

A geometric construction of colored HOMFLYPT homology

BEN WEBSTER
GEORDIE WILLIAMSON

The aim of this paper is twofold. First, we give a fully geometric description of the HOMFLYPT homology of Khovanov and Rozansky. Our method is to construct this invariant in terms of the cohomology of various sheaves on certain algebraic groups, in the same spirit as the authors' previous work on Soergel bimodules. All the differentials and gradings which appear in the construction of HOMFLYPT homology are given a geometric interpretation.

In fact, with only minor modifications, we can extend this construction to give a categorification of the colored HOMFLYPT polynomial, *colored HOMFLYPT homology*. We show that it is in fact a knot invariant categorifying the colored HOMFLYPT polynomial and that it coincides with the categorification proposed by Mackaay, Stošić and Vaz.

17B10, 57T10

1 Introduction

The *colored HOMFLYPT polynomial* is an invariant of links together with a labeling or “coloring” of each component with a positive integer; in particular, for knots, there is an invariant for each positive integer. Its most important properties are that

- it reduces to the usual HOMFLYPT polynomial when all labels are 1, and
- colored HOMFLYPT encapsulates all Reshetikhin–Turaev invariants for the link labeled with wedge powers of the standard representation of \mathfrak{sl}_n , just as the HOMFLYPT polynomial does for the standard representation alone.

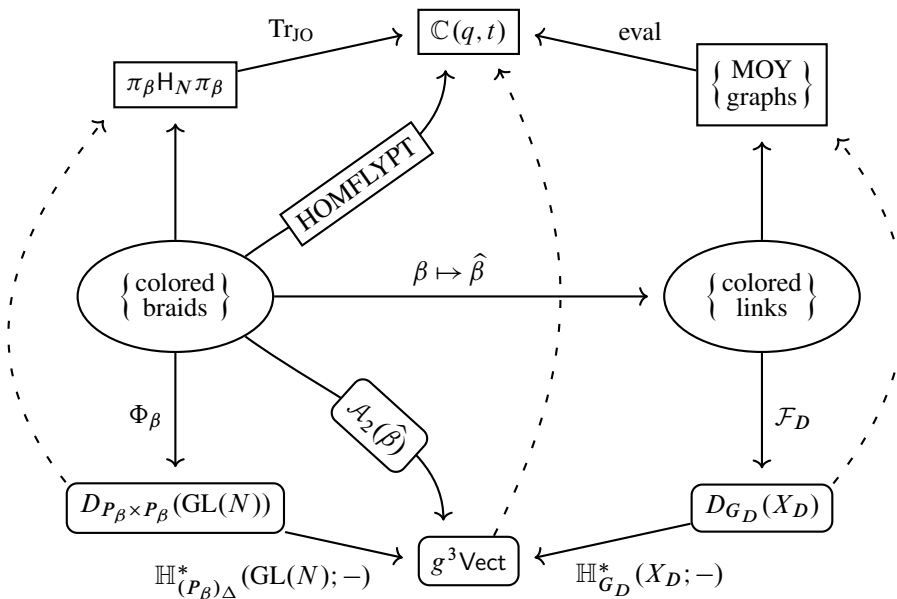
In this paper we give a geometric construction of a categorification of this invariant, *colored HOMFLYPT homology*. Like the HOMFLYPT homology of Khovanov and Rozansky [13], this associates a triply graded vector space to each colored link such that the bigraded Euler characteristic is the colored HOMFLYPT polynomial. In fact, we produce an infinite sequence of such invariants, one for each page of a spectral sequence, but only the first and second pages are connected via an Euler characteristic to a known classical invariant.

Our construction and proofs of invariance and categorification are algebro-geometric in nature. As a special case we obtain a new and entirely geometric interpretation of Khovanov’s Soergel bimodule construction of HOMFLYPT homology [12].

We also show that this invariant has a purely combinatorial description via the Hochschild homology of bimodules analogous to that of Khovanov. In fact, it coincides with the link homology proposed from an algebraic perspective by Mackaay, Stošić and Vaz [17]. Thus, the main result of our paper has an entirely algebraic statement:

Theorem 1.1 *The colored HOMFLYPT homology defined in [17] is a knot invariant, and its Euler characteristic is the colored HOMFLYPT polynomial.*

Our definition also has the advantage of categorifying essentially all algebraic objects involved in the definition of colored HOMFLYPT homology. Let us give a schematic diagram for the pieces here, with actual operations given by solid arrows, and (de)categorifications given by dashed ones:



The top half of the diagram shows two different definitions of the colored HOMFLYPT polynomial:

- The path through $\{\text{MOY graphs}\}$ is the description of the colored HOMFLYPT polynomial by Murakami, Ohtsuki and Yamada [19]: one associates to a link diagram a sum of weighted trivalent graphs, and then defines an evaluation function on such

graphs, which in turn gives a state sum interpretation of the colored HOMFLYPT polynomial. While the paper [19] only considers certain specializations of the HOMFLYPT polynomial, their technique is easily extended to the polynomial itself.

- The path through $\pi_\beta H_N \pi_\beta$ is described by Lin and Zheng [16]: to each closable colored braid β , we have an associated element of the Hecke algebra H_N , where N is the colored braid index of β (the sum of the colorings of the strands). In fact, this element lies in a certain subalgebra $\pi_\beta H_N \pi_\beta$, where π_β is a projection which depends on the coloring of β . The colored HOMFLY polynomial is obtained by applying a certain trace Tr_{JO} defined by Ocneanu [11] on H_N .

In this paper, we show how to categorify both of these paths, as is schematically indicated in the bottom half of the diagram, and briefly summarized in Section 1.2. The final result of this construction is a knot invariant $\mathcal{A}_2(\hat{\beta})$; we show that this invariant is well-defined in Theorem 1.2 and that it agrees with HOMFLYPT homology in Theorem 1.4.

- The leftmost dashed arrow is the isomorphism of $\pi_\beta H_N \pi_\beta$ with the Grothendieck group of sheaves on $GL(N)$ which are bi-equivariant for the left and right multiplication of a subgroup of block upper triangular matrices P_β .
- The rightmost dashed arrow can be described as follows: to each link diagram D , we associate a group G_D , a G_D -variety X_D and a G_D -equivariant perverse sheaf whose the composition factors are in bijection with the MOY graphs arising from this link diagram.
- The central dashed arrow simply indicates taking bigraded Euler characteristic of a trigraded vector space with respect to one of its gradings.

We must also show that this diagram, including the dashed arrows, “commutes”. This follows from a result of the authors giving a similar construction of a Markov trace for the Hecke algebra of any semisimple Lie group, shown in the paper [27].

As should be clear from the above, the techniques we use are those of algebraic geometry and geometric representation theory. While these are not familiar to the average topologist, we have striven to make this paper accessible to the novice, at least if they are willing to accept a few deep results as black boxes. As a general rule, our actual calculations are simple and quite geometric in nature; however, we must cite rather serious machinery to show that these calculations are meaningful.

1.1 The geometric machinery

Let us briefly indicate the geometric setting in which we work. All material covered here is discussed in greater detail in Section 3.

Let X be an algebraic variety defined by equations with integer coefficients. (In this paper, our varieties are built from copies of the general linear group, so we can always describe them in terms of integral equations.) To X one may associate a derived category $D^b(X)$ of sheaves with constructible cohomology. There are numerous technicalities in the construction of this category, but we postpone discussion of these until Section 3.

The category $D^b(X)$ behaves similarly to the bounded derived category of constructible sheaves on the complex algebraic variety defined by these equations. However, since we used integral equations, we have an alternate perspective on these varieties; one can also reduce modulo a prime p , and work over the finite field \mathbb{F}_p . The objects in $D^b(X)$ can also be interpreted as sheaves on these varieties in characteristic p , and for technical reasons, this is the perspective we will take. In this situation, there is an extra structure which helps us to understand our complexes of sheaves: an action of the Frobenius Fr on our variety.

The category $D^b(X)$ contains a remarkable abelian subcategory $P(X)$ of “mixed perverse sheaves”. For us the most important feature of $P(X)$ is that every object of $P(X)$ has a canonical “weight filtration” with semisimple subquotients, which is defined using the Frobenius.

As with any filtration, this leads to a spectral sequence

$$E_1^{p,q} = \mathbb{H}^{p+q}(\text{gr}_{-p}^W \mathcal{F}) \implies \mathbb{H}^{p+q}(\mathcal{F}).$$

Each term on the left hand side also carries an action of Frobenius induced by that on the variety. Considering the norms of the eigenvalues of Frobenius may be used to give an additional grading to each page of the spectral sequence. It follows that each page of the spectral sequence is *triple* graded.

We will need to consider a generalization of this category, which is a version of equivariant sheaves for the action of an affine algebraic group on X . While in principle, the technical difficulties of understanding such a category could be resolved by working in the category of stacks, we have found it less burdensome to give a careful definition of the mixed equivariant derived category from a more elementary perspective. For the sake of brevity, this has been done in a separate note [25].

1.2 Application to knot theory

In order to apply the above machinery to knot theory, we must define a sheaf associated to a link. More precisely, as we discuss in Section 2, to any projection D of a colored link, we associate the natural graph with vertices given by crossings and edges by arcs.

To this graph, we associate a variety X_D together with the action of a reductive group G_D . Remembering the crossings in D allows us to construct a G_D -equivariant mixed shifted perverse sheaf $\mathcal{F}_D \in D_{G_D}^b(X_D)$. We then show that \mathcal{F}_D may be used to construct a series of knot invariants.

Associated to any filtration on \mathcal{F}_D (as a perverse sheaf), we have a canonical spectral sequence converging to $\mathbb{H}_{G_D}^*(X_D; \mathcal{F}_D)$. Furthermore, we can endow $\mathbb{H}_{G_D}^*(X_D; -)$ of any mixed sheaf with the weight grading. This is a grading which is preserved by all spectral sequence differentials, so we can think of any page of this spectral sequence as a triply graded vector space, where two gradings are given by the usual spectral sequence structure, and the third by weight.

The sheaf \mathcal{F}_D carries a natural weight filtration. This is easily confused with, but distinct from, the weight grading discussed above.¹ We call the spectral sequence associated to this weight filtration *chromatographic*.

Theorem 1.2 *If D is the diagram of a closed braid, then every page E_i for $i \geq 2$ of the spectral sequence computing $\mathbb{H}_{G_D}^*(X_D; \mathcal{F}_D)$ associated to the weight filtration is an invariant of the underlying link L , up to an overall shift in the grading. We let $\mathcal{A}_i(\hat{\beta})$ be the i^{th} page of this spectral sequence.*

If D is not a closed braid, then this theorem fails, since $\mathcal{A}_i(\hat{\beta})$ can be changed by the Reidemeister IIb move; above we are using the fact that by the Markov theorem, any two braid closure diagrams for the same knot can be related without using this move.

This description has a similar flavor to that of Khovanov and Rozansky [13] or Bar-Natan [2]: it begins by assigning a simple object to a single crossing, and then an algebraic rule for gluing crossings together (this process can be formalized as an object called a *canopolis* as introduced by Bar-Natan [2]; we will discuss this perspective in Section 6.2). However, other papers, such as [12] or [17], have used a description which depended on the link diagram chosen being a closed braid. In order to show that our invariants coincide with those of [17], we must find a geometric description of this form.

Assume that β is a closable colored braid with coloring given by positive integers, $\hat{\beta}$ its closure and let N be the colored braid index (the sum of the colorings over the strands of the braid). Let P_β be the block upper triangular matrices inside $G_N := \text{GL}(N)$ with the sizes of the blocks given by the coloring of the strands of β . Using left and right multiplication, we obtain a natural $P_\beta \times P_\beta$ action on G_N . We let $(P_\beta)_\Delta$ be

¹The weight grading mentioned above comes from the weight filtration of the pushforward of \mathcal{F}_D to a point.

the diagonal subgroup, which acts on G_N by conjugation. By the classical theory of characteristic classes, we have a canonical isomorphism of $H^*(BP_\beta)$ to partially symmetric polynomials corresponding to the block sizes of P_β , which we will use freely from now on.

Theorem 1.3 *For each β , there is a $(P_\beta \times P_\beta)$ -equivariant complex of sheaves Φ_β on G_N with a natural filtration, such that the associated spectral sequence computing $\mathbb{H}_{(P_\beta)\Delta}^*(G_N; \Phi_\beta)$ is canonically isomorphic to the spectral sequence obtained from the weight filtration for $\mathbb{H}_{G_\beta}^*(X_{\hat{\beta}}; \mathcal{F}_{\hat{\beta}})$.*

Moreover, we have an isomorphism of the E_1 page $\mathcal{A}_1(\hat{\beta})$ of the spectral sequence for the hypercohomology $\mathbb{H}_{P_\beta \times P_\beta}^*(G_N; \Phi_\beta)$, as a complex of bimodules over $H^*(BP_\beta)$, to the complex of singular Soergel bimodules considered by Mackaay, Stošić and Vaz in [17, Section 8].

Singular Soergel bimodules have been defined and classified in the thesis of the second author [28] and in the context of Harish-Chandra bimodules by Stroppel [23]. Since previous work of the authors [26] has related Hochschild homology to conjugation equivariant cohomology, we can identify our geometric knot invariant in terms of such bimodules.

Theorem 1.4 *If D is a closed braid, then the E_2 page of our spectral sequence is the categorification of the colored HOMFLYPT polynomial proposed in [17].*

If all the labels on the components of D are 1, then this agrees with the triply graded link homology as defined by Khovanov and Rozansky in [13].

The higher pages of this spectral sequence are not easy to compute, and it is not known what their Euler characteristics are. Whether they correspond to any classical link invariant is unknown.

Acknowledgements We would like to thank: Wolfgang Soergel for his observation that “Komplexe von Bimoduln sind die Gewichtsfiltrierung des armen Mannes” (“Complexes of bimodules are the poor man’s weight filtration”), which formed a starting point for this work; Marco Mackaay for suggesting that it could be generalized to the colored case and explaining the constructions of [17]; Raphaël Rouquier and Olaf Schnürer for illuminating discussions; and Catharina Stroppel, Noah Snyder and Carl Mautner for comments on an earlier version of this paper. Part of this research was conducted whilst Williamson took part in the program “Algebraic Lie Theory” at the Isaac Newton Institute, Cambridge. Webster was supported by an NSF Postdoctoral Fellowship.

2 Description of the varieties

We start by recalling the steps involved in our categorification, beginning with a braidlike diagram D of an oriented colored link L :

- To D (with its coloring) we associate a reductive group G_D together with a G_D -variety X_D , which only depends on the graph obtained from the diagram D by forgetting the distinction between under and overcrossings.
- The crossing data allows us to define a G_D -equivariant sheaf \mathcal{F}_D on X_D .
- This sheaf \mathcal{F}_D has a chromatographic spectral sequence converging to the G_D -equivariant hypercohomology of \mathcal{F}_D .
- Each page E_i of this spectral sequence for $i \geq 2$ is a knot invariant (up to overall shift) and the E_2 page categorifies the colored HOMFLYPT polynomial.

In this section we discuss the first step.

First let us fix some notation. We fix a chain of vector spaces $0 \subset V_1 \subset V_2 \subset V_3 \subset \dots$ over \mathbb{F}_q such that $\dim V_i = i$ for all i . Let

$$G_{i_1, \dots, i_n} := \mathrm{GL}(i_1) \times \dots \times \mathrm{GL}(i_n),$$

and let P_{i_1, \dots, i_n} be the block upper triangular matrices with blocks $\{i_1, \dots, i_n\}$. We may identify P_{i_1, i_2, \dots, i_n} with the stabilizer in $G_{i_1 + \dots + i_n}$ of the standard partial flag


$$\{0 \subset V_{i_1} \subset V_{i_1 + i_2} \subset \dots \subset V_{i_1 + \dots + i_n}\}.$$

Let D be a diagram of an oriented tangle with marked points, with no marked points occurring at a crossing. Let Γ be the oriented graph obtained by the diagram's projection, with vertices corresponding to crossings and marked points in D . That is, we simply forget the over and undercrossings in D . We deal with the exterior ends of the tangle in a somewhat unconventional manner; we do not think of them as vertices in the graph, so we think of the arcs connecting to the edge as connecting to 1 or 0 vertices. By adding marked points to D if necessary, we may assume that every component of Γ contains at least one vertex.

Recall that to the diagram D we wish to associate a variety X_D acted on by an algebraic group G_D . Let us write $\mathcal{E}(D)$ and $\mathcal{V}(D)$ for the edges and vertices of Γ respectively. Given an edge $e \in \mathcal{E}(D)$ write G_e for G_i , where i is the label on e . Similarly, given $v \in \mathcal{V}(D)$ write G_v for G_i , where i is the sum of the labels on the incoming vertices at v (which necessarily equals the sum of the labels on the outgoing vertices). We define

$$X_D := \prod_{v \in \mathcal{V}(D)} G_v \quad \text{and} \quad G_D := \prod_{e \in \mathcal{E}(D)} G_e.$$

It remains to describe how G_D acts on X_D . Locally, near any crossing, Γ is isotopic to a graph of the form:



We will call e_1 and e_2 *upper* and e_3 and e_4 *lower* edges with respect to the vertex v . Whenever a vertex v lies on an edge e we define an inclusion map $i_e: G_e \rightarrow G_v$, which is the identity if v corresponds to a marked point, and is the composition of the canonical inclusions

$$\begin{cases} G_i \hookrightarrow G_{i,j} \hookrightarrow G_{i+j} & \text{if } e \text{ is upper,} \\ G_i \hookrightarrow G_{j,i} \hookrightarrow G_{i+j} & \text{if } e \text{ is lower.} \end{cases}$$

That is, G_e is included as the upper left or lower right block matrices in G_v , according to whether e is upper or lower.

We now describe how G_D acts on X_D by describing the action componentwise. Let $g \in G_e$ and $x \in G_v$. We have


$$g \cdot x = i_e(g)^\omega x i_e(g)^{-\alpha},$$

where $\omega = 1$ if e is incoming at v , and 0 otherwise, and $\alpha = 1$ if e is outgoing at v and 0 otherwise. Thus, we have

$$g \cdot x = \begin{cases} x & \text{if } v \text{ does not lie on } e, \\ x i_e(g)^{-1} & \text{if } e \text{ is only outgoing at } v, \\ i_e(g) x & \text{if } e \text{ is only incoming at } v. \\ i_e(g) x i_e(g)^{-1} & \text{if } e \text{ forms a loop at the vertex } v. \end{cases}$$

Example 2.1 Here are two examples of X_D and G_D .


- Let D be the standard diagram of the unknot labeled i with one marked point:



Then we have $X_D = G_D = G_i$ and G_D acts on X_D by conjugation.

- If D is an oriented arc with a single marked point, then we have $X_D = G_i \times G_i$, and $G_D = G_i$, with the action $g \cdot (a, b) = (a g^{-1}, g b)$, where a corresponds to the arc leaving the marked point, and b to the arc entering it.

- Let D be the diagram of an (i, j) -crossing:



Here $X_D = G_{i+j}$ and $G_D = G_i \times G_j \times G_j \times G_i$ and (a, b, c, d) acts on $x \in G_{i+j}$ by

$$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix} x \begin{pmatrix} ccc^{-1} & 0 \\ 0 & d^{-1} \end{pmatrix}.$$

Note that if we glue two open edges with label i of the tangle diagram D together to make a new diagram D' by adding a marked point, then the spaces $X_D = X_{D'}$ are isomorphic, but $G_{D'} = G_D \times G_i$, with the new factor acting on the factors in X_D corresponding to the glued edges.

The group G_i attached to a marked point acts freely if the two connected edges don't close into a loop, and removing this point simply quotients both X_D and G_D by this G_i , leaving the equivariant geometry unchanged. Combining these observations with the examples above is enough to construct X_D and G_D for any tangle diagram.

This is the variety and group that we shall use in our construction. But before defining our invariant, we must first cover some generalities on categories of sheaves on these varieties.

3 Mixed and equivariant sheaves

In the rest of this paper, we will be using the machinery of mixed equivariant sheaves. In this section we intend to summarize the essential features of the theory that are necessary for us, and to indicate to the reader where the details can be found.

3.1 Weight grading

An important point underlying what follows is that cohomology of a complex algebraic variety (as well as most variations, such as equivariant cohomology, or intersection cohomology) has an additional natural grading, the *weight grading*. This grading is difficult to describe explicitly without using methods over characteristic p (as we will later), but is best understood by two simple properties:

- The weight grading is preserved by cup products, pullback and all maps in long exact sequences (in fact, by all differentials in any Serre spectral sequence).
- This weight grading is equal to the cohomological grading on smooth projective varieties.

Example 3.1 (cohomology of \mathbb{C}^*) If we write $\mathbb{C}P^1$ as the union of \mathbb{C} and $\mathbb{C}P^1 - \{0\}$, then in the Mayer–Vietoris sequence we have an isomorphism $H^2(\mathbb{C}P^1) \cong H^1(\mathbb{C}^*)$. Thus, the cohomological and weight gradings do not agree on $H^1(\mathbb{C}^*)$.

We plan to describe homological knot invariants using the equivariant cohomology of varieties, and the weight grading will be necessary to give all the gradings we expect on our knot homology.

3.2 Sheaves and perverse sheaves

We must use a generalization of the weight grading, the weight filtration on a mixed perverse sheaf. References for this section include [1], [9], [5] and [14]. Although we will not use it below, we should also point out that there is a way to understand mixed perverse sheaves which only uses characteristic-0 methods (Saito’s mixed Hodge modules [22]; see the book of Peter and Steenbrink [21]).

Let $q = p^e$ be a prime power. We consider throughout a finite field \mathbb{F}_q with q elements, and an algebraic closure \mathbb{F} of \mathbb{F}_q . Unless we state otherwise all varieties and morphisms will be defined over \mathbb{F}_q . Given a variety X we will write $X \otimes \mathbb{F}$ for its extension of scalars to \mathbb{F} .

We fix a prime number $\ell \neq p$ and let \mathbb{k} denote the algebraic closure $\overline{\mathbb{Q}}_\ell$ of the field of ℓ -adic numbers. We should note that the choices of p , q and ℓ are purely auxiliary; the resulting knot homology will be independent of all of these choices. Throughout we fix a square root of q in \mathbb{k} and denote it by $q^{1/2}$. Given a variety Y defined over \mathbb{F}_q or \mathbb{F} we denote by $D^b(Y)$ (resp. $D^+(Y)$) the bounded (resp. bounded below) derived category of constructible \mathbb{k} -sheaves on Y ; see [9]. By abuse of language we also refer to objects in $D^b(X)$ or $D^+(X)$ as sheaves. Given a sheaf \mathcal{F} on X we denote by $\mathcal{F} \otimes \mathbb{F}$ its extension of scalars to a sheaf on $X \otimes \mathbb{F}$. Given a sheaf \mathcal{F} on X we abuse notation and write

$$\mathbb{H}^*(\mathcal{F}) := \mathbb{H}^*(X \otimes \mathbb{F}, \mathcal{F} \otimes \mathbb{F}) = \mathbb{H}^*(\mathcal{F} \otimes \mathbb{F}).$$

We *never* consider hypercohomology before extending scalars.

On the category $D^b(X)$, we have the Verdier duality functor $\mathbb{D}: D^b(X) \rightarrow D^b(X)^{op}$, and for each map $f: X \rightarrow Y$, we have Verdier dual pushforward functors

$$f_*, f_!: D^b(X) \rightarrow D^b(Y)$$

(often denoted Rf_* and $Rf_!$), and Verdier dual pullback functors

$$f^*, f^!: D^b(Y) \rightarrow D^b(X).$$

In $D^b(X)$ we have the full abelian subcategory $P(X)$ of *perverse sheaves*; see [5]. We will call a sheaf \mathcal{F} *shifted perverse* if $\mathcal{F}[n]$ is perverse for some $n \in \mathbb{Z}$.

3.3 The Frobenius and its action on sheaves

Given any variety X defined over \mathbb{F}_q we have the Frobenius morphism

$$\text{Fr}_q: X \rightarrow X,$$

which for affine $X \subset \mathbb{A}^m$ is given by $(x_1, \dots, x_m) \mapsto (x_1^q, \dots, x_m^q)$. The fixed points of $\text{Fr}_{q^n} := (\text{Fr}_q)^n$ are precisely $X(\mathbb{F}_{q^n})$, the points of X defined over \mathbb{F}_{q^n} .

Given any $\mathcal{F} \in D^b(X)$ we have an isomorphism (see [5, Chapter 5])

$$F_q^*: \text{Fr}_q^* \mathcal{F} \xrightarrow{\sim} \mathcal{F},$$

and obtain an induced action of $F_{q^n}^* := (F_q^*)^n$ on the stalk of \mathcal{F} at any point $x \in X(\mathbb{F}_{q^n})$. By considering the eigenvalues of the action of $F_{q^n}^*$ on the stalks of \mathcal{F} at all points $x \in X(\mathbb{F}_{q^n})$ for all $n \geq 1$, one defines the subcategory of *mixed sheaves* $D_m^b(X)$ as well as the full subcategories of sheaves of *weight* $\leq w$ and *weight* $\geq w$ for $w \in \mathbb{Z}$, which we denote $D_{\leq w}^b(X)$ and $D_{\geq w}^b(X)$ respectively; see [5, Chapter 5], [10] or the first chapter of [14]. An object is called *pure of weight* i if it lies in both $D_{\leq i}^b(X)$ and $D_{\geq i}^b(X)$.

Given any mixed sheaf \mathcal{F} on X , all eigenvalues $\alpha \in \mathbb{k}$ of Fr_q^* on $\mathbb{H}^*(\mathcal{F})$ are algebraic integers such that all complex numbers with the same minimal polynomial have the same complex norm, which by abuse of notation we denote by $|\alpha|$. As \mathcal{F} is assumed mixed, all such norms will be $q^{i/2}$ for some i . Let $\mathbb{H}_\alpha^*(\mathcal{F}) \subset \mathbb{H}^*(\mathcal{F})$ be the generalized eigenspace of α , and let

$$(1) \quad \mathbb{H}^{*,i}(\mathcal{F}) := \bigoplus_{|\alpha|=q^{i/2}} \mathbb{H}_\alpha^*(\mathcal{F}).$$

If X is proper and \mathcal{F} pure, then the decomposition (1) will agree with the cohomological grading of $\mathbb{H}^*(\mathcal{F})$ by the Riemann hypothesis for X ; that is, $\mathbb{H}^{*,i}(\mathcal{F}) = \mathbb{H}^i(\mathcal{F})$. Since we are not assuming that X is proper, this can fail even if \mathcal{F} is pure. For example, if $X = \mathbb{A}^1 \setminus \{0\} \cong G_1$ and $\mathcal{F} = \mathbb{k}_X$, then as in Example 3.1, we have $\mathbb{H}^{*,2}(\mathcal{F}) = \mathbb{H}^1(\mathcal{F})$.

Definition 3.2 The grading on $\mathbb{H}^*(\mathcal{F})$ where the elements of $\mathbb{H}^{*,i}(\mathcal{F})$ are defined to have degree i is called the *weight grading*.

Remark 1 The constant sheaf on X has a unique mixed structure for which the Frobenius acts trivially on all stalks, and its hypercohomology is the étale cohomology of X . The i^{th} graded component of $H^*(X; \mathbb{k})$ for the weight grading is $H^{*,i}(X; \mathbb{k})$. So our previous discussion was a reflection of some of the properties of the Frobenius action on the cohomology of algebraic varieties.

If $X = \text{Spec } \mathbb{F}_q$, then a perverse sheaf on X is the same as a finite-dimensional \mathbb{k} -vector space together with a continuous action of the absolute Galois group of \mathbb{F}_q . In particular we have the *Tate sheaf* $\underline{\mathbb{k}}(1)$, which under the above equivalence corresponds to \mathbb{k} , with action of F_q^* given by q^{-1} . Recall that we have fixed a square root $q^{1/2}$ of q in \mathbb{k} , allowing us to define the *half Tate sheaf* $\underline{\mathbb{k}}(\frac{1}{2})$, with F_q^* acting by $q^{-1/2}$.

Given any X with structure morphism $X \xrightarrow{a} \text{Spec } \mathbb{F}_q$ and any sheaf \mathcal{F} on X , we define

$$\mathcal{F}(\frac{m}{2}) := \mathcal{F} \otimes a^* \underline{\mathbb{k}}(\frac{1}{2})^{\otimes m}.$$

The following notation will prove useful:

$$\mathcal{F}\langle d \rangle = \mathcal{F}[d]\left(\frac{d}{2}\right).$$

If \mathcal{F} is pure of weight w , then $\mathcal{F}[d]$ is pure of weight $w + d$, and $\mathcal{F}(d/2)$ pure of weight $w - d$, so $\mathcal{F}\langle d \rangle$ is again pure of weight 0. The natural isomorphism $\mathbb{H}^*(F) \cong \mathbb{H}^*(F)\langle d \rangle$ has degree d for both the weight and cohomological gradings.

The most important fact about mixed sheaves for our purposes is that every mixed perverse sheaf \mathcal{F} on X admits a unique increasing filtration W , called the *weight filtration*, such that, for all i ,

$$\text{gr}_i^W \mathcal{F} := W_i \mathcal{F} / W_{i-1} \mathcal{F}$$

is pure of weight i . Any morphism of mixed sheaves preserves this filtration. The decomposition (1) comes from the weight filtration applied to the pushforward sheaf $p_* \mathcal{F}$ to a point.

In fact, after extension of scalars to the algebraic closure, the extensions in this filtration are the only way that mixed perverse sheaves can fail to be semisimple.

Theorem 3.3 (Gabber; [5, Théorème 5.3.8]) *If \mathcal{F} is a pure perverse sheaf on X then $\mathcal{F} \otimes \mathbb{F}$ is semisimple.*

3.4 The function–sheaf dictionary

The eigenvalues of the Frobenius on stalks are also valuable for analyzing the structure of a given perverse sheaf. To any mixed perverse sheaf \mathcal{F} (or more generally, any mixed sheaf) one may associate a function on $X(\mathbb{F}_{q^n})$ for each n given by the supertrace of the Frobenius on the stalks of the cohomology sheaves at those points:

$$[\mathcal{F}]_n: X(\mathbb{F}_{q^n}) \rightarrow \mathbb{k}, \quad x \mapsto \text{Tr}(F_{q^n}^*, \mathcal{F}_x) := \sum (-1)^j \text{Tr}(F_{q^n}^*, \mathcal{H}^j(\mathcal{F}_x)).$$

Proposition 3.4 *These functions give a map from the Grothendieck group of the category of mixed perverse sheaves to the abelian group of functions on $X(\mathbb{F}_q^n)$ for each n , and these maps are jointly injective. That is, if \mathcal{F} and \mathcal{G} are semisimple and $[\mathcal{F}]_n = [\mathcal{G}]_n$ for all n , then \mathcal{F} and \mathcal{G} are isomorphic.*

Proof The fact that these functions give a map of Grothendieck groups is just that all maps in the long exact sequence must respect the action of the Frobenius, so the supertrace is additive under extensions. The proof that this map is injective may be found in [15, Théorème 1.1.2]; see also [14, Theorem 12.1]. \square

This reduces the calculation of the constituents of a weight filtration to a problem of computing $[\mathcal{F}]_n$ for simple perverse sheaves, followed by linear algebra. Indeed, suppose that $\mathcal{F}, \mathcal{G} \in D_m^b(X)$ are such that $[\mathcal{F}]_n$ and $[\mathcal{G}]_n$ agree for all n , with \mathcal{G} semisimple. As $[\mathcal{F}]_n = \sum [\text{gr}_i^W \mathcal{F}]_n$ for all n , we conclude that $\text{gr}_i^W \mathcal{F}$ is isomorphic to the largest direct summand of \mathcal{G} of weight i .

3.5 The chromatographic complex

We want to explain how to move between the weight filtration and a complex, which we term *the chromatographic complex*, composed of its pure constituents. For background, the reader is referred to [8, Section 1.4] and [5, Section 3.1].

Let \mathcal{A} be an abelian category with enough injectives and let $D^+(\mathcal{A})$ denote its bounded below derived category. We may also consider the filtered derived category $DF^+(\mathcal{A})$ whose objects consist of $K \in D^+(\mathcal{A})$ together with a finite increasing filtration

$$\cdots \subset W_{i-1}K \subset W_iK \subset W_{i+1}K \subset \cdots,$$

where finite means that $W_iK = 0$ for $i \ll 0$ and $W_iK = W_{i+1}K$ for $i \gg 0$.

For all p we define

$$\text{gr}_p^W K := W^p K / W^{p-1} K.$$

More generally, for $q \leq p$, let

$$(W^p / W^q)(K) := W^p K / W^q K.$$

For all p we have a distinguished triangle

$$\text{gr}_p^W K \rightarrow (W^{p+1} / W^{p-1})(K) \rightarrow \text{gr}_{p+1}^W K \xrightarrow{[1]}$$

and in particular a “boundary” morphism $\text{gr}_W^{p+1} \rightarrow \text{gr}_W^p K[1]$. Shifting, we obtain a sequence

$$(2) \quad \cdots \rightarrow \text{gr}_{p+1}^W K[-(p+1)] \rightarrow \text{gr}_p^W K[-p] \rightarrow \text{gr}_{p-1}^W K[-(p-1)] \rightarrow \cdots.$$

Lemma 3.5 *The morphisms in (2) define a complex.*

Proof After completing the (commuting) triangle

$$\begin{array}{ccc} & (W^{p+1}/W^{p-1})(K) & \\ & \nearrow & \searrow \\ \mathrm{gr}_p^W K & \xrightarrow{\quad\quad\quad} & (W^{p+2}/W^{p-1})(K) \end{array}$$

to an octahedron, one sees that the morphism

$$\mathrm{gr}_{p+2}^W K \rightarrow \mathrm{gr}_{p+1}^W K[1] \rightarrow \mathrm{gr}_p^W K[2]$$

may be factored as

$$\mathrm{gr}_{p+2}^W \rightarrow W^{p+1}/W^{p-1}(K) \rightarrow \mathrm{gr}_{p+1}^W K[1] \rightarrow \mathrm{gr}_p^W K[2].$$

However, the second two morphisms form part of a distinguished triangle, and so their composition is zero. □

Given any left exact functor $T: \mathcal{A} \rightarrow \mathcal{B}$ between abelian categories we can consider the hypercohomology objects $R^i T(K) \in \mathcal{B}$, obtained by applying T to an injective resolution of K . One has a spectral sequence

$$(3) \quad E_1^{p,q} = R^{p+q} T(\mathrm{gr}_{-p}^W K) \Rightarrow R^{p+q} T(K)$$

(see [18, Theorem 2.6] or [8, Section 1.4.5]), and a diagram chase shows that the first differential of this spectral sequence (ie the differential on the E_1 page) is the same as the differential obtained by applying $R^q T(-)$ to the complex (2).

We now apply these ideas to $D_m^b(X)$, where X and $D_m^b(X)$ are as in Section 3.3.

Lemma 3.6 *Any $\mathcal{G} \in D_m^b(X)$ admits a “filtration” $\dots \rightarrow \mathcal{G}_{\leq i} \rightarrow \mathcal{G}_{\leq i+1} \rightarrow \dots$ such that:*

(1) *If we define $\mathrm{gr}_i(X)$ via the distinguished triangle*

$$\mathcal{G}_{\leq i-1} \rightarrow \mathcal{G}_{\leq i} \rightarrow \mathrm{gr}_i(\mathcal{G}) \xrightarrow{[1]},$$

then $\mathrm{gr}_i(\mathcal{G})$ is pure of weight i .

(2) $\mathrm{gr}_i(\mathcal{G}) = 0$ for $|i| \gg 0$.

We will refer to any sequence of maps satisfying the conditions of the lemma as a *weight filtration* on \mathcal{G} . As the choice of article emphasizes, this is *not* unique. (For example, the reader may convince themselves easily that the zero object admits many nonequivalent weight filtrations.)

Proof It is enough to show that for any $\mathcal{G} \in D_m^b(X)$ there exists a distinguished triangle

$$\mathcal{G}_{\leq 0} \rightarrow \mathcal{G} \rightarrow \mathcal{G}_{> 0} \xrightarrow{[1]}$$

with $\mathcal{G}_{\leq 0}$ (resp. $\mathcal{G}_{> 0}$) of weight ≤ 0 (resp. > 0). If \mathcal{G} is perverse then the statement is an immediate consequence of the existence of weight filtrations on perverse sheaves [5, Théorème 5.3.5].

By induction on the perverse filtration it is enough to show the following claim: if

$$\mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{K} \xrightarrow{[1]}$$

is a distinguished triangle of sheaves, and there exist distinguished triangles

$$\mathcal{F}_{\leq 0} \rightarrow \mathcal{F} \rightarrow \mathcal{F}_{> 0} \xrightarrow{[1]} \quad \text{and} \quad \mathcal{K}_{\leq 0} \rightarrow \mathcal{K} \rightarrow \mathcal{K}_{> 0} \xrightarrow{[1]}$$

with $\mathcal{F}_{\leq 0}$ and $\mathcal{K}_{\leq 0}$ (resp. $\mathcal{F}_{> 0}$ and $\mathcal{K}_{> 0}$) of weights ≤ 0 (resp. > 0), then there exists a filtration

$$\mathcal{G}_{\leq 0} \rightarrow \mathcal{G} \rightarrow \mathcal{G}_{> 0} \xrightarrow{[1]}$$

satisfying the same conditions. For the rest of the proof the following notation will be useful. Given a commutative triangle

$$\begin{array}{ccc} & \mathcal{B} & \\ \nearrow & & \searrow \\ \mathcal{A} & \longrightarrow & \mathcal{C} \end{array}$$

we denote by $O(\mathcal{A}, \mathcal{B}, \mathcal{C})$ the corresponding octahedron (the maps will be clear from the context).

By considering $O(\mathcal{G}, \mathcal{K}, \mathcal{K}_{> 0})$ we deduce the existence of distinguished triangles

$$\mathcal{A} \rightarrow \mathcal{G} \rightarrow \mathcal{K}_{> 0} \xrightarrow{[1]} \quad \text{and} \quad \mathcal{F} \rightarrow \mathcal{A} \rightarrow \mathcal{K}_{\leq 0} \xrightarrow{[1]} .$$

By considering $O(\mathcal{F}_{\leq 0}, \mathcal{F}, \mathcal{A})$ we deduce the existence of distinguished triangles

$$\mathcal{F}_{\leq 0} \rightarrow \mathcal{A} \rightarrow \mathcal{B} \xrightarrow{[1]} \quad \text{and} \quad \mathcal{F}_{> 0} \rightarrow \mathcal{B} \rightarrow \mathcal{K}_{\leq 0} \xrightarrow{[1]} .$$

Because $\text{Hom}(\mathcal{K}_{\leq 0}, \mathcal{F}_{> 0}[1]) = 0$ by [5, Proposition 5.1.15(ii)], we deduce that the “connecting” map $\mathcal{K}_{\leq 0} \rightarrow \mathcal{F}_{> 0}[1]$ in the second triangle is zero, and hence that $\mathcal{B} \cong \mathcal{F}_{> 0} \oplus \mathcal{K}_{\leq 0}$. (This decomposition is not canonical; we fix one.) By considering $O(\mathcal{A}, \mathcal{B}, \mathcal{F}_{> 0})$ (the map $\mathcal{B} = \mathcal{F}_{> 0} \oplus \mathcal{K}_{\leq 0} \rightarrow \mathcal{F}_{> 0}$ is the projection) we deduce the existence of distinguished triangles

$$\mathcal{C} \rightarrow \mathcal{A} \rightarrow \mathcal{F}_{> 0} \xrightarrow{[1]} \quad \text{and} \quad \mathcal{F}_{\leq 0} \rightarrow \mathcal{C} \rightarrow \mathcal{K}_{\leq 0} \xrightarrow{[1]} .$$

In particular, \mathcal{C} has weight ≤ 0 . Finally, by considering $O(\mathcal{C}, \mathcal{A}, \mathcal{G})$ we deduce the existence of distinguished triangles

$$\mathcal{C} \rightarrow \mathcal{G} \rightarrow \mathcal{D} \xrightarrow{[1]} \quad \text{and} \quad \mathcal{F}_{>0} \rightarrow \mathcal{D} \rightarrow \mathcal{K}_{>0} \xrightarrow{[1]} .$$

In particular, \mathcal{D} has weight > 0 . It follows that we can take $\mathcal{G}_{\leq 0} := \mathcal{C}$ and $\mathcal{G}_{>0} := \mathcal{D}$. \square

Applying the above considerations to \mathcal{F} together with its weight filtration we obtain:

Definition 3.7 The *local chromatographic complex* of a mixed sheaf $\mathcal{F} \in D_m^b(X)$ is the complex of objects in $D_m^b(X)$ given by

$$\cdots \rightarrow \text{gr}_{p+1}^W \mathcal{F}[-(p+1)] \rightarrow \text{gr}_p^W \mathcal{F}[-p] \rightarrow \text{gr}_{p-1}^W \mathcal{F}[-(p-1)] \rightarrow \cdots .$$

Applying $T = \mathbb{H}^*(-)$ we obtain the *global chromatographic complex*,

$$\cdots \rightarrow \mathbb{H}^*(\text{gr}_{i+1}^W \mathcal{F}[-(i+1)]) \rightarrow \mathbb{H}^*(\text{gr}_i^W \mathcal{F}[-i]) \rightarrow \mathbb{H}^*(\text{gr}_{i-1}^W \mathcal{F}[-(i-1)]) \rightarrow \cdots .$$

The spectral sequence (3) with $T = \mathbb{H}^*(-)$ is the *chromatographic spectral sequence*.

Note that the global chromatographic complex sends $\langle d \rangle$ to simultaneous grading shift on terms of the complex, and the Tate twist $(d/2)$ to homological shift of the complex. Unfortunately, this definition is not entirely an invariant of the object \mathcal{G} , but the dependence on choice of filtration is not very strong.

Proposition 3.8 The chromatographic complexes associated to two different weight filtrations on a single object $\mathcal{G} \in D^b(X)$ are homotopy equivalent, after extending scalars to \mathbb{F} .

In particular, this shows that all pages of the chromatographic spectral sequence after the first are independent of the choice of filtration.

Proof We note that if \mathcal{G} is quasi-isomorphic to a complex $\cdots \rightarrow \mathcal{F}_i \rightarrow \cdots$, then we obtain a natural bicomplex by writing the chromatographic complexes of \mathcal{F}_i vertically, and then the maps induced by the original differentials horizontally. By Gabber’s theorem, we note that after passing to \mathbb{F} every term in this bicomplex is semisimple, and the horizontal maps go between objects pure of the same degree, and thus split.

Now assume perverse sheaves \mathcal{F}'_i form another complex isomorphic in the derived category to \mathcal{G} . For simplicity, we may assume there is a quasi-isomorphism $\phi_i: \mathcal{F}_i \rightarrow \mathcal{F}'_i$ between these complexes. This induces a map $\phi^\#$ between our bicomplexes, which is an isomorphism after taking horizontal cohomology, since this will give us the chromatographic complexes of the perverse cohomology of \mathcal{G} .

Consider the kernel of $\phi^\#$. This is itself a bicomplex, and each of its rows has trivial cohomology, and is split. Thus, each row is homotopic to 0. Furthermore, we can choose these homotopies so that they commute with the vertical differentials, and thus when applied to the total complex of the kernel, they show that this total complex is nullhomotopic.

We now use the result that any surjective chain map whose kernel is homotopic to the zero complex and is a summand of the chain complex *with the differentials forgotten* is a homotopy equivalence (this is a consequence of Gaussian elimination). Thus, the chromatographic complex from the \mathcal{F}_i is homotopy equivalent to the total complex of the image of $\phi^\#$, and the dual result applied to the inclusion of the image shows that the chromatographic complex for \mathcal{F}'_i is also equivalent to this image. \square

Proposition 3.9 *The global chromatographic complex is preserved (up to homotopy) by proper pushforward.*

Proof Proper pushforward preserves purity, and thus sends weight filtrations to weight filtrations. Furthermore, pushforward always preserves hypercohomology. \square

Corollary 3.10 *If we let $E_*^{*,*}$ be the chromatographic spectral sequence, then all differentials preserve the weight grading on hypercohomology. Furthermore, we have:*

- $E_1^{i,j} = \mathbb{H}^{i+j}(\text{gr}_{-j}^W \mathcal{F})$ is the global chromatographic complex.
- E_2 is the cohomology of the global chromatographic complex.
- The chromatographic spectral sequence converges to the hypercohomology $\mathbb{H}^{i+j}(\mathcal{F})$.

Remark 2 It seems likely that it is possible to interpret the results of this section in terms of “weight structures”, introduced by Bondarko [6] and Pauksello [20]. In particular, Bondarko shows the existence of a functor from a derived category equipped with a suitable weight structure, to the homotopy category of pure complexes in a very general framework.

3.6 The equivariant derived category

We have thus far discussed the theory of perverse sheaves on schemes, but we will require a generalization of schemes that includes the quotient of a scheme X by the action of an algebraic group G , which can be understood as G -equivariant geometry on X .

This quotient can be understood as a stack, but the theory of perverse sheaves on stacks is not straightforward, and it proved more suitable to give a treatment of the equivariant derived category similar to that of Bernstein and Lunts [4], but with an eye to working over characteristic p with the action of the Frobenius (that is, “in the mixed setting”). We have done this in a separate note [25].

The result is the *bounded below equivariant derived category* $D_G^+(X)$ and its subcategory $D_G^b(X)$ of bounded complexes for a variety X acted on by an affine algebraic group G . The resulting formalism is essentially identical to that of Bernstein and Lunts. We now summarize the essential points.

We have a forgetful functor

$$\text{For}: D_G^+(X) \rightarrow D^+(X)$$

which preserves the subcategories of bounded complexes and, given any $\mathcal{F} \in D_G^+(X)$, the cohomology sheaves of $\text{For}(\mathcal{F})$ are locally constant along the G -orbits on X .

Given an equivariant map $f: X \rightarrow Y$ of G -varieties we have functors

$$f_*, f_!: D_G^+(X) \rightarrow D_G^+(Y) \quad \text{and} \quad f^*, f^!: D_G^+(Y) \rightarrow D_G^+(X)$$

for equivariant maps $f: X \rightarrow Y$ of G -varieties. These functors commute with the forgetful functor.

If $H \subset G$ is a closed subgroup and X is a G -space, we have an adjoint pair $(\text{res}_H^G, \text{ind}_H^G)$ of restriction and induction functors

$$\text{res}_H^G: D_G^+(X) \rightarrow D_H^+(X) \quad \text{and} \quad \text{ind}_H^G: D_H^+(X) \rightarrow D_G^+(X).$$

These functors preserve the subcategories of bounded complexes, and one has an isomorphism $\text{res}_{\{1\}}^G \cong \text{For}$.

More generally, given a map $\phi: H \rightarrow G$, a G -variety X , an H -variety Y and a ϕ -equivariant map $m: X \rightarrow Y$, we have an adjoint pair $({}^G_H m^*, {}^G_H m_*)$ of functors

$${}^G_H m^*: D_H^+(Y) \rightarrow D_G^+(X) \quad \text{and} \quad {}^G_H m_*: D_G^+(X) \rightarrow D_H^+(Y).$$

As a special case, we have ${}^G_H \text{id}^* \cong \text{res}_H^G$ and ${}^G_H \text{id}_* \cong \text{ind}_H^G$. The functor ${}^G_H m^*$ preserves the subcategory of bounded complexes, but this is not true in general for ${}^G_H m_*$. In fact, this is the reason that we are forced to consider complexes of sheaves which are not bounded above.

If $G = G_1 \times G_2$ and G_1 acts freely on X with quotient X/G_1 , one has the *quotient equivalence*

$$(4) \quad D_G^+(X) \cong D_{G_2}^+(X/G_1),$$

which restricts to an equivalence between the subcategories of bounded complexes. If we let $\phi: G_1 \times G_2 \rightarrow G_2$ denote the projection, then the quotient map $X \rightarrow X/G_1$ is ϕ -equivariant and the above equivalence is realized by ${}_{G_1 \times G_2} G_2 m^*$ and ${}_{G_1 \times G_2} G_2 m_*$.

Using the forgetful functor $\text{For}: D_G^+(X) \rightarrow D^+(X)$ many notions carry over immediately. For example, we call an object \mathcal{F} in $D_G^+(X)$ *perverse* if and only if $\text{For } \mathcal{F}$ is perverse.

Moreover, if X is defined over \mathbb{F}_q , then we can also incorporate the action of the Frobenius. In particular, perverse objects in $D_G^+(X)$ still have weight filtrations, which are preserved by the restriction functor and we can extend Proposition 3.4 to the equivariant setting.


4 Description of the invariant

Equipped with these geometric tools, we continue the construction of our invariant.

4.1 The sheaf associated to a diagram

In this subsection we describe the sheaf \mathcal{F}_D on X_D .

We first discuss the case of a single (i, j) -crossing:



As we have seen, $X_D = G_{i+j}$. Consider the big Bruhat cell

$$(5) \quad U := \{g \in G_{i+j} \mid V_i \cap gV_j = 0\},$$

and let $k: U \hookrightarrow G_{i+j}$ denote its inclusion. As U is an orbit under $P_{i,j} \times P_{j,i}$ it is certainly G_D -invariant. We now define $\mathcal{F}_v = \mathcal{F}_D \in D_{G_D}(X_D)$ as follows:

$$\begin{array}{ccc}
 \begin{array}{c} i \\ \diagdown \\ \diagup \\ j \end{array} & \longmapsto & k_* \underline{\mathbb{K}}_U \langle ij \rangle, & \begin{array}{c} i \\ \diagdown \\ \diagup \\ j \end{array} & \longmapsto & k_! \underline{\mathbb{K}}_U \langle ij \rangle.
 \end{array}$$

As U is the complement of a divisor in G_{i+j} , both these sheaves are shifted perverse.

We now consider the case of a general diagram D of an oriented colored tangle. After forgetting equivariance, \mathcal{F}_D is simply the exterior product of the above sheaves associated to each crossing. To take care of the equivariant structure we need to proceed a little more carefully.

Let D be the diagram of an oriented colored tangle and Γ its underlying graph. Let D' be the diagram obtained from D by cutting each strand connecting two vertices in Γ (so that D' is a disjoint union of (i, j) -crossings). Let Γ' be the graph corresponding to D' . Obviously we have $X_D = X_{D'}$. Note also that for every e with two vertices in Γ , we have two edges, which we denote by e_1 and e_2 , in Γ' . We have a natural map $G_D \rightarrow G_{D'}$, which is the identity on factors corresponding to external edges, and is the diagonal $G_e \rightarrow G_{e_1} \times G_{e_2}$ on the remaining factors.

We define

$$\mathcal{F}_D := \text{res}_{G'}^G \left(\bigotimes_{v \in \mathcal{V}(D')} \mathcal{F}_v \right) \in D_{G_D}^b(X_D).$$

Of course, this sheaf depends on the link diagram used; different diagrams correspond to sheaves on different spaces. Instead, we will study the hypercohomology of these sheaves, and the corresponding chromatographic spectral sequence.

Definition 4.1 We let $\mathcal{A}_i(L)$ denote the i^{th} page of the chromatographic spectral sequence (as given by Definition 3.7) for \mathcal{F}_D . This is triply graded, where by convention subquotients of $\mathbb{H}^{j-\ell; j-k}(\text{gr}_\ell^W \mathcal{F}_D)$ lie in $\mathcal{A}_i^{j;k;\ell}(L)$.

Remark 3 These grading conventions may seem strange, but they are an attempt to match those already in use in the field. These conventions are almost those of [17], though we will not match perfectly since we have different grading shifts in our definition of the complex for a single crossing. We hope the reader finds these choices defensible on grounds of geometric naturality. This simply changes the shift we must apply to our invariant to assure it is a true knot invariant.

It is these spaces for $i > 1$ which we intend to show are knot invariants (up to shift).

4.2 Braids and sheaves on groups

As we mentioned in Section 1, in the special case of a braid β , there is a different perspective on this construction.

Let β be the diagram of a colored braid on n strands with labels $\mathbf{n} = (i_1, i_2, \dots, i_n)$ and underlying labeled graph Γ . Let $N = \sum_{j=1}^n i_j$ denote the colored braid index. We assume our braid is in generic position, so reading from start to finish, we fix an order on the vertices v_1, v_2, \dots, v_p of Γ . This corresponds to an expression for β in the standard generators of the braid group.

In the previous section we described how to associate to β a group G_β and a G_β -variety X_β . We can decompose G_β as

$$G_\beta = G_\beta^+ \times G_\beta^l \times G_\beta^-,$$

where G_β^+ , G_β^l and G_β^- denote the factors of G_β corresponding to incoming, interior and outgoing edges of Γ respectively.

In what follows we will describe an action of $G_\beta^+ \times G_\beta^-$ on G_N and a map

$$m: X_\beta \rightarrow G_N$$

equivariant with respect to the natural projection $\phi: G_\beta \rightarrow G_\beta^+ \times G_\beta^-$. We will study our sheaf \mathcal{F}_β by considering its equivariant pushforward under this map.

First we describe an embedding $\alpha_v: G_v \rightarrow G_N$ corresponding to each vertex $v \in \Gamma$. Let us fix a basis e_1, \dots, e_N of V_N and let W_1, W_2, \dots, W_n be vector spaces (again with fixed bases) of dimensions i_1, i_2, \dots, i_n respectively.

Definition 4.2 Given any permutation $w \in S_n$, we let

$$h_w: W = \bigoplus_{j=1}^n W_j \xrightarrow{\sim} V$$

be the isomorphism defined by mapping the basis vectors of $W_{w^{-1}(1)}$ to the first $w^{-1}(1)$ basis vectors of V in their natural order, the basis vectors of $W_{w^{-1}(2)}$ to the next $w^{-1}(2)$ basis vectors, etc.

For any braid β , we have an induced permutation, and by abuse of notation, we let h_β be the map corresponding to this permutation.

In the obvious basis, this map is a permutation matrix. The corresponding permutation is a shortest coset representative for the Young subgroup preserving the partition of $[1, N]$ of sizes i_1, \dots, i_n , corresponding to the ‘‘cabling’’ of the permutation w .

Now choose a vertex v in Γ and let e' and e'' denote the two incoming edges which are in the strands connected to the incoming vertices labeled j' and j'' respectively, so $i_{j'}$ and $i_{j''}$ are the labels on e' and e'' . Because we have ordered the vertices of Γ , we may factor β into braids $\alpha_v \cdot \beta_v \cdot \omega_v$, with β_v consisting of a simple crossing corresponding to v . The procedure described in the previous paragraph yields an embedding

$$W_{j'} \oplus W_{j''} \hookrightarrow W \xrightarrow{h_{\alpha_v}} V_N.$$

This induces an embedding

$$\iota_v: G_v \hookrightarrow G_N.$$

We let braids on n strands act on sequences of n elements on the right by the usual association of a permutation to each braid. We may then identify

$$G_\beta^+ \cong G_n, \quad G_\beta^- \cong G_{n\beta},$$

and therefore obtain an action of $G_\beta^+ \times G_\beta^-$ on G_N by left and right multiplication. We let $P_\beta^+ = P_n$ and $P_\beta^- = P_{n\beta}$. We denote by $\phi: G_\beta \rightarrow P_\beta^+ \times P_\beta^-$ the composition of the natural projection with the inclusion $G_\beta^\pm \hookrightarrow P_\beta^\pm$.

Consider the map

$$m: X_\beta \rightarrow G_N, \quad (g_{v_1}, \dots, g_{v_p}) \mapsto \iota_{v_1}(g_{v_1})\iota_{v_2}(g_{v_2}) \cdots \iota_{v_p}(g_{v_p}).$$

It is easy to see that this map is equivariant with respect to ϕ .

Definition 4.3 Let $\Phi_\beta = \frac{P_\beta^+ \times P_\beta^-}{G_\beta} m_* \mathcal{F}_\beta$.

This definition is useful, since it is compatible with braid multiplication. We have a diagram of equivariant maps of spaces:

$$\begin{array}{ccccc} G_N & \xleftarrow{\pi_1} & & & \\ & & G_N \times G_N & \xrightarrow{\mu} & G_N \\ G_N & \xleftarrow{\pi_2} & & & \end{array}$$

As usual, this diagram can be used to construct the functor of sheaf convolution:

$$\begin{aligned} - \star - : D_{P_n \times P_{n\beta}}^b(G_N) \times D_{P_{n\beta} \times P_{n\beta\beta'}}^b(G_N) &\rightarrow D_{P_n \times P_{n\beta\beta'}}^b(G_N), \\ \mathcal{F}_1 \star \mathcal{F}_2 &\cong \frac{P_n \times P_{n\beta\beta'}}{P_n \times P_{n\beta} \times P_{n\beta\beta'}} \mu_* \left(\text{res}_{\frac{P_n \times P_{n\beta}^2 \times P_{n\beta\beta'}}{P_n \times P_{n\beta} \times P_{n\beta\beta'}}}(\mathcal{F}_1 \boxtimes \mathcal{F}_2) \right). \end{aligned}$$

Theorem 4.4 We have a canonical isomorphism $\Phi_\beta \star \Phi_{\beta'} \cong \Phi_{\beta\beta'}$.

We should note that here we are simply claiming that this holds for the composition of diagrams. We will prove in Sections 8 and 9 that the sheaf we associate to a braid doesn't depend on the choice of presentation.

Proof Immediate from the definition of Φ . □

As G_β^l acts freely on X_β , we may factor m as

$$X_\beta \rightarrow X_\beta / G_\beta^l \rightarrow G_N.$$

One may verify that the second map is the composition of an affine bundle along which \mathcal{F}_β is smooth, and a proper map. It follows that

$$\frac{P_\beta^+ \times P_\beta^-}{G_\beta} m_*$$

preserves the weight filtration on \mathcal{F}_β . Hence the chromatographic spectral sequences for \mathcal{F}_β and Φ_β are isomorphic.

Note that if β is closable, then $n\beta = n$, and P_β^\pm have the same image in the group. Thus these subgroups are canonically isomorphic. Let $(P_\beta)_\Delta \subset P_\beta^+ \times P_\beta^-$ be the diagonal and let $\hat{\beta}$ be the colored link diagram given by the closure of β .

Theorem 4.5 We have a canonical isomorphism between

- the chromatographic spectral sequence of $\mathcal{F}_{\hat{\beta}}$ as a $G_{\hat{\beta}}$ -sheaf, and
- the chromatographic spectral sequence of Φ_β as a $(P_\beta)_\Delta$ -sheaf.

Proof Since P_* and G_* are homotopy equivalent, the functor $\text{res}_{G_*}^{P_*}$ is fully faithful, so we may work with their restrictions instead. We have already observed that the weight filtration on Φ_β and the pushforward of the weight filtration on \mathcal{F}_β agree. Thus the equivariant chromatographic spectral sequences of

$$\text{res}_{\phi^{-1}(H)}^{G_\beta} \mathcal{F}_\beta \quad \text{and} \quad \text{res}_H^{G_\beta^+ \times G_\beta^-} \Phi_\beta$$

are canonically isomorphic for any subgroup $H \subset G_\beta^+ \times G_\beta^-$.

On the other hand, we have a canonical identification $G_{\hat{\beta}} \cong \phi^{-1}((G_\beta)_\Delta)$, and $X_\beta = X_{\hat{\beta}}$, with

$$\mathcal{F}_{\hat{\beta}} = \text{res}_{G_{\hat{\beta}}}^{G_\beta} \mathcal{F}_\beta.$$


The result follows. □

5 Analyzing an (m, n) -crossing

5.1 Preliminary details

In this section we work out all the details for an (m, n) -crossing. This will be of use in expressing the invariant in terms of bimodules.

We consider an (m, n) -crossing. Its underlying graph is



and the variety in question is G_{m+n} acted on by $P_{m,n} \times P_{n,m}$ by left and right multiplication: $(p, q) \cdot g = pgq^{-1}$ for $g \in G_{m+n}$ and $(p, q) \in P_{m,n} \times P_{n,m}$. The orbits under this action are

$$\mathcal{O}_i = \{g \in G_{m+n} \mid \dim V_m \cap gV_n = i\},$$

for $0 \leq i \leq \min(n, m)$. Clearly $\mathcal{O}_j \subset \overline{\mathcal{O}}_i$ if and only if $j > i$. For all $0 \leq i \leq \min(n, m)$ we denote the inclusion of the orbit \mathcal{O}_i by $f_i: \mathcal{O}_i \hookrightarrow G_{n+m}$.

For each orbit \mathcal{O}_i we have the corresponding intersection cohomology complex. It will prove natural to normalize them by requiring

$$\mathrm{IC}(\overline{\mathcal{O}}_i)|_{\mathcal{O}_i} \cong \mathbb{k}_{\mathcal{O}_i}\langle nm - i^2 \rangle.$$

Under this normalization each $\mathrm{IC}(\overline{\mathcal{O}}_i)$ is pure of weight 0.

We first describe resolutions for the closures $\overline{\mathcal{O}}_i \subset G_{m+n}$. Consider the variety

$$\tilde{\mathcal{O}}_i = \{(W, g) \in \mathrm{Gr}_i^m \times G_{m+n} \mid W \subset V_m \cap gV_n\}.$$

We have an action of $P_{m,n} \times P_{n,m}$ on $\tilde{\mathcal{O}}_i$ given by $(p, q) \cdot (W, g) = (pW, pgq^{-1})$. The second projection induces an equivariant map:

$$\pi_i: \tilde{\mathcal{O}}_i \rightarrow \overline{\mathcal{O}}_i.$$

Proposition 5.1 *This is a small resolution of singularities.*

Proof The morphism π_i is patently an isomorphism over \mathcal{O}_i . Since \mathcal{O}_i is exactly the subset of G_{n+m} where the induced map $V_n \rightarrow V/V_m$ has rank $n - i$, we have that \mathcal{O}_i has the same codimension in G_{m+n} as the space of rank $n - i$ matrices in G_n , which is i^2 . Hence, for $j < i$, \mathcal{O}_i is of codimension $i^2 - j^2$ in $\overline{\mathcal{O}}_j$. Over any $x \in \mathcal{O}_j$ the fiber is the Grassmannian Gr_i^j . Thus

$$2 \dim \pi_i^{-1}(x) = 2i(j - i) < (j + i)(j - i) = \mathrm{codim}_{\overline{\mathcal{O}}_j} \mathcal{O}_j. \quad \square$$

Corollary 5.2 $\mathrm{IC}(\overline{\mathcal{O}}_i) \cong \pi_{i*} \mathbb{k}_{\tilde{\mathcal{O}}_i}\langle nm - i^2 \rangle.$

Proof Proposition 5.1 implies that $\pi_{i*} \mathbb{k}_{\tilde{\mathcal{O}}_i}$ is a shift and twist of $\mathrm{IC}(\overline{\mathcal{O}}_i)$, since pushforward by a small resolution sends the constant sheaf to a shift of the intersection cohomology sheaf on the target. The restriction of $\pi_{i*} \mathbb{k}_{\tilde{\mathcal{O}}_i}\langle nm - i^2 \rangle$ to \mathcal{O}_i is isomorphic to $\mathbb{k}_{\mathcal{O}_i}\langle nm - i^2 \rangle$, which is our choice of normalization. \square

Given sheaves $\mathcal{F}, \mathcal{G} \in D_G^b(X)$ let us write

$$\mathrm{Hom}^\bullet(\mathcal{F}, \mathcal{G}) := \bigoplus_m \mathrm{Hom}(\mathcal{F}, \mathcal{G}[m]).$$

This is a graded vector space.

Proposition 5.3 *In $D_{P_{m,n} \times P_{n,m}}^b(G)$ we have an isomorphism*

$$\mathrm{Hom}^\bullet(\mathrm{IC}(\mathcal{O}_i), \mathrm{IC}(\mathcal{O}_{i'})) \cong \bigoplus_j \mathrm{Hom}^\bullet(f_j^* \mathrm{IC}(\mathcal{O}_i), f_j^! \mathrm{IC}(\mathcal{O}_{i'})).$$

Proof For flag varieties this is [3, Theorem 3.4.1]. One may reduce to this situation using the quotient equivalence. \square

The space $\text{Hom}^\bullet(\text{IC}(\mathcal{O}_i), \text{IC}(\mathcal{O}_{i'}))$ itself has a weight grading, when thought of as sections of the sheaf-Hom from $\mathcal{H}om^\bullet(\text{IC}(\mathcal{O}_i), \text{IC}(\mathcal{O}_{i'}))$, which has a natural mixed structure. The decomposition of Proposition 5.3 is compatible with the Frobenius structure, and so the purity of the cohomology of \mathcal{O}_i (which is an affine bundle over a partial flag variety) and the pointwise purity of $\text{IC}(\mathcal{O}_i)$ shows that the weight grading of $\text{Hom}^\bullet(\text{IC}(\mathcal{O}_i), \text{IC}(\mathcal{O}_{i'}))$ agrees with the cohomological grading.

This shows that:

Proposition 5.4 *In the mixed equivariant derived category $D_{m, P_{m,n} \times P_{n,m}}^b(G)$, there are no higher Exts between $\text{IC}(\mathcal{O}_i)$ and $\text{IC}(\mathcal{O}_{i'})\langle d \rangle$.*

Proof By the purity discussed above, all of the eigenvalues of Frobenius on the space $\text{Ext}^i(\text{IC}(\mathcal{O}_i), \text{IC}(\mathcal{O}_{i'})\langle d \rangle)$ have complex norm $p^{i/2}$, so they are not 1. Thus, there are no invariants of Frobenius in this space. \square

This shows immediately that:

Corollary 5.5 *Any mixed $(P_{m,n} \times P_{n,m})$ -equivariant sheaf \mathcal{F} on G is the iterated cone of its local chromatographic complex (in any dg-refinement). In particular, \mathcal{F} is indecomposable if and only if the same is true of its local chromatographic complex.*

5.2 Calculating the weight filtration

Our aim in this section is to calculate the weight filtration on the sheaves associated to positive and negative crossings. We set

$$\begin{aligned} [n]_q &= 1 + q + \dots + q^{n-1}, \\ [n]_q! &= [n]_q [n-1]_q \dots [1]_q, \\ \begin{bmatrix} j \\ i \end{bmatrix}_q &= \frac{[j]_q!}{[j-i]_q! [i]_q!}. \end{aligned}$$

In order to understand the constituents via the function–sheaf correspondence discussed in Section 3.4, we must calculate the trace of the Frobenius on the stalks of $\text{IC}(\overline{\mathcal{O}}_i)$. Base change combined with the Grothendieck–Lefschetz fixed point formula yields:

Corollary 5.6 *If $j > i$ and $x \in \mathcal{O}_j(\mathbb{F}_{q^a})$, we have*

$$\text{Tr}(F_{q^a}^*, (\pi_{i*} \mathbb{K}_{\overline{\mathcal{O}}_i})_x) = \# \text{Gr}_i^j(\mathbb{F}_{q^a}) = \begin{bmatrix} j \\ i \end{bmatrix}_{q^a}.$$

In the following proposition W denotes the weight filtration.

Proposition 5.7 *One has isomorphisms*

$$\mathrm{gr}_{-i}^W j_! \mathbb{k}_{\mathcal{O}_0} \langle nm \rangle \cong \mathrm{IC}(\bar{\mathcal{O}}_i)(\tfrac{i}{2}), \quad \mathrm{gr}_i^W j_* \mathbb{k}_{\mathcal{O}_0} \langle nm \rangle \cong \mathrm{IC}(\bar{\mathcal{O}}_i)(-\tfrac{i}{2}).$$

Proof Because taking weight filtrations commutes with forgetting equivariance, it is enough to handle the nonequivariant case. Note also that $\mathrm{IC}(\mathcal{O}_i)(i/2)$ is pure of weight $-i$. Thus, by the remarks in Section 3.4, the first statement of the proposition follows from the equality of the functions

$$[j_! \mathbb{k}_{\mathcal{O}_0} \langle nm \rangle]_{q^a} = \sum_i [\mathrm{IC}(\mathcal{O}_i)(\tfrac{i}{2})]_{q^a}$$

for all $a \geq 1$. Evaluating at a point $x \in \mathcal{O}_j(\mathbb{F}_{q^a})$ we need to verify that

$$(-1)^{nm/2} \delta_{0j} q^{-anm/2} = \sum_{0 \leq i \leq j} (-1)^{nm-i^2} q^{a(i^2-nm-i)/2} \begin{bmatrix} j \\ i \end{bmatrix}_{q^a},$$

or equivalently,

$$\delta_{0j} = \sum_{0 \leq i \leq j} (-1)^i q^{i(i-1)/2} \begin{bmatrix} j \\ i \end{bmatrix}_q,$$

which is a standard identity on q -binomial coefficients. The second statement follows from the first by Verdier duality. \square

Proposition 5.8 *We have equalities*

$$\dim \mathrm{Ext}^1(\mathrm{IC}(\mathcal{O}_i), \mathrm{IC}(\mathcal{O}_{i+1})) = \dim \mathrm{Ext}^1(\mathrm{IC}(\mathcal{O}_{i+1}), \mathrm{IC}(\mathcal{O}_i)) = 1.$$

Proof By the Verdier self-duality of IC sheaves, we have an equality of dimensions

$$\dim \mathrm{Ext}^1(\mathrm{IC}(\mathcal{O}_i), \mathrm{IC}(\mathcal{O}_{i+1})) = \dim \mathrm{Ext}^1(\mathrm{IC}(\mathcal{O}_{i+1}), \mathrm{IC}(\mathcal{O}_i)),$$

so we need only give a proof for one.

Using Proposition 5.3, we have that

$$\begin{aligned} \dim \mathrm{Ext}^1(\mathrm{IC}(\mathcal{O}_i), \mathrm{IC}(\mathcal{O}_{i+1})) &= \dim \mathrm{Hom}(\mathrm{IC}(\mathcal{O}_i), \mathrm{IC}(\mathcal{O}_{i+1})[1]) \\ &= \dim \mathrm{Hom}(f_{i+1}^* \mathrm{IC}(\mathcal{O}_i), f_{i+1}^! \mathrm{IC}(\mathcal{O}_{i+1})[1]), \end{aligned}$$

since no other terms that appear in Proposition 5.3 can contribute in this degree (by the conditions for being an IC sheaf).

Recall our small resolution $\pi_i: \tilde{\mathcal{O}}_i \rightarrow \bar{\mathcal{O}}_i$ from earlier. We have

$$f_{i+1}^* \mathrm{IC}(\mathcal{O}_i) = f_{i+1}^* \pi_{i*} \mathbb{k}_{\tilde{\mathcal{O}}_i} [nm - i^2] = H^*(\mathbb{P}^i) \otimes \mathbb{k}_{\bar{\mathcal{O}}_i} [nm - i^2]$$

by the proper base change theorem and the fact that π_i is a fiber bundle with fiber \mathbb{P}^i over \mathcal{O}_{i+1} . Thus

$$\begin{aligned} \text{Hom}(f_{i+1}^* \text{IC}(\mathcal{O}_i), f_{i+1}^! \text{IC}(\mathcal{O}_{i+1})[1]) \\ = \text{Hom}(H^*(\mathbb{P}^i) \otimes \mathbb{k}_{\mathcal{O}_i}[nm - i^2], \mathbb{k}_{\mathcal{O}_i}[nm - (i + 1)^2 + 1]) \\ = \text{Hom}(H^*(\mathbb{P}^i) \otimes \mathbb{k}_{\mathcal{O}_i}, \mathbb{k}_{\mathcal{O}_i}[-2i]) = \mathbb{k} \end{aligned}$$

because $\text{Hom}(\mathbb{k}_{\mathcal{O}_i}, \mathbb{k}_{\mathcal{O}_i}[a]) = 0$ for $a < 0$, and $H^{2i}(\mathbb{P}^i) = \mathbb{k}$. □

Corollary 5.9 *The local chromatographic complex of $j_! \mathbb{k}_{\mathcal{O}_0} \langle nm \rangle$ is the unique complex of the form*


$$0 \rightarrow \text{IC}(\mathcal{O}_0) \rightarrow \text{IC}(\mathcal{O}_1)\langle 1 \rangle \rightarrow \dots \rightarrow \text{IC}(\mathcal{O}_i)\langle i \rangle \rightarrow \dots$$

where all differentials are nonzero. Similarly, that for $j_* \mathbb{k}_{\mathcal{O}_0} \langle nm \rangle$ is the unique complex of the form

$$\dots \rightarrow \text{IC}(\mathcal{O}_i)\langle -i \rangle \rightarrow \dots \rightarrow \text{IC}(\mathcal{O}_1)\langle -1 \rangle \rightarrow \text{IC}(\mathcal{O}_0) \rightarrow 0$$

also where all differentials are nonzero.

Remark 4 This corollary shows that this chromatographic complex categorifies the MOY expansion of a crossing in terms of trivalent graphs, with $\text{IC}(\mathcal{O}_i)$ corresponding to the following MOY graph:



Proof The terms in the complex are determined by Proposition 5.7, and Proposition 5.8 implies that the isomorphism type of the complex is just determined by which maps are nonzero. Since $j_! \mathbb{k}_{\mathcal{O}_0}$ and $j_* \mathbb{k}_{\mathcal{O}_0}$ are indecomposable, all these maps must be nonzero by Corollary 5.5. □

6 The invariant via bimodules

6.1 The global chromatographic complex of a crossing

The following lemma gives a description of $\tilde{\mathcal{O}}_i$ as a “Bott–Samelson” type space.

Lemma 6.1 We have an isomorphism of $(P_{m,n} \times P_{n,m})$ -equivariant varieties

$$\tilde{\mathcal{O}}_i \cong P_{m,n} \times_{P_{i,m-i,n}} P_{i,m+n-i} \times_{P_{i,n-i,m}} P_{n,m}.$$

Proof The map sending $[g, h, k]$ to (gV_i, ghV_n, ghk) defines a closed embedding

$$P_{m,n} \times_{P_{i,m-i,n}} P_{i,m+n-i} \times_{P_{i,n-i,m}} P_{n,m} \hookrightarrow \mathrm{Gr}_i^m \times \mathrm{Gr}_n^{n+m} \times G_{m+n}.$$

Its image is given by triples (W, V, g) satisfying $W \subset V$ and $V = gV_n$, which is isomorphic to $\tilde{\mathcal{O}}_i$ under the map forgetting V . □

Definition 6.2 We let $R_{i_1, \dots, i_m} = \mathbb{k}[x_1, \dots, x_m]^{S_{i_1} \times \dots \times S_{i_m}}$ be the rings of partially symmetric functions corresponding to Young subgroups. We will use without further mention the canonical isomorphism $R_{i_1, \dots, i_m} \cong H^*(BG_{i_1, \dots, i_m})$ sending Chern classes of tautological bundles to elementary symmetric functions.

Given a graded module or bimodule M over any ring R , we let $M(n)$ be the same module with the grading decreased by n .

Corollary 6.3 As $(R_{m,n} \otimes R_{n,m})$ -modules, we have natural isomorphisms

$$\begin{aligned} H_{P_{m,n} \times P_{n,m}}^*(\tilde{\mathcal{O}}_i) &\cong M_i \stackrel{\text{def}}{=} R_{i,m-i,n} \otimes_{R_{i,m+n-i}} R_{i,n-i,m}, \\ \mathbb{H}_{P_{m,n} \times P_{n,m}}^*(\mathrm{IC}(\mathcal{O}_i)) &\cong M_i(nm - i^2). \end{aligned}$$

Proof The first equality follows immediately from the main theorem of [4] (which we restated in the most convenient form for our work in our earlier paper [26, Theorem 3.3]) and Lemma 6.1. The second is a consequence of Corollary 5.2. □

Now have a global version of Proposition 5.8:

Proposition 6.4 The spaces of bimodule maps

$$\mathrm{Hom}_{R_{m,n} \otimes R_{n,m}}(M_i(-2i), M_{i-1}) \quad \text{and} \quad \mathrm{Hom}_{R_{m,n} \otimes R_{n,m}}(M_i(2i), M_{i+1})$$

are trivial in degrees < 1 , and one-dimensional in degree 1.

Proof This follows from [28, Theorem 5.4.1]. In fact, combined with Proposition 5.3, the theorem cited above implies that we have isomorphisms

$$\begin{aligned} \mathrm{Hom}_{R_{m,n} \otimes R_{n,m}}(M_i(-2i), M_{i-1}) &\cong \mathrm{Hom}^\bullet(\mathrm{IC}(\mathcal{O}_i), \mathrm{IC}(\mathcal{O}_{i-1})), \\ \mathrm{Hom}_{R_{m,n} \otimes R_{n,m}}(M_i(2i), M_{i+1}) &\cong \mathrm{Hom}^\bullet(\mathrm{IC}(\mathcal{O}_i), \mathrm{IC}(\mathcal{O}_{i+1})), \end{aligned}$$

with grading degree on module maps matching the homological grading. Thus, this result is equivalent to Proposition 5.8. □

Corollary 6.5 *The global chromatographic complex of $j_1 \mathbb{k}_{\mathcal{O}_0} \langle nm \rangle$ is the unique complex of the form*

$$(6) \quad M^- = \dots \xrightarrow{\partial_{i+1}^-} M_{i+1}(nm - i(i + 1)) \xrightarrow{\partial_i^-} M_i(nm - i(i - 1)) \xrightarrow{\partial_{i-1}^-} \dots$$

where all differentials are nonzero. Similarly, that for $j_* \mathbb{k}_{\mathcal{O}_0} \langle nm \rangle$ is the unique complex of the form

$$(7) \quad M^+ = \dots \xrightarrow{\partial_{i-1}^+} M_i(nm - i(1 + i)) \xrightarrow{\partial_i^+} M_{i+1}(nm - (i + 1)(i + 2)) \xrightarrow{\partial_{i+1}^+} \dots$$


also where all differentials are nonzero.

We note that these are the complexes defined in [17, Section 8], with slight change in grading shift, since they have the same modules, and there is only one such complex up to isomorphism.

We note that these maps have a geometric origin. Consider the correspondence

$$\tilde{\mathcal{O}}_{i+1,i} = \{(U, W, g) \in \text{Gr}_{i+1}^n \times \text{Gr}_i^n \times G_{n+m} \mid gV_n \cap V_m \supset U \supset W\}.$$

Obviously, we have natural maps:



Proposition 6.6 *Up to scaling, we have equalities*

$$\partial_i^- = (p_i^2)_*(p_i^1)^*, \quad \partial_i^+ = (p_i^1)_*(p_i^2)^*.$$

Proof We note that $(p_i^2)_*(p_i^1)^*$ has the expected degree and is nonzero. Thus it must be ∂_i^- . Similarly with $(p_i^1)_*(p_i^2)^*$. □

6.2 Building the global chromatographic complex, I: via canopolises

Now, we are faced with the question of how to build the global chromatographic complex of an arbitrary braid fragment (by which we mean a tangle that can be completed to a closed braid by planar algebra operations).

While the operations we describe are nothing complicated or mysterious, it can be a bit difficult to both be precise and not pile on unnecessary notation. In an effort to give an understandable account for all readers, we give two similar, but slightly different, expositions of how to build the complex for a knot: one quite analogous to

Khovanov’s exposition in [12] using braids and their closures, and one in the language of planar algebras and canopolises, in the vein of the work of Bar-Natan [2] and the first author [24].

This approach is based around planar diagrams in sense of planar algebra: a planar diagram is a crossingless tangle diagram in a planar disk with holes. A canopolis is a way of formalizing the process of building up a tangle by gluing smaller tangles into planar diagrams.

Our definition of our geometric invariant can be phrased in this language. Given a tangle T written as a union of smaller tangles T_i in a planar diagram D , the space X_T has a product decomposition $X_T \cong \prod_i X_{T_i}$, and G_T is a subgroup of $\prod_i G_{T_i}$, given by taking the diagonal inside the factors corresponding to the edges on T_i and T_j identified by D .

That is, the sheaf \mathcal{F}_D can be built from the sheaves corresponding to crossings by successive applications of exterior product and restriction of groups. It is easy to understand how each of these affects chromatographic complexes, and our desired invariant can be built piece by piece.

Formally, to each oriented colored tangle diagram in a disk with boundary points $\{p_1, \dots, p_m\}$, we will associate a complex of modules over $R_\Pi = H^*(\prod_i BG_{p_i})$, where we use Π to denote all the boundary data of the tangle (the points, their coloring, their orientation).

The association of the category $\mathcal{K}(R_\Pi\text{-mod})$ of complexes up to homotopy over R_Π to the boundary data Π (with their colorings) is a canopolis \mathcal{K} , where the functor associated to a planar diagram is an analogue of that used in the canopolis \mathcal{M}_0 in [24]. The canopolis functor

$$\tilde{\eta}: \mathcal{K}(R_{\Pi_1}\text{-mod}) \times \dots \times \mathcal{K}(R_{\Pi_k}\text{-mod}) \rightarrow \mathcal{K}(R_{\Pi_0}\text{-mod}),$$

associated to a planar diagram with outer circle labeled with Π_0 and k inner circles labeled with Π_1, \dots, Π_k , will be given by tensoring with a complex of (R_{Π_0}, R_{Π_*}) -bimodules, where $R_{\Pi_*} = R_{\Pi_1} \otimes \dots \otimes R_{\Pi_k}$.

Let $\mathcal{A}(\eta)$ be the set of arcs in η , let α_a, ω_a be the tail and head of $a \in \mathcal{A}(\eta)$, and let n_a be the integer a is colored with. Associated to each arc is the sequence

$$(e_1(\omega_a) - e_1(\alpha_a), \dots, e_{n_a}(\omega_a) - e_{n_a}(\alpha_a)),$$

which identifies the classes $e_i \in H^*(BG_n)$ corresponding to the elementary symmetric polynomials (geometrically, these are the Chern classes of the tautological bundle on BG_n) for the endpoints connected by the arc. To our diagram, we associate the concatenation of these sequences.

Let $\kappa(\eta)$ be the Koszul complex over $R_{\Pi_0} \otimes \cdots \otimes R_{\Pi_k}$ of this concatenated sequence for our diagram η , which we think of as a bimodule with the R_{Π_0} action on the left and the R_{Π_*} on the right.

Definition 6.7 The canopolis functor $\tilde{\eta}$ associated to the diagram η is $\kappa(\eta) \otimes_{R_{\Pi_*}} -$.

Proposition 6.8 The map sending a tangle T to the global chromatographic complex of \mathcal{F}_T is a canopolis map.

Proof We simply need to justify why tensoring with such a Koszul resolution (which is a free resolution of the diagonal bimodule for $H^*(BG_{p_i})$) is the same as changing G_T to only include the diagonal subgroup of $G_{\omega_a} \times G_{\alpha_a}$. This is one of the basic results of [4]; as we mentioned earlier, this is rephrased most conveniently for us in [26, Theorem 3.3]. □

Remark 5 We note that this construction at no point used the fact that our diagram should be a braid fragment; unfortunately, it is unclear whether our construction will be invariant under the oppositely oriented Reidemeister II move, as with Khovanov and Rozansky’s original construction (see, for example, [24, Section 3]), though we will note that proving invariance under this move for the labeling with all labels 1 is sufficient to imply it for every labeling, by the same cabling arguments we will use later.

6.3 Building the global chromatographic complex, II: via bimodules

A less flexible, but perhaps more familiar, perspective is to associate to each braid a complex of bimodules, in a manner similar to [12] (though the same complex had previously appeared in other works on geometric representation theory). In the case where all labels are 1, our construction will coincide with Khovanov’s.

As in Section 4.2, we let β be a braid with n strands, and $\mathbf{n} = (i_1, \dots, i_m)$ be the labels of the top end of the strands (so $\mathbf{n}\beta$ is the labeling of the bottom end). In that section, we showed that our invariant can also be described in terms of the chromatographic complex of a sheaf Φ_β on G_N .

This sheaf has the advantage that it can be built from the sheaves for smaller braids by convolution of sheaves. However, convolution of sheaves is a geometric operation which is not always easy to understand. Thus, we will give a description of it using the tensor product of bimodules. Let $F(\beta)$ be the $P_{\mathbf{n}} \times P_{\mathbf{n}\beta}$ -equivariant global chromatographic complex of Φ_β , considered as a complex of bimodules over $H^*(BP_{\mathbf{n}})$ and $H^*(BP_{\mathbf{n}\beta})$.

Proposition 6.9 *We have natural isomorphisms*

$$F(\beta\beta') \cong F(\beta) \otimes_{H^*(BP_{n\beta})} F(\beta').$$

Proof Form the exterior product $\Phi_\beta \boxtimes \Phi_{\beta'}$ on $G_N \times G_N$. The $P_n \times P_{n\beta} \times P_{n\beta} \times P_{n\beta\beta'}$ -equivariant chromatographic complex of this is $F(\beta) \otimes_{\mathbb{C}} F(\beta')$. If we restrict to the diagonal $P_{n\beta}$, then this complex is

$$F(\beta) \overset{L}{\otimes}_{H^*(BP_{n\beta})} F(\beta').$$

By the equivariant formality of all simple, Schubert-smooth perverse sheaves on a partial flag variety, $F(\beta)$ is free as a right module, so it is not necessary to take derived tensor product.

By the convolution description, we have

$$\Phi_{\beta\beta'} \cong_{P_n \times P_{n\beta} \times P_{n\beta\beta'}}^{P_n \times P_{n\beta\beta'}} \mu_*(\Phi_{\beta,\beta'}),$$

where $\mu: G_N \times G_N \rightarrow G_N$. Since $G/P_{n\beta}$ is projective, this map simply has the effect of forgetting the $H^*(BP_{n\beta})$ action on each page of the chromatographic spectral sequence. □

Thus, we can construct $F(\beta)$ just by knowing the complex $F(\sigma_i^{\pm 1})$ for the elementary twists $\sigma_i^{\pm 1}$. However, first we must compute the corresponding sheaves. Given \mathbf{n} , we let $Q_j = P_{i_1, \dots, i_j + i_{j+1}, \dots, i_n}$ and $\mathring{Q}_j = Q_j - Q_0$.

Proposition 6.10 *We have isomorphisms*

$$\Phi_{\sigma_i} = j_* \mathbb{k}_{\mathring{Q}_i} \langle i_i i_{i+1} \rangle, \quad \Phi_{\sigma_i^{-1}} = j! \mathbb{k}_{\mathring{Q}_i} \langle i_i i_{i+1} \rangle,$$

where $j: \mathring{Q}_i \hookrightarrow G_N$ is the obvious inclusion.

The global complex of this is very close to the complex M^+ described in (6), considered as a complex of $(R_{i_i, i_{i+1}}, R_{i_{i+1}, i_i})$ -bimodules. However, we must extend scalars to get a complex of $(R_{\mathbf{n}}, R_{\sigma_i \mathbf{n}})$ -bimodules:

Proposition 6.11 $F(\sigma_i^{\pm 1}) = R_{i_1, \dots, i_{i-1}} \otimes_{\mathbb{Q}} M^{\pm} \otimes_{\mathbb{Q}} R_{i_{i+2}, \dots, i_k}$.

Again, this is precisely the complex given in [17, Section 8] up to grading shift.

If $n\beta = \mathbf{n}$, then we can close this braid to a link. Our definition of the knot invariant for this link is the equivariant chromatographic complex for the diagonal P_n action. By the authors' previous work [26, Theorem 1.2], this coincides with the Hochschild homology $\text{HH}^*(F(\beta))$, applied termwise, of the complex $F(\beta)$.

Proposition 6.12 *The cohomology of the complex $\mathrm{HH}_{R_n}^*(F(\beta))$ coincides with the invariant $\mathcal{A}_2(\widehat{\beta})$ of the closure of the braid.*

In fact, the chromatographic spectral sequence is exactly the natural spectral sequence

$$\mathcal{H}^i(\mathrm{HH}^j(F(\beta))) \Rightarrow \mathcal{H}^{i+j}(K \otimes_{R_n \otimes R_n} F(\beta)),$$

where K is a free resolution of R_n as a $R_n \otimes R_n$ -module.

Proof Let $\pi: G_N \rightarrow \mathrm{pt}$, and consider the object $\pi_*\Phi_\beta$ in the equivariant derived category $D_{P_n \times P_n}(\mathrm{pt})$. Under the equivalence to R_n -dg-bimodules given in [25, Theorem 7], this is sent to the complex $F(\sigma)$. Similarly, the weight filtration is sent to that induced by thinking of $F(\beta)$ as a complex. Thus, the spectral sequences match under this equivalence. \square

Since $\mathcal{H}^*(\mathrm{HH}^*(F(\beta)))$ is precisely the invariant proposed by [17], Theorem 1.4 follows immediately.

7 Decategorification

We also wish to show that our knot invariant is, in fact, a categorification of the HOMFLYPT polynomial.

7.1 A categorification of the Hecke algebra

This requires a few basic results about the relationship between sheaves on G_n and the Hecke algebra H_n . As usual, $B = P_{1, \dots, 1}$ is the standard Borel.

Definition 7.1 The Hecke algebra H_n is the algebra over $\mathbb{Z}[q^{1/2}, q^{-1/2}]$ given by the quotient of the group algebra of the braid group \mathcal{B}_n by the quadratic relation

$$(\sigma_i + q^{1/2})(\sigma_i - q^{-1/2}) = 0$$

for each elementary twist σ_i .

Proposition 7.2 [14] *The Grothendieck group $K^0(D_{B \times B}^b(G_n))$ of the equivariant derived category $D_{B \times B}^b(G_n)$ is isomorphic to the Hecke algebra H_n , with the convolution product decategorifying to the algebra product in H_n .*

This map is fixed by the assignment

$$[j_* \underline{\mathbb{k}}_{B s_i B}] \mapsto q^{1/2} \sigma_i,$$

where $j: B s_i B \hookrightarrow G_n$ is the obvious inclusion.

Let \mathcal{F} be a $B \times B$ -equivariant sheaf on G_n . Then we have a map

$$\mathcal{E}_B(G; \mathcal{F}) = \sum_{i,j,k} (-1)^\ell q^{j/2} t^k \dim \mathbb{H}_{B_\Delta}^{j-\ell; j-k}(\mathrm{gr}_\ell^W \mathcal{F}),$$

sending the class of \mathcal{F} in the Grothendieck group to the bigraded Euler characteristic of its global chromatographic complex, often called the *mixed Hodge polynomial*.

This map agrees with a previously known trace on the Hecke algebra, a fact that the authors have proven in a separate note, due to its independent interest and separate connection to the question of constructing Markov traces on general Hecke algebras.

Proposition 7.3 [27, Theorem 1] *The map $\mathcal{E}_B(G_n; -)$ is the Jones–Ocneanu trace Tr (see [11]) on H_n with appropriate normalization factors.*

Remark 6 This geometric definition applies equally well to any simple Lie group, and defines a canonical trace on the Hecke algebra for any type. In fact, our construction can be modified in a straightforward way to a “triply graded homology” invariant on all Artin braid groups. In type B, this can be interpreted as a homological knot invariant for knots in the complement of a solid torus.

7.2 Decategorification for colored HOMFLYPT

To apply this result, we must relate our construction to the categorification of the Hecke algebra above. Recall that if σ is a braid with all labels 1, then Φ_σ is an object of $D_{B \times B}^b(G_n)$.

Proposition 7.4 *The class $[\Phi_\sigma] \in H_n$ is the image of σ under the natural map $B_n \rightarrow H_n$.*

This, combined with Proposition 7.3, gives a new proof of the result of Khovanov [12] that when all components are labeled with 1, the invariant

$$\mathcal{E}(L) = \mathcal{E}_{G_D}(X_D; \mathcal{F}_D) = \sum_{i,j,k} (-1)^\ell q^j t^k \dim \mathcal{A}_2^{j;k;\ell}(L)$$

is the appropriately normalized HOMFLYPT polynomial of the link L underlying the diagram D . We wish to extend this to the colored case. For this, we must use a “cabling/projection” formula.

Consider a closable colored braid σ , and let $P = P_n$ and $G = G_N$. We have defined a $P \times P$ -equivariant sheaf Φ_σ on G by the multiplication map $m: X_\sigma \rightarrow G$.

Theorem 7.5 *For any colored link L , the Euler characteristic $\mathcal{E}(D)$ is the (suitably normalized) colored HOMFLYPT polynomial for any diagram D of L .*

In order to prepare for the proof, we show a pair of lemmata. Let σ_{cab} denote the cabling of σ in the blackboard framing with multiplicities given by the colorings, thought of as colored with all labels 1.

Lemma 7.6 *We have an isomorphism of $P \times B$ -equivariant sheaves,*

$$\text{res}_{P \times B}^{P \times P} \Phi_\sigma \cong \text{ind}_{B \times B}^{P \times B} \Phi_{\sigma_{\text{cab}}}.$$

Proof The proof is a straightforward induction on the length of σ , left to the reader. \square

Let λ_n be the partition given by arranging the parts of n in decreasing order, and let λ_n^t be its transpose. Let π_n be the projection in the Hecke algebra to the representations indexed by Young diagrams less than λ_n^t in dominance order. Alternatively, if we identify H_N with the endomorphisms of $V^{\otimes N}$, where V is the standard representation of $U_q(\mathfrak{sl}_m)$ for $m \geq n$, then this is the projection to $\wedge^{i_1} V \otimes \cdots \otimes \wedge^{i_n} V$.

Let $q_P = \sum_{W_P} q^{\ell(w)}$ be the Poincaré polynomial of the flag variety P/B .

Lemma 7.7 *For every complex Φ in $D_{B \times B}^b(G)$, we have*

$$[\text{res}_{P \times B}^{B \times B} \text{ind}_{B \times B}^{P \times B} \Phi] = q_P \pi_P[\Phi].$$

Proof First consider the case where $P = G$. In this case, the sheaf $\text{res}_{G \times B}^{B \times B} \text{ind}_{B \times B}^{G \times B} \Phi$ has a filtration whose successive quotients are of the form $\mathbb{H}^i(\Phi) \otimes \underline{\mathbb{k}}_G$. Thus we have

$$[\text{res}_{G \times B}^{B \times B} \text{ind}_{B \times B}^{G \times B} \Phi] = \dim_q \mathbb{H}^*(\Phi) \cdot [\underline{\mathbb{k}}_G].$$

It is a classical fact that $[\underline{\mathbb{k}}_G] = q_G \pi_G$; here π_G is just the projection to $\wedge^N V$. This computation immediately extends to the general case. \square

Remark 7 This proposition shows why our approach works for colored HOMFLYPT polynomials, but would need to be modified to approach the HOMFLY polynomials for more general type A representations; we lack a good categorification of most of the projections in the Hecke algebra, but π_P has a beautiful geometric counterpart. This may be related to the fact that π_P is the projection not just to a subrepresentation, but in fact to a cellular ideal in H_n .

Proof of Theorem 7.5 Immediately from Lemmata 7.6 and 7.7, we have the equality of Grothendieck classes $[\text{res}_{B \times B}^{P \times P} \Phi_\sigma] = q_P \pi_P[\Phi_{\sigma_{\text{cab}}}]$. Thus

$$\begin{aligned} \mathcal{E}_P(G; \Phi_\sigma) &= q_P^{-1} \mathcal{E}_B(G; \text{res}_{B \times B}^{P \times P} \Phi_\sigma) \\ &= \text{Tr}(q_P^{-1} [\text{res}_{B \times B}^{P \times P} \Phi_\sigma]) \\ &= \text{Tr}(\pi_P[\Phi_{\sigma_{\text{cab}}}]). \end{aligned}$$

By the “projection/cabling” formula (see, for example, [16, Lemma 3.3]), this is precisely the colored HOMFLYPT polynomial. \square

8 The proof of invariance: the 1-colored case

We first concentrate on the simpler case of $GL(2)$ before attacking the general case. In this case, we will obtain an invariant which matches the HOMFLYPT homology of Khovanov and Rozansky [13; 12], so the section below can be thought of as a geometric proof of the invariance of this homology theory.

Recall that if σ is a braidlike diagram on n strands, we described in Section 4.2 a map

$$m: X_\sigma \rightarrow G_n$$

which is equivariant with respect to $\phi: G_\sigma \rightarrow T \times T$, where $T \times T$ acts on G_n by left and right multiplication. This map gives rise to a functor

$${}^{B \times B}_{G_\sigma} m_*: D_{G_D}^+(X_D) \rightarrow D_{T \times T}^+(G_n),$$

and we denoted the image of \mathcal{F}_σ by Φ_σ . We saw that this functor preserves weight filtrations.

Now suppose that w is an element of the symmetric group on n letters (which we regard as permutation matrices in G_n) and that $\sigma = \sigma_{i_1} \sigma_{i_2} \cdots \sigma_{i_p}$ is a (positive) braid in the standard generators corresponding to a reduced expression $s_{i_1} \cdots s_{i_p}$ for w .

It is straightforward to see that if we restrict m to the open set \tilde{U} in G_D consisting of tuples (g_1, \dots, g_p) with each $g_i \in U$, where U denotes the open Bruhat cell in G_2 , then we may factor m as

$$(8) \quad \tilde{U} \rightarrow \tilde{U} / \ker \phi \rightarrow G_n,$$

where the first map is a quotient by a free action, and the second map is an isomorphism.

Moreover, if we denote by B the subgroup of upper triangular matrices, then the image of the restriction of m to \tilde{U} is contained in Schubert cell BwB . It follows that

$$(9) \quad \Phi_\sigma = j_{w*} \mathbb{k}_{BwB} \langle \ell(w) \rangle,$$

where j_w denotes the inclusion of the Bruhat cell BwB into G_n .


Proposition 8.1 *Theorem 1.2 holds in the case where all strands are labeled by 1.*

Proof As usual with proofs that knot invariants defined in terms of a projection are really invariants, we check that our description is unchanged by the Reidemeister moves. Since we only consider closed braids, we only need to check Reidemeister II and III in the braid-like case, when all strands are coherently oriented. Those who prefer to use the Markov theorem can consider the proof of Reidemeister I as a proof of the Markov 1 move, and the Reidemeister II and III calculations as proving the independence of the presentation of our braid in terms of elementary twists *and* of the Markov 2 move (which only uses Reidemeister IIa).

In each case, we will use the fact that while we wish to compare the pushforwards of sheaves corresponding to diagrams D and D' from X_D/G_D and $X_{D'}/G_{D'}$ to a point, we can accomplish this by showing that their pushforwards by any pair of maps to any common space coincide. Being able to use these techniques is one of the principal advantages of a geometric definition over a purely algebraic one.

In each case, the calculation we need to do is local in terms of diagrams. Proposition 6.8 implies that if we show that we have an isomorphism of global chromatographic complexes of two diagrams as modules over the polynomial rings attached to external edges, then “pasting” these into a fixed larger diagram again gives an isomorphism of global chromatographic complexes.

Reidemeister I Consider the following tangles:



To simplify notation we denote the associated varieties by X and X' and groups by G and G' , respectively. We have $X = G_2$ and $X' = G_1$, $G = G_1^3$ and $G' = G_1^2$. The determinant gives a map

$$d: X \rightarrow X',$$

which is equivariant with respect to the map $\phi: G \rightarrow G'$ forgetting the factor corresponding to the internal edge. We wish to exhibit an isomorphism

$$(11) \quad \mathbb{G}' d_* \mathcal{F}_D \cong \mathcal{F}_{D'}$$

compatible with the weight filtrations on both sheaves. Note that the weight filtration on $\mathcal{F}_{D'}$ is trivial, whereas that on \mathcal{F}_D is not.

Let $B \xrightarrow{a} X \xleftarrow{b} BsB$ be the decomposition of $X = G_2$ into its two Bruhat cells. We have a distinguished triangle

$$a_! a^! \underline{\mathbb{k}}_X \langle 1 \rangle \rightarrow \underline{\mathbb{k}}_X \langle 1 \rangle \rightarrow b_* b^* \underline{\mathbb{k}}_X \langle 1 \rangle \xrightarrow{[1]} .$$

Because a is the inclusion of a smooth divisor, $a^! \underline{\mathbb{k}}_X = \underline{\mathbb{k}}_X \langle -2 \rangle = \underline{\mathbb{k}}_X[-2](-1)$. Hence

$$a^! \underline{\mathbb{k}}_X \langle 1 \rangle = \underline{\mathbb{k}}_X[-1](-\frac{1}{2}).$$

Turning the triangle gives the weight filtration on $b_* \underline{\mathbb{k}}_{BsB} \langle 1 \rangle$:

$$(12) \quad \underline{\mathbb{k}}_X \langle 1 \rangle \rightarrow b_* \underline{\mathbb{k}}_{BsB} \langle 1 \rangle \rightarrow a_* \underline{\mathbb{k}}_B(-\frac{1}{2}) \xrightarrow{[1]} .$$

The left (resp. right) hand term is pure of weight 0 (resp. 1). In the following we analyze the effect of $G' d_*$ on this triangle.

The restriction of d to $BsB \subset X$ is a trivial $G_1 \times \mathbb{A}^2$ -bundle over X' . One may easily check that $\ker \phi$ acts freely on the multiplicative group in the fiber. It follows that

$$G' d_* b_* \underline{\mathbb{k}}_{BsB} \cong \underline{\mathbb{k}}_{X'} .$$

On the other hand, the restriction of d to $B \subset X$ yields a trivial $G_1 \times \mathbb{A}^1$ -bundle, with $\ker \phi$ only acting on \mathbb{A}^1 . It follows that


$$G' d_* a_* \underline{\mathbb{k}}_B = H^\bullet(\mathbb{P}^\infty) \otimes H^\bullet(G_1) \otimes \underline{\mathbb{k}}_{X'} .$$

Applying $G' d_*$ to (12) and using the above isomorphisms, we obtain

$$G' d_* \underline{\mathbb{k}}_X \langle 1 \rangle \rightarrow \underline{\mathbb{k}}_{X'} \langle 1 \rangle \rightarrow H^\bullet(\mathbb{P}^\infty) \otimes H^\bullet(G_1) \otimes \underline{\mathbb{k}}_{X'}(-\frac{1}{2}) \xrightarrow{[1]} .$$

As $\text{Hom}(\underline{\mathbb{k}}_{X'}, \underline{\mathbb{k}}_{X'}[i]) = H^i_{G'}(X')$ is zero for $i < 0$ we conclude that the second arrow above is zero. Thus, the induced weight filtration on $\underline{\mathbb{k}}_{X'}$ is trivial. Thus, we have the desired Equation (11). As discussed before, the general case follows from Section 6.2, where we think of adding the rest of the diagram as a canopolis operation.

Reidemeister IIa Here we are concerned with the following two tangles:



We denote the associated varieties and groups X, X', G, G' . We denote by m the multiplication map $X \rightarrow G_2$ considered at the start of this section. We regard X' as the diagonal matrices inside G_2 .

We have seen that $G' m_*$ preserves weight filtrations, and hence we may ignore weight filtrations when comparing $G' m_* \mathcal{F}_D$ and $\mathcal{F}_{D'}$. The map $B \rightarrow X'$ forgetting the off-diagonal entry is acyclic, and therefore it is enough to show that $G' m_* \mathcal{F}_D \cong \underline{\mathbb{K}}_B$.

We decompose G_2 into its Bruhat cells $B \xrightarrow{a} G_2 \xleftarrow{b} B s B$ as before. We claim we have isomorphisms

$$(13) \quad G' m_*(a_* \underline{\mathbb{K}}_B \boxtimes b_! \underline{\mathbb{K}}_{BsB}) \cong b_! \underline{\mathbb{K}}_{BsB},$$

$$(14) \quad G' m_*(\underline{\mathbb{K}}_G \boxtimes a_* \underline{\mathbb{K}}_B) \cong \underline{\mathbb{K}}_G,$$

$$(15) \quad G' m_*(\underline{\mathbb{K}}_G \boxtimes \underline{\mathbb{K}}_G) \cong \underline{\mathbb{K}}_G \oplus \underline{\mathbb{K}}_G \langle -2 \rangle,$$

$$(16) \quad G' m_*(\underline{\mathbb{K}}_G \boxtimes b_! \underline{\mathbb{K}}_{BsB}) \cong \underline{\mathbb{K}}_G \langle -2 \rangle.$$

(As always we regard the exterior tensor product of equivariant sheaves on G_2 as an equivariant sheaf on X via restriction.)

Indeed, (13) and (14) follow from the fact that the restriction of m to $B \times G$ or $G \times B$ is a trivial B -bundle, with $\ker \phi$ acting freely on the multiplicative groups in the fiber. The factorization (8) of m as “essentially a \mathbb{P}^1 -bundle” implies (15). Then (16) follows from the others by taking the exterior tensor product of $\underline{\mathbb{K}}_G$ with the distinguished triangle $b_! \underline{\mathbb{K}}_{BsB} \rightarrow \underline{\mathbb{K}}_G \rightarrow a_* \underline{\mathbb{K}}_B \rightarrow$ and applying $G' m_*$.

Now B is smooth of codimension 1 inside G_2 so $a^! \underline{\mathbb{K}}_G = \underline{\mathbb{K}}_B \langle -2 \rangle$ and we have an exact triangle

$$a_* \underline{\mathbb{K}}_B \langle -2 \rangle \rightarrow \underline{\mathbb{K}}_G \rightarrow b_* \underline{\mathbb{K}}_{BsB} \xrightarrow{[1]}.$$

Taking the exterior tensor product with $b_! \underline{\mathbb{K}}_{BsB}$, applying $G' m_*$ and using the above isomorphisms we obtain a distinguished triangle

$$(17) \quad b_! \underline{\mathbb{K}}_{BsB} \langle -2 \rangle \rightarrow \underline{\mathbb{K}}_G \langle -2 \rangle \rightarrow G' m_*(b_* \underline{\mathbb{K}}_{BsB} \boxtimes b_! \underline{\mathbb{K}}_{BsB}) \xrightarrow{[1]}.$$

Note that $\text{Hom}(b_! \underline{\mathbb{K}}_{BsB}, \underline{\mathbb{K}}_G)$ is one-dimensional and contains the adjunction morphism $b_! b^! \underline{\mathbb{K}}_G \rightarrow \underline{\mathbb{K}}_G$. By considering its dual, one may show that the first arrow in (17) is nonzero. It follows that this arrow is the adjunction morphism (up to a nonzero scalar) and we have an isomorphism

$$G' m_*(b_* \underline{\mathbb{K}}_{BsB} \boxtimes b_! \underline{\mathbb{K}}_{BsB}) \cong \underline{\mathbb{K}}_B \langle -2 \rangle.$$

Finally note that by definition \mathcal{F}_D is $b_* \underline{\mathbb{K}}_{BsB} \boxtimes b_! \underline{\mathbb{K}}_{BsB} \langle 2 \rangle$ and so

$$G' m_* \mathcal{F}_D \cong \underline{\mathbb{K}}_B,$$

which finishes the proof of invariance under Reidemeister II.

Reidemeister III This follows immediately from the considerations at the beginning of this section. Indeed, if σ and σ' are the diagrams corresponding to the words $\sigma_1\sigma_2\sigma_1$ and $\sigma_2\sigma_1\sigma_2$ we have maps

$$X_\sigma \xrightarrow{m} G_3 \xleftarrow{m'} X_{\sigma'}$$

and we have

$${}^{T \times T}_{G_\sigma} m_* \mathcal{F}_\sigma \cong j_{w_0} \mathbb{k}_{B w_0 B} \cong {}^{T \times T}_{G_\sigma} m'_* \mathcal{F}_{\sigma'},$$

where w_0 indicates the longest element in S_3 . □

9 The proof of invariance: arbitrary labels

Now, we expand to the full case of all possible positive integer labels.

Proof of Theorem 1.2 All of the Reidemeister moves can simply be reduced to the corresponding statement for the cabling with all labels 1. Interestingly, the same trick was used in [17] to prove invariance in a special case. Almost certainly our proof could be rephrased in a purely algebraic language like their paper, though at the moment it is unclear how.

Reidemeister IIa & III Here we need only establish the isomorphisms of $P \times P$ -equivariant sheaves

$$\Phi_{\sigma_i} \star \Phi_{\sigma_i^{-1}} \cong \mathbb{k}_P \quad \text{and} \quad \Phi_{\sigma_i} \star \Phi_{\sigma_{i+1}} \star \Phi_{\sigma_i} \cong \Phi_{\sigma_{i+1}} \star \Phi_{\sigma_i} \star \Phi_{\sigma_{i+1}}.$$

Lemma 7.6 implies that these hold as $P \times B$ -equivariant sheaves, applying the invariance for the cabling with all labels 1.

In fact, both are the $*$ -inclusion of a local system on a $P \times P$ -orbit: P itself in the first case, and the $P \times P$ orbit of the permutation corresponding to the cabling of $\sigma_i \sigma_{i+1} \sigma_i$ in the second. Since the stabilizer of any point under $P \times P$ is connected, any $P \times B$ -equivariant local system on an orbit has at most one $P \times P$ -equivariant structure, and this equality holds as $P \times P$ -equivariant sheaves.

Reidemeister I We again use the “cabling/projection” philosophy, but this argument requires a bit more subtlety. We are interested in the chromatographic complex of a single crossing with its right ends capped off; that is, the tangle projection denoted by D in (10). To construct the sheaf \mathcal{F}_D , we take $U \subset G_{2n}$, as defined in (5), and consider $j_* \mathbb{k}_U \langle n^2 \rangle$ or $j! \mathbb{k}_U \langle n^2 \rangle$, depending on whether our crossing is positive or negative. These cases are Verdier dual, and the proofs of invariance are essentially identical, so we will treat the positive case, and only note where the negative differs. If we consider this sheaf equivariantly for the action of $G_{n,n}$ on the left and the right, then we obtain the sheaf attached to a single crossing with label n on both strands.

By convention, we let G^1 denote the first copy of $G_n \subset G_{n,n}$ and G^2 the second. As before, we let T_n be diagonal matrices in G_n , and we use T^1, T^2 for the inclusions into the two factors. We let $G^{1,1,2}$ denote $G^1 \times G^1 \times (G^2)_\Delta$; that is, the left and right action of G^1 , and the conjugation action of G^2 . This is the group G_D for the diagram labeled D in (10). The sheaf \mathcal{F}_D for this diagram is thus $j_! \underline{\mathbb{K}}_U \langle n^2 \rangle$ (or $j_* \underline{\mathbb{K}}_U \langle n^2 \rangle$ if D is taken with a positive crossing) considered equivariantly for $G^{1,1,2}$.

Thus in order to prove the theorem, what we must do is consider the $G^{1,1,2}$ -equivariant global chromatographic complex of \mathcal{F}_D as a $H^*(BG^1)$ -bimodule, and show that it matches that of an untwisted strand (the diagram denoted D' in (10)).

Note that for any G_n sheaf \mathcal{F} on any G_n -space X , the inclusion of the symmetric group as permutation matrices normalizing T_n gives an action of S_n on $\mathbb{H}_{T_n}^*(X; \text{res}_{T_n}^{G_n} \mathcal{F})$.


Lemma 9.1 *The natural transformation of functors*

$$\mathbb{H}_{G^{1,1,2}}^*(G_{2n}; -) \rightarrow \mathbb{H}_{G^{1,1} \times T^2}^*(G_{2n}; \text{res}_{G^{1,1} \times T^2}^{G^{1,1,2}} -)$$

is the inclusion of the S_n -invariants for the permutation action on T^2 .

Proof This is the abelianization theorem for equivariant cohomology; see, for example, [7, Proposition 1]. □

Let \hat{U} be the Bruhat cell $Bw_{2n}^{n,n}B$, where $w_{2n}^{n,n}$ is the permutation which switches i and $i \pm n$, and let \hat{j} be its inclusion to G_{2n} . We note that $\hat{j}_* \underline{\mathbb{K}}_{\hat{U}}$ is Φ_σ where σ is the braid given by the n -cabling of a single crossing:



Lemma 9.2 *The $G^{1,1} \times T^2$ -equivariant global chromatographic complex of $j_* \underline{\mathbb{K}}_U$ is isomorphic to the $T^{1,1} \times T^2$ -equivariant one for $\hat{j}_* \underline{\mathbb{K}}_{\hat{U}}$, with the bimodule structure restricted to $H^*(BG^{1,1}) \subset H^*(BT^{1,1})$.*

Proof Let $Q = G^1 \cap B$ be the upper triangular matrices in G_n . Then

$$\text{ind}_{T^{1,1} \times T^2}^{G^{1,1} \times T^2} j_* \underline{\mathbb{K}}_{\hat{U}} \cong \text{ind}_{Q \times Q \times T^2}^{G^{1,1} \times T^2} \text{ind}_{T^{1,1} \times T^2}^{Q \times Q \times T^2} j_* \underline{\mathbb{K}}_{\hat{U}} \cong \text{res}_{T^{1,1} \times T^2}^{G^{1,1,2}} j_* \underline{\mathbb{K}}_U.$$

The first induction leaves chromatographic complexes unchanged, since Q and T^1 are homotopy equivalent, and $j_* \underline{\mathbb{K}}_{\hat{U}}$ is smooth on $Q \times Q$ -orbits.

For the second, we have a projective map

$$\mu: G_n \times_Q \widehat{U} \times_Q G_n \rightarrow G_{2n},$$

which induces an isomorphism

$$G_n \times_Q \widehat{U} \times_Q G_n \cong U.$$

By [25, Theorem 5], under taking equivariant cohomology, induction of sheaves corresponds to the restriction of scalars, and since G_n/Q is projective, pushforward preserves purity and thus commutes with taking local chromatographic complex. This means that the result extends to all terms in the chromatographic spectral sequence. \square

By definition, the $T^{1,1} \times T^2$ -equivariant chromatographic complex for $\widehat{j}_* \underline{\mathbb{k}} \widehat{U}$ is just the complex of bimodules for the tangle diagram D_{cab} corresponding to closing the right half of the strands in the braid above. Applying the invariance result for labelings with all labels 1, this is the same as the complex corresponding to a full twist of n strands. Since $\widehat{j}_* \underline{\mathbb{k}} \widehat{U}$ is in fact equivariant for $T^{1,1} \times G^2$, this has an S_n action, which is compatible with its module structure over $H^*(BT^2) \cong \mathbb{k}[x_1, \dots, x_n]$. Doing this straightening one strand at a time, we see that the actions of $H^*(BT^2)$ and $H^*(BT^1) \cong \mathbb{k}[y_1, \dots, y_n]$ are intertwined by the map sending x_i to y_{n+1-i} . Thus, the S_n action discussed above is compatible with the standard S_n -module action on $H^*(BT_n)$ acting on the left and right after conjugation by the longest element w_0 .

Note that if we consider a negative crossing, we will have to include n times the usual shift for removing a negative stabilization, but this is easily accounted for in the normalization.

Restricted to symmetric polynomials (that is, $H^*(BG_n)$), every Soergel bimodule is a number of copies of the regular bimodule, and every map in the complex for a single crossing splits, so restricted to $H^*(BG_n)$, the complex attached to a braid with all labels 1 is homotopic to a single copy of $H^*(BT_n)$ with the regular bimodule action and standard S_n action (conjugated by the longest element w_0).

By Lemma 9.1, to obtain the $G^{1,1,2}$ -equivariant global chromatographic complex we simply take S_n -invariants and thus we obtain a single copy of the regular bimodule for $H^*(BG_n)$, as desired. \square

References

- [1] **M Artin, A Grothendieck, J L Verdier**, *Théorie des topos et cohomologie étale des schémas, Tome 3: Exposés IX–XIX (SGA 4₃)*, Lecture Notes in Math. 305, Springer (1973) MR

- [2] **D Bar-Natan**, *Khovanov's homology for tangles and cobordisms*, *Geom. Topol.* 9 (2005) 1443–1499 MR
- [3] **A Beilinson, V Ginzburg, W Soergel**, *Koszul duality patterns in representation theory*, *J. Amer. Math. Soc.* 9 (1996) 473–527 MR
- [4] **J Bernstein, V Lunts**, *Equivariant sheaves and functors*, *Lecture Notes in Math.* 1578, Springer (1994) MR
- [5] **A A Beilinson, J Bernstein, P Deligne**, *Faisceaux pervers*, from “Analysis and topology on singular spaces, I”, *Astérisque* 100, Soc. Math. France, Paris (1982) 5–171 MR
- [6] **M V Bondarko**, *Weight structures and motives; comotives, coniveau and Chow-weight spectral sequences, and mixed complexes of sheaves: a survey*, preprint (2009) arXiv
- [7] **M Brion**, *Equivariant cohomology and equivariant intersection theory*, from “Representation theories and algebraic geometry” (A Broer, G Sabidussi, editors), *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.* 514, Kluwer Acad., Dordrecht (1998) 1–37 MR
- [8] **P Deligne**, *Théorie de Hodge, II*, *Inst. Hautes Études Sci. Publ. Math.* (1971) 5–57
- [9] **P Deligne**, *Cohomologie étale (SGA 4^{1/2})*, *Lecture Notes in Math.* 569, Springer (1977) MR
- [10] **P Deligne**, *La conjecture de Weil, II*, *Inst. Hautes Études Sci. Publ. Math.* (1980) 137–252 MR
- [11] **V F R Jones**, *Hecke algebra representations of braid groups and link polynomials*, *Ann. of Math.* 126 (1987) 335–388 MR
- [12] **M Khovanov**, *Triply-graded link homology and Hochschild homology of Soergel bimodules*, *Internat. J. Math.* 18 (2007) 869–885 MR
- [13] **M Khovanov, L Rozansky**, *Matrix factorizations and link homology, II*, *Geom. Topol.* 12 (2008) 1387–1425 MR
- [14] **R Kiehl, R Weissauer**, *Weil conjectures, perverse sheaves and l -adic Fourier transform*, *Ergeb. Math. Grenzgeb.* 42, Springer (2001) MR
- [15] **G Laumon**, *Transformation de Fourier, constantes d'équations fonctionnelles et conjecture de Weil*, *Inst. Hautes Études Sci. Publ. Math.* (1987) 131–210 MR
- [16] **X-S Lin, H Zheng**, *On the Hecke algebras and the colored HOMFLY polynomial*, *Trans. Amer. Math. Soc.* 362 (2010) 1–18 MR
- [17] **M Mackaay, M Stošić, P Vaz**, *The 1, 2-coloured HOMFLY-PT link homology*, *Trans. Amer. Math. Soc.* 363 (2011) 2091–2124 MR
- [18] **J McCleary**, *A user's guide to spectral sequences*, 2nd edition, *Cambridge Studies in Advanced Math.* 58, Cambridge Univ. Press (2001) MR
- [19] **H Murakami, T Ohtsuki, S Yamada**, *Homfly polynomial via an invariant of colored plane graphs*, *Enseign. Math.* 44 (1998) 325–360 MR

- [20] **D Pauksztello**, *Compact corigid objects in triangulated categories and co-t-structures*, Cent. Eur. J. Math. 6 (2008) 25–42 MR
- [21] **C A M Peters, J H M Steenbrink**, *Mixed Hodge structures*, Ergeb. Math. Grenzgeb. 52, Springer (2008) MR
- [22] **M Saito**, *Mixed Hodge modules*, Proc. Japan Acad. Ser. A Math. Sci. 62 (1986) 360–363 MR
- [23] **C Stroppel**, *A structure theorem for Harish-Chandra bimodules via coinvariants and Golod rings*, J. Algebra 282 (2004) 349–367 MR
- [24] **B Webster**, *Khovanov–Rozansky homology via a canopolis formalism*, Algebr. Geom. Topol. 7 (2007) 673–699 MR
- [25] **B Webster, G Williamson**, *The bounded below equivariant derived category*, in preparation
- [26] **B Webster, G Williamson**, *A geometric model for Hochschild homology of Soergel bimodules*, Geom. Topol. 12 (2008) 1243–1263 MR
- [27] **B Webster, G Williamson**, *The geometry of Markov traces*, preprint (2009) arXiv
- [28] **G Williamson**, *Singular Soergel bimodules*, Int. Math. Res. Not. 2011 (2011) 4555–4632 MR

*Department of Pure Mathematics, University of Waterloo
and Perimeter Institute for Theoretical Physics
Waterloo ON, Canada*

*School of Mathematics and Statistics, University of Sydney
Sydney NSW, Australia*

ben.webster@uwaterloo.ca, g.williamson@sydney.edu.au

Proposed: Peter S. Ozsváth
Seconded: Ciprian Manolescu, Haynes Miller

Received: 28 September 2010
Revised: 25 June 2016