3)$ are not dual canonical. *Pacific J. Math.*, 188(1):129–153, (1999).

[6] D. Kim. Graphical calculus on representations of quantum Lie algebras. *Ph.D thesis, UC Davis*, (2003).

[7] A. Knutson, T. Lam, and D. Speyer. Positroid varieties: juggling and geometry. *Compos. Math*, 149(10):1710–1752, (2013).

[8] E. H. Kuo. Applications of graphical condensation for enumerating matchings and tilings. *Theoret. Comput. Sci.*, 319(1–3):29–57, (2004).

[9] G. Kuperberg. Spiders for rank 2 Lie algebras. *Commun. Math. Phys.*, 1(180): 109–151, (1996).

[10] T. Lam. Dimers, webs, and positroids. *J. Lond. Math. Soc.*, 92(3):633–656, (2015).

[11] T. Lam. Totally nonnegative Grassmannian and Grassmann polytopes. *Current developments in mathematics 2014*, pages 51–152, (2016).

[12] G. Lusztig. Canonical bases in tensor products. *Proc. Nat. Acad. Sci. U.S.A*, 89(17): 8177–8179, (1992).

[13] S. Morrison. A diagrammatic catgeory for the representation theory of $U_q(\mathfrak{sl}_n)$. *Ph.D thesis, UC Berkeley*, (2007).

[14] S. Oh. Positroids and Schubert matroids. *J. Combin. Theory Ser. A*, 118(8):2426–2435, (2011).

[15] K. T. Peterson, P. Pylyavksyy, and B. Rhoades. Promotion and cyclic sieving via webs. *J. Algebraic Comb.*, 30(1):19–41, (2009).

[16] A. Postnikov. Total positivity, Grassmannians, and Networks. *arXiv preprint*, (2006). `arXiv:0609764` [math.CO].

[17] A. Postnkov, L. Williams, and D. Speyer. Matching polytopes, toric geometry, and the totally non-negative Grassmannian. *J. Algebraic Combin.*, 30(2):173–191, (2009).

[18] J. Tymoczko. A simple bijection between standard $3 \times n$ tableaux and irreducible webs for $\mathfrak{sl}_3$. *J. Algebraic Combin.*, 35(4):611–632, (2012).

[19] B. Westbury. Web bases for the general linear groups. *J. Algebraic Combin.*, 35(1): 5753–5780, (2012).

DEPARTMENT OF MATHEMATICAL SCIENCES, INDIANA UNVERSITY–PURDUE UNIVERSITY INDIANAPOLIS, 402 N. BLACKFORD ST., INDIANAPOLIS, 46202.
E-mail address: `chfraser@iupui.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, 530 CHURCH ST., ANN ARBOR MI 48109.
E-mail address: `tfylam@umich.edu`

PERIMETER INSTITUTE FOR THEORETICAL PHYSICS, 31 CAROLINE STREET NORTH, WATERLOO, CANADA, N2L 2Y5.
E-mail address: `ile@perimeterinstitute.ca`