Quantum Satake in type A. Part I

Ben Elias[*]

Abstract. We give an interpretation of $\mathfrak{sl}_n$-webs as morphisms between certain singular Soergel bimodules. We explain how this is a combinatorial, algebraic version of the geometric Satake equivalence (in type A). We then q-deform the construction, giving an equivalence between representations of $U_q(\mathfrak{sl}_n)$ and certain singular Soergel bimodules for a q-deformed Cartan matrix.

In this paper, we discuss the general case but prove only the case $n = 2, 3$. In the sequel we will prove $n \geq 4$.

1. Introduction

1.1. Equivalent equivalences. The *geometric Satake equivalence* (or just *geometric Satake* for short) is an equivalence between two symmetric monoidal abelian categories which can be attached to a reductive algebraic group. In this paper we state a *Soergel Satake equivalence*, an equivalence of (strict) additive 2-categories associated to a pair of Langlands dual lie algebras. Furthermore, in type A we state an *algebraic Satake equivalence*, an equivalence between additive 2-categories living inside the Soergel Satake equivalence. The fact that these three equivalences imply each other is reasonably straightforward (given the results of Soergel and Härterich). The real meat of this paper is computational: an explicit construction of the algebraic Satake equivalence in type A, coming from a presentation of both additive 2-categories by generators and relations. This gives a new, simple proof of geometric Satake in type A (but see Remark 1.1). Finally, we present a q-deformation of the Soergel and algebraic Satake equivalences in type A. This q-deformation has no known geometric source at present.

Let $\mathfrak{g}$ be a simple Lie algebra with Langlands dual $\mathfrak{g}^\vee$. Let Ω denote the fundamental group of $\mathfrak{g}^\vee$, a finite abelian group realized as the weight lattice

modulo the root lattice. Representations of $\mathfrak{g}^\vee$ form a semisimple Ω-graded-monoidal category, in the sense that each irreducible object has a character (or highest weight) in Ω, and these characters add under taking tensor products and their summands. This additive Ω-graded-monoidal category can also be encoded as an additive 2-category Rep^Ω with one object for each element of Ω. We give an introduction to these constructions in §4.2.

Associated to the affine Dynkin diagram $\widetilde{\Gamma}$ of $\mathfrak{g}$, one can construct an additive 2-category of *singular Soergel bimodules* $\mathcal{S}\mathbb{S}\mathrm{Bim}$ as in [45], having one object for each proper subgraph of $\widetilde{\Gamma}$. Chapter 5 contains an introduction to singular Soergel bimodules. Within $\mathcal{S}\mathbb{S}\mathrm{Bim}$ lies the full sub-2-category of *maximally singular Soergel bimodules $m\mathcal{S}\mathbb{S}\mathrm{Bim}$*, which has objects for each subgraph of $\widetilde{\Gamma}$ isomorphic to the original Dynkin diagram Γ. Such subgraphs are parametrized naturally by the set Ω.

The Soergel Satake equivalence is an equivalence between Rep^Ω and $m\mathcal{S}\mathbb{S}\mathrm{Bim}$. Actually, 2-morphisms in $m\mathcal{S}\mathbb{S}\mathrm{Bim}$ are graded vector spaces, so to be more precise, Soergel Satake is a 2-functor from Rep^Ω to $m\mathcal{S}\mathbb{S}\mathrm{Bim}$ which is essentially surjective up to grading shift, faithful, and full onto degree 0 maps.

The original geometric Satake equivalence (see, for instance, [18]) can be similarly rephrased as an equivalence of 2-categories between Rep^Ω and some 2-category of perverse sheaves on affine partial flag varieties. Maximally singular Soergel bimodules should be (roughly) thought of as the equivariant hypercohomologies of these perverse sheaves, though thankfully they have an independent algebraic definition.

Inside Rep^Ω one can consider the additive sub-2-category Fund^Ω whose 1-morphisms are tensor products of fundamental representations. There are certain 1-morphisms in $m\mathcal{S}\mathbb{S}\mathrm{Bim}$ which correspond to fundamental representations in Rep^Ω, which could be called *fundamental* singular Soergel bimodules. Clearly, Fund^Ω should be equivalent to the monoidal sub-2-category of $m\mathcal{S}\mathbb{S}\mathrm{Bim}$ generated by fundamental singular Soergel bimodules[1], a statement equivalent to Soergel Satake. Sadly, neither 2-category is well-understood in general.

In type A, however, all fundamental representations are miniscule. In this case, the fundamental singular Soergel bimodules are actually easy to describe, being the generating 1-morphisms in the 2-category of *maximally singular Bott–Samelson bimodules $m\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim}$* inside $m\mathcal{S}\mathbb{S}\mathrm{Bim}$. Moreover, in type A both Fund^Ω and $\mathcal{S}\mathbb{S}\mathrm{Bim}$ have a presentation by generators and relations using planar diagrams ([5] and [10] respectively), which is why we consider these categories to be "algebraic." The equivalence between Fund^Ω and $m\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim}$ in type A will be called the algebraic Satake equivalence, and will be proven by exploiting these presentations.

This paper proves the algebraic Satake equivalence for $\mathfrak{sl}_2$ and $\mathfrak{sl}_3$. We require several results about singular Soergel bimodules, which for $\mathfrak{sl}_2$ are all available in [8],

---

[1]More precisely, there should be a functor from Fund^Ω to this sub-2-category which is essentially surjective up to grading shift, faithful, and full onto degree 0 maps.

and for $\mathfrak{sl}_3$ are proven in the appendix. For $\mathfrak{sl}_{n+1}$ with $n \geq 3$, the requisite background will eventually appear in joint work with Williamson [10], though we do not know when this manuscript will become available. Moreover, the $n \geq 3$ case also requires a great deal of additional calculation. For these reasons, we postpone the $n \geq 3$ case to a followup paper.

Remark 1.1. The feature which is most obscured by the transformation from geometric Satake to algebraic Satake is the symmetric structure on the original monoidal category. The symmetric structure on perverse sheaves arises in a complicated way from the study of the Beilinson–Drinfeld Grassmannian [35], and it is unclear how this is translated into the language of Soergel bimodules. One could make the argument that geometric Satake without the symmetric structure is a significantly weaker theorem (we discuss this in more detail in Remark 6.1). However, this argument seems to rely on the Decomposition Theorem [3], while this paper gives a proof of geometric Satake (or rather, Soergel Satake) without any reliance on the Decomposition Theorem or the related Soergel conjecture. Moreover, making the equivalence explicit on 2-morphisms is a non-trivial result, and the ability to discuss the equivalence in a new language (without Tannakian formalism) is useful. In addition, the q-deformation below is new.

Remark 1.2. One can upgrade algebraic Satake to an equivalence of symmetric monoidal categories, using recent work of Lusztig. It is easy to equip Fund^{Ω} with a natural symmetric structure, using the computational formulas (well known to knot theorists) for the symmetric structure in representation theory. When this paper was first written, there was no natural way to construct a natural symmetric structure on singular Soergel bimodules. However, Lusztig has since devised such a structure in an update to [34]. Technically, his symmetric structure lives on the Soergel bimodules which arise from our singular Soergel bimodules, though it is not hard to adapt the definition to the singular Soergel bimodules themselves. As far as we can tell, no connection has yet been made between Lusztig's symmetric structure and the Beilinson–Drinfeld Grassmannian. See also Remark 1.6 below.

1.2. The q-deformation. The 2-categories Rep^{Ω} and Fund^{Ω} admit natural q-deformations Rep_q^{Ω} and Fund_q^{Ω}, describing representations of the quantum group $U_q(\mathfrak{g}^{\vee})$. Let us restrict henceforth to type A, where Fund_q^{Ω} also has a known presentation by generators and relations. The surprising fact is that $\mathcal{SS}\mathrm{Bim}$ also admits a q-deformation $\mathcal{SS}\mathrm{Bim}_q$ compatible with the Soergel Satake equivalence. We have a functor $\mathrm{Fund}_q^{\Omega} \to m\mathcal{SB}\mathcal{SS}\mathrm{Bim}_q$ which is fully faithful onto degree 0 maps, which we call the *quantum algebraic Satake equivalence*. At the moment, there is no known geometric source for this q-deformation of $\mathcal{SS}\mathrm{Bim}$. Gaitsgory [17] also has a notion of a quantum Satake equivalence, but there is currently no connection known between his theory and ours.

We now provide a brief description of the q-deformation of Soergel bimodules, and a number of remarks. There are some technicalities which we ignore in the introduction; §5 has a more accurate discussion.

Let W be a Coxeter group with simple reflections S. In [41] Soergel defines a *reflection faithful representation* of W to be a vector space $\mathfrak{h}$ over a field $\Bbbk$ on which W acts faithfully, such that an element of W acts by a reflection on $\mathfrak{h}$ (it has one eigenvalue -1 and fixes a codimension 1 hyperplane) if and only if it is a reflection in W. To such a representation, Soergel associates a monoidal category of Soergel bimodules, and Williamson [45] a 2-category $\mathcal{SS}\mathrm{Bim}$ of singular Soergel bimodules. The lightning definition is this: consider the coordinate ring $R = \mathrm{Sym}(\mathfrak{h}^*)$ equipped with its action of W. For any finite parabolic subgroup associated to $I \subset S$ one has a subring R^I of invariant polynomials, and when $I \subset J$ one has $R^J \subset R^I$. Singular Soergel bimodules are defined to be the summands of (grading shifts of) iterated tensor products of the induction and restriction bimodules between these various rings R^I.

In [12] it is explained how to generalize this construction beyond reflection faithful representations. A *realization* of W is a representation $\mathfrak{h}$ which is free over an arbitrary commutative ring $\Bbbk$, together with a choice of simple coroots $\Delta^\vee \subset \mathfrak{h}$ and simple roots $\Delta \subset \mathfrak{h}^*$, satisfying some natural conditions. The pairing between Δ and $\Delta^\vee$ is encoded in a *(generalized) Cartan matrix A* with entries $(a_{s,t})_{s,t \in S}$ valued in $\Bbbk$. It need not be the case that Δ and $\Delta^\vee$ form bases for their respective spaces, but when they do, the realization is determined by the Cartan matrix. To any realization, [12] and [10] provide a 2-category analogous to $\mathcal{SS}\mathrm{Bim}$.

To obtain $\mathcal{SS}\mathrm{Bim}_q$ one begins with a q-deformed version of the $\widetilde{A}_n$ Cartan matrix, defined over a base ring $\Bbbk = \mathbb{Z}[q, q^{-1}]$, called the *exotic affine $\mathfrak{sl}_{n+1}$ Cartan matrix*. For $n \geq 2$ we use

$$\begin{pmatrix} 2 & -1 & 0 & \cdots & 0 & -q^{-1} \\ -1 & 2 & -1 & \cdots & 0 & 0 \\ 0 & -1 & 2 & \cdots & & 0 \\ \vdots & \vdots & \vdots & \ddots & -1 & 0 \\ 0 & 0 & & -1 & 2 & -q \\ -q & 0 & 0 & \cdots & -q^{-1} & 2 \end{pmatrix}. \tag{1.1}$$

For $n = 1$, we use

$$\begin{pmatrix} 2 & -(q + q^{-1}) \\ -(q + q^{-1}) & 2 \end{pmatrix}. \tag{1.2}$$

The parameter q can be specialized to a non-zero complex number, so that these exotic matrices also yield one-parameter families of Cartan matrices over $\Bbbk = \mathbb{C}$. These exotic Cartan matrices do not seem to appear in the literature, nor is there yet a satisfactory geometric explanation for them. (However, the corresponding 1-parameter family of representations of the affine Weyl group does appear in the

literature, see the discussion in §5.3). The author came upon them while hunting for a possible quantum Satake equivalence, inspired by the $\mathfrak{sl}_2$ case which is studied in detail in [8]. An *exotic realization* of affine $\mathfrak{sl}_{n+1}$ is a realization of $\widetilde{A}_n$ over $\mathbb{Z}[q, q^{-1}]$ or $\mathbb{C}$ having such a Cartan matrix, and $\mathcal{SS}\mathrm{Bim}_q$ will be the associated 2-category.

Remark 1.3. Suppose that $\Bbbk$ is a field containing $\mathbb{R}$. For any two $s \neq t \in S$, let m_{st} denote the order of st in W. A *standard* realization of W is one with a symmetric Cartan matrix valued in $\mathbb{R}$, satisfying $a_{s,s} = 2$ and $a_{s,t} = -\cos(\frac{\pi}{m_{st}})$ for $s \neq t$. When the Coxeter graph of W is a tree, it is easy to argue that any Cartan matrix is *standardizable*, i.e. conjugate by a diagonal matrix to a standard Cartan matrix. Conjugating by a diagonal matrix corresponds to rescaling Δ and $\Delta^\vee$, and does not alter the 2-category of singular Soergel bimodules.

However, when the Coxeter graph of W is not a tree, such as for $\widetilde{A}_n$, a realization need not be standardizable. The reader can check that for $n \geq 2$ the exotic Cartan matrices above are standardizable if and only if $q = \pm 1$. In particular, when $q \in \mathbb{C}^\times \setminus \mathbb{R}^\times$, the complex realization $\mathfrak{h}$ has no real form. It is not difficult to deduce that the exotic family over $\mathbb{C}$ exhausts the possible complex Cartan matrices for $\widetilde{A}_n$, up to conjugation by diagonal matrices, from the fact that the Dynkin diagram of $\widetilde{A}_n$ has fundamental group isomorphic to $\mathbb{Z}$.

Because Coxeter graphs of affine Weyl groups in other types are trees, it seems unlikely that one could use a similar construction to produce a quantum algebraic Satake equivalence in other types.

Remark 1.4. The exotic Cartan matrix has determinant $2 - q^2 - q^{-2}$. Therefore, outside of the classical cases $q = \pm 1$, Δ and $\Delta^\vee$ are each linearly independent. Let w_1 denote the reflection corresponding to the longest root in the finite Weyl group (the part with the usual Cartan matrix) and let s_0 denote the affine reflection. When q is specialized to a primitive $2m$-th root of unity, the element $(s_0 w_1)$ has finite order m on the span of $\Delta^\vee$. Thus the action of W on the span of $\Delta^\vee$ is not faithful. Whenever a Cartan matrix is nondegenerate, $\mathfrak{h}$ can be decomposed as a direct sum of the span of $\Delta^\vee$ and a trivial representation, so that no exotic realization is faithful at a non-real root of unity. We will discuss this remark further in Section 5.3.

Remark 1.5. Soergel conjectured that for any reflection faithful representation $\mathfrak{h}$ over a field $\Bbbk$ of characteristic zero, the indecomposable Soergel bimodules descend in the Grothendieck group to the Kazhdan–Lusztig basis of (in this case) the affine Hecke algebra. Williamson has shown [45] that when this is true, the classes of the indecomposable singular Soergel bimodules descend to the Kazhdan–Lusztig basis of the affine Hecke algebroid. Soergel's conjecture is proven in [11] for the standard reflection representation, but the arguments also apply to other realizations possessing a real form, and satisfying a positivity property.

Now consider the exotic realization when $n \geq 1$, and set $q \neq \pm 1$ to be a root of unity. Quantum algebraic Satake relates maximally singular Soergel bimodules to the representation theory of $U_q(\mathfrak{sl}_{n+1})$, which is no longer semisimple. The

lack of semisimplicity prevents the indecomposable bimodules from having the same character they have generically. Therefore indecomposable bimodules do not descend to the Kazhdan–Lusztig basis. This does not contradict the results of [11], as the realization has no real form. It also does not contradict the Soergel conjecture, as no exotic realization is reflection faithful at a root of unity.

Remark 1.6. We continue the discussion from Remark 1.2 above. One might ask whether there is a q-deformation of Lusztig's symmetric structure that produces a natural braiding on maximally singular Soergel bimodules, agreeing under the quantum algebraic Satake equivalence with the standard braiding on $U_q(\mathfrak{sl}_n)$-representations. We consider this a very interesting question (the answer did not appear obvious)!

1.3. Addendum: further study of the q-deformation.

As mentioned above, the exotic realization of affine $\mathfrak{sl}_{n+1}$ is something of a mystery. There is currently no known geometric or representation-theoretic explanation for this q-deformed Cartan matrix; it was fabricated by the author solely to make quantum algebraic Satake work. However, since this paper originally appeared, the exotic realization has taken on importance in other ways, and we would like to advocate that it is a fundamental object well worth studying, above and beyond the results of this paper. We briefly discuss two recent applications.

The geometric Satake equivalence is a categorical explanation for a numerical concurrence first observed by Lusztig: that certain multiplicities attached to the Kazhdan–Lusztig basis of (a singular version of) the Hecke algebra in affine type agreed with multiplicities in the category of representations of the Langlands dual Lie group. Geometric Satake solves this mystery by declaring that the categories which govern these multiplicities are in fact equivalent.

A similar numerical concurrence is related to the parabolic Kazhdan–Lusztig basis of the affine Hecke algebra, rather than the singular Kazhdan–Lusztig basis. The multiplicities here are related to the rational representation theory of the corresponding algebraic group in finite characteristic, at least for sufficiently large characteristic. This relationship goes under many names: Lusztig's conjecture, Andersen's conjecture, James' conjecture. The categorical equivalence which would explain this numerical concurrence has appeared (outside of type A, only conjecturally) in recent work of Riche–Williamson [38], which contains an excellent introduction to the topic and many additional references. These same multiplicities are also related to the representation theory of quantum groups at roots of unity. In this case, the numerical equality has been a theorem for some time, due to the work of many: Kazhdan, Lusztig, Kashiwara, Tanisaki, Soergel, etcetera. A quantum Riche–Williamson conjecture would give an equivalence of categories between tilting modules in the trivial block of the quantum group $U_q(\mathfrak{sl}_{n+1})$ at a root of unity, and a parabolic version of the diagrammatic category of Soergel bimodules coming from

the exotic realization of affine $\mathfrak{sl}_{n+1}$. For $\mathfrak{sl}_2$ this was proven already by Andersen–Tubbenhauer [1], to which we refer the reader for a further introduction to this topic. The general case is under investigation.

We were careful to refer to the "diagrammatic category of Soergel bimodules" in the previous paragraph, instead of just the category of Soergel bimodules. This diagrammatic category was constructed in [12], and agrees with Soergel's algebraic definition when Soergel bimodules behave well (e.g. for reflection faithful representations). For arbitrary realizations, it is the diagrammatic category that is the appropriate generalization. The reflection representation of affine $\mathfrak{sl}_{n+1}$ in finite characteristic, and its exotic realization at a root of unity, are certainly not faithful, so that the diagrammatic category is the category whose multplicities agree with those in the affine Hecke algebra, not Soergel's category. However, it is still well worth asking: what does the algebraic category of Soergel bimodules look like? What does it categorify? When q is specialized to a $2m$-th root of unity, the exotic realization factors through the finite complex reflection group $G(m, m, n)$, and Soergel bimodules in this setting appear to be related to this complex reflection group.

An early investigation of this story is currently underway between the author and Benjamin Young. We study the *exotic NilCoxeter algebra*, the algebra generated by Demazure operators when q is specialized to a $2m$-th root of unity. Unlike the case of generic q, this is a finite dimensional algebra, which we can present by generators and relations (at least for small n, so far). It appears to be an entirely new algebra, with an unfamiliar Poincaré polynomial and even a surprising dimension! For example, the exotic NilCoxeter algebra for $G(2, 2, 3)$ is 36 dimensional, even though $G(2, 2, 3) \cong S_4$ only has size 24. We expect these exotic NilCoxeter algebras to be extremely interesting. They have desirable properties: a unique longest element, for instance. This longest element can be used to define a Frobenius extension structure between the polynomial ring of the exotic realization and the subring of polynomials invariant under the action of $G(m, m, n)$. This Frobenius extension structure, in turn, is essential to the study of Soergel bimodules.

1.4. Organization of the paper. In §2 we define a 2-category $\mathrm{Fund}^{\Omega}(\mathfrak{sl}_2)$ by generators and relations, using the language of planar diagrammatics. We similarly define a graded 2-category $m\mathcal{SBS}\mathrm{Bim}(\mathfrak{sl}_2)$. Then we define a 2-functor $\mathrm{Fund}^{\Omega}(\mathfrak{sl}_2) \to m\mathcal{SBS}\mathrm{Bim}(\mathfrak{sl}_2)$, and prove that it is essentially surjective (up to grading shift) and fully faithful (onto degree zero 2-morphisms). In §3, we repeat the same process for $\mathfrak{sl}_3$. These two sections constitute the real mathematical content of this paper (which is quite easy). They are presented in a vacuum, as it were, without any reference to Soergel bimodules, perverse sheaves, or representation theory, and thus are accessible without any background.

In the subsequent chapters we fill in the details, eventually connecting this computational result to geometric Satake. In §4 we give an introduction to the

diagrammatic approach to $\text{Rep}(\mathfrak{g}^\vee)$. We also describe the group Ω and the notion of Ω-graded-monoidal categories and Ω-2-categories. In §5 we give background on singular Soergel bimodules, and state the Soergel Satake equivalence. Between these two chapters, both sides of the algebraic Satake equivalence described in §2 and §3 will be explained.

Finally, in §6 we explain how to connect Soergel Satake to geometric Satake. This takes place in three steps, with two being very straightforward, and the last relying upon the difficult work of Soergel and Härterich.

With the exception of §6, everything is already written with the q-deformation built in. This leads to a host of complications when defining Soergel bimodules. As a result, §5 is not the easiest introduction to Soergel bimodules in general, though it does provide an in-depth introduction to the subtleties involved when dealing with "odd-unbalanced" realizations, such as the exotic affine realization of $\mathfrak{sl}_n$.

This paper is organized with two audiences in mind: the neophyte and the expert. For the neophyte with little to no experience with geometric Satake, we suggest reading the paper in order. We hope that the concrete, combinatorial approach will make it readable, and the hands-on experience of §2 and §3 will help when trying to understand the more abstract approaches. The paper should be entirely accessible until §6. We do assume the reader is familiar with the technology of planar diagrammatics for 2-categories with adjunction: an introduction can be found in [30, Section 2]. The expert who is more interested in the connection from geometric Satake to Soergel Satake and then to algebraic Satake is welcome to skip directly to §6, only backtracking when indicated for several definitions.

2. The diagrammatic approach: $\mathfrak{sl}_2$

Most of the material in this chapter can be found in [8]. We include it as a warm-up for the more complicated case of $\mathfrak{sl}_3$, and $\mathfrak{sl}_n$ in the sequel.

In this chapter, let $\Omega = \{\overline{0}, \overline{1}\}$ denote the group $\mathbb{Z}/2\mathbb{Z}$. Let the set $S = \{b, r\}$ be an Ω-torsor, and assign the elements colors: blue to b and red to r.

We also assign names and colors to proper subsets $I \subsetneq S$. We abusively let r (resp. b) denote the singleton subset $\{r\} \subsetneq S$. Let $\emptyset$ denote the empty set, to which we assign the color white.

Notation 2.1. In this chapter and the next, we will define several 2-categories with similar presentations. The objects will form a finite set, and each object will have an assigned color. There will be a set of generating 1-morphisms, with the property that

for any two colors s_1, s_2 there is at most one generating 1-morphism from s_1 to s_2. Therefore, a 1-morphism can be represented uniquely by the (non-empty) sequence $\underline{x} = s_1 s_2 \ldots s_m$ of colors through which it passes (though not all sequences of colors are permitted). We read 1-morphisms from right to left, so that $\underline{x}$ has source s_m and target s_1. The identity 1-morphism of an object s is thus also denoted s. We represent a composition of 1-morphisms diagrammatically as a sequence of dots on the line, separating regions of different colors.

Example 2.2. The 1-morphism $brbrbr$ in the 2-category Fund_q^Ω below:

For the rest of this chapter we fix a $\mathbb{Z}[q^{\pm 1}]$-algebra $\Bbbk$. In order to connect these results to representation theory we will set $\Bbbk = \mathbb{Q}(q)$, or $\Bbbk = \mathbb{C}$ with $q = 1$. However, the algebraic Satake equivalence is defined in more generality.

2.1. Webs. Background on this material can be found in [19, 44]. More specifics on this particular 2-colored version can be found in [9, Section 2] or [8, Section 4].

Definition 2.3. Let $\text{Fund}_q^\Omega = \text{Fund}_q^\Omega(\mathfrak{sl}_2)$ (also known as the 2-colored Temperley–Lieb category) be the $\Bbbk$-linear 2-category defined as follows, using Notation 2.1. It has objects $S = \{r, b\}$, and has generating 1-morphisms rb and br. Thus Example 2.2 gives a 1-morphism in Fund_q^Ω.

The 2-morphisms are generated by colored cups and caps. More precisely, there is a cap map $brb \to b$ and a cup map $b \to brb$, as pictured below, as well as the corresponding maps with the colors switched.

There are two types of relations, which also hold with the colors switched. Recall that $[2] = q + q^{-1}$.

The Isotopy relation:

 (2.1)

The Circle relation:

 (2.2)

This ends the definition.

The existence of cups and caps, together with (2.1), is equivalent to the statement that the 1-morphisms rb and br are biadjoint. There is an action of Ω on Fund^Ω which permutes the colors.

Definition 2.4. A *crossingless matching* is an isotopy class of colored 1-manifold with boundary embedded in the planar strip, without any closed components, providing a matching of boundary points.

Example 2.5. A crossingless matching in $\mathrm{Hom}(brbrbrbrbrb, brbrbrbrb)$:
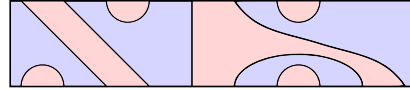
Claim 2.6. *Each morphism space* $\mathrm{Hom}(\underline{x}, \underline{y})$ *has a basis over* $\Bbbk$ *given by crossingless matchings.*

Now we define another 2-category.

Definition 2.7. Let $\mathrm{Rep}_q^\Omega = \mathrm{Rep}_q^\Omega(\mathfrak{sl}_2)$ denote the following full sub-2-category of **Cat**. The objects are in bijection with Ω, where $\overline{0}$ is identified with the category of even $U_q(\mathfrak{sl}_2)$ representations, and $\overline{1}$ with odd $U_q(\mathfrak{sl}_2)$ representations. The 1-morphisms are (functors arising from) tensor products with finite dimensional $U_q(\mathfrak{sl}_2)$ representations, and the 2-morphisms are all natural transformations (i.e. $U_q(\mathfrak{sl}_2)$-morphisms between tensor products).

Now we fix an identification of Ω with S. There is a 2-functor from Fund_q^Ω to Rep_q^Ω which on objects sends S to Ω. It sends both 1-morphisms rb or br to the tensor product with the standard representation V. The cups and caps are sent to inclusions and projections respectively between $V \otimes V$ and the trivial summand $\Lambda^2 V$. These inclusions and projections are unique up to scalar; one chooses the scalars so that (2.1) holds. For any such choice of scalars, it turns out that (2.2) will also hold.

Claim 2.8. *Suppose that* $\Bbbk = \mathbb{Q}(q)$. *For any identification of* S *with* Ω, *this 2-functor is well defined and fully faithful.*

The proofs of these claims are well known, and are reasonable exercises for the uninitiated. The result also holds when $\Bbbk = \mathbb{C}$, $q = 1$, and $U_q(\mathfrak{sl}_2)$ is replaced by $\mathfrak{sl}_2$. A similar claim holds when $\Bbbk = \mathbb{Z}[q, q^{-1}]$, for the correct integral form of $U_q(\mathfrak{sl}_2)$.

2.2. Singular Soergel diagrams.
For another introduction to this material, see [8, Section 5]. We now make a mild assumption on $\Bbbk$: that the ideal generated by 2 and [2] is the unit ideal (c.f. Demazure Surjectivity in [8, Section 3.3]).

Definition 2.9. Let $\mathcal{SBSBim}_q = \mathcal{SBSBim}_q(\mathfrak{sl}_2)$ be the graded $\Bbbk$-linear 2-category defined as follows, using Notation 2.1. The objects are proper subsets $I \subsetneq S$, and the generating 1-morphisms are $\{r\emptyset, \emptyset r, b\emptyset, \emptyset b\}$.

Example 2.10. The 1-morphism $b\emptyset r\emptyset r\emptyset b\emptyset$:

In the drawing above, we placed an orientation on the dots consistent with the rule that the right-hand color contains the left-hand color. This orientation is redundant

information, but helps to make 2-morphisms easier to describe. Similarly, we will color the strands in our pictures by the color which is added or subtracted by the 1-morphism.

The 2-morphisms are generated by oriented cups and caps (again, the orientation itself is redundant). There are cap maps $\emptyset b\emptyset \to \emptyset$ and $b\emptyset b \to b$, as well as cup maps in the other direction, as pictured below. There are also maps with blue and red switched.
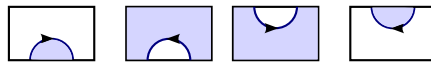
We place a grading on the 2-morphisms, where clockwise cups and caps have degree $+1$, and anticlockwise cups and caps have degree -1. The relations between 2-morphisms are listed below (with the color-switched versions assumed). They are all homogeneous.

The Isotopy relation:

$$(2.3)$$

The Empty Circle relation:

 $= 0$ $$(2.4)$$

The Cartan relations:

 $= 2$ $$(2.5a)$$

 $= -[2]$ $$(2.5b)$$

The Forcing relations:


$$\qquad (2.6a)$$

$$\qquad (2.6b)$$

This ends the definition.

The existence of cups and caps, together with (2.3), is equivalent to the statement that the 1-morphisms $b\emptyset$ and $\emptyset b$ are biadjoint. There is an action of Ω on Fund^{Ω} which permutes the colors; however, this is special to the $\mathfrak{sl}_2$ case.

2.3. The equivalence.

Definition 2.11. Let $m\mathcal{SBSBim}_q$ denote the full sub-2-category of $\mathcal{SBSBim}_q$ whose objects are $\{r, b\}$. In other words, we allow 1-morphisms and 2-morphisms which contain the color white, but their right-hand and left-hand colors must be either red or blue.

Definition 2.12. Let $\mathcal{F}$ be the 2-functor $\mathrm{Fund}_q^{\Omega} \to m\mathcal{SBSBim}_q$ defined as follows. On objects, it sends red to red and blue to blue. On 1-morphisms, it sends $rb \mapsto r\emptyset b$ and $br \mapsto b\emptyset r$. On 2-morphisms, it acts as below. Visually, the map on 2-morphisms takes a crossingless matching and widens each strand into a white region, with its boundary oriented anticlockwise.
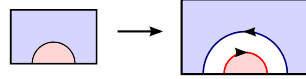
Claim 2.13. *The 2-functor $\mathcal{F}$ is well defined, and its image consists of degree 0 maps.*

Proof. The isotopy relation in Fund_q^{Ω} follows from the isotopy relations in $\mathcal{SBSBim}_q$. The circle relation in Fund_q^{Ω} follows from the Cartan relations in $\mathcal{SBSBim}_q$. $\square$

Definition 2.14. A 2-functor $\mathcal{G}$ from a $\Bbbk$-linear 2-category $\mathcal{C}$ to a $\Bbbk$-linear graded 2-category $\mathcal{D}$ is a *degree-zero equivalence* if the following properties hold.

- $\mathcal{G}$ induces a bijection between the objects.
- $\mathcal{G}$ induces an isomorphism $\mathrm{Hom}_{\mathcal{C}}(\underline{x}, \underline{y}) \to \mathrm{Hom}_{\mathcal{D}}^0(\mathcal{G}(\underline{x}), \mathcal{G}(\underline{y}))$ to the morphisms of degree 0.
- $\mathrm{Hom}_{\mathcal{D}}^k(\mathcal{G}(\underline{x}), \mathcal{G}(\underline{y})) = 0$ for all $k < 0$.

- Every 1-morphism in $\mathcal{D}$ is isomorphic to a direct sum of grading shifts of $\mathcal{G}(x)$ for some 1-morphism x in the Karoubi envelope $\mathbf{Kar}(\mathcal{C})$; i.e. $\mathcal{G}$ is *essentially surjective up to grading shift*.

Theorem 2.15. *Suppose that the ideal generated by 2 and $[2]$ in $\Bbbk$ is the unit ideal. Then the 2-functor $\mathcal{F}$ is a degree-zero equivalence.*

This is one of the main results of [8, Section 4], where its proof is dispersed throughout many pages. We will sketch the proof after some more material.

2.4. Singular Soergel diagrams and polynomials. Let us introduce some important notation for certain 2-morphisms in $\mathcal{SBS}\mathrm{Bim}_q$. Let α_b (resp. α_r) denote a clockwise blue (resp. red) circle. These are endomorphisms of the identity 1-morphism $\emptyset$. Therefore we have a homomorphism from the polynomial ring $R = \Bbbk[\alpha_b, \alpha_r]$ to $\mathrm{End}(\emptyset)$, which is a graded homomorphism provided we set the degree of α_b to be $+2$. We use this homomorphism to replace disjoint unions of circles by "boxes" (labeled with polynomials).

$$
\text{(diagram of two blue circles and one red circle)} \quad = \quad \boxed{\alpha_b^2 \alpha_r}
$$

We define endomorphisms b and r of R by the formulas

$$
\begin{array}{llr}
b(\alpha_b) = -\alpha_b, & b(\alpha_r) = \alpha_r + [2]\alpha_b & \text{(2.7a)} \\
r(\alpha_b) = \alpha_b + [2]\alpha_r, & r(\alpha_r) = -\alpha_r. & \text{(2.7b)}
\end{array}
$$

In particular, $b^2 = r^2 = 1$. This action on linear terms is a q-deformation of the reflection representation of the affine Weyl group of $\mathfrak{sl}_2$.

We can define subrings R^b and R^r of invariant polynomials. In [8, Section 5.2.1] it is shown how to define a map from R^b to the endomorphism ring of the identity 1-morphism b. Thus we can place a box labeled by a polynomial within a blue region as well, provided the polynomial is invariant under the operator b.

One crucial fact about these rings is that $R^b \subset R$ is a *Frobenius extension*. Roughly, this means that R is free as an R^b-module, and that the functors of induction and restriction between R-modules and R^b-modules are biadjoint. Part of the data of a Frobenius extension is an R^b-linear map $\partial_b \colon R \to R^b$, in this case defined by the formula

$$
\partial_b(f) = \frac{f - b(f)}{\alpha_b}.
$$

Moreover, R has dual bases (as an R^b-module) $\{a_i\}$ and $\{a_i^*\}$ with respect to ∂_b, in the sense that

$$
\partial_b(a_i a_j^*) = \delta_{ij}.
$$

(These dual bases are not part of the data of a Frobenius extension.) There is a canonical "coproduct" element $\Delta_b = \sum_i a_i \otimes a_i^*$ living in $R \otimes_{R^b} R$, independent of the choice of dual bases. We often use Sweedler notation $\Delta_b = \Delta_{b(1)} \otimes \Delta_{b(2)}$. One invariant of a Frobenius extension is the product $\mu_b = \Delta_{(1)}\Delta_{(2)} = \sum_i a_i a_i^* \in R$, which in this case is equal to α_b. We provide more detail on Frobenius extensions in §5.4.

The values of $\partial_s(\alpha_t)$ for $s, t \in S$ are encoded in a *Cartan matrix*, yielding precisely the q-deformed affine Cartan matrix of $\mathfrak{sl}_2$, given in (1.2). These values are also evident from the Cartan relation (2.5).

Note that ∂_b is surjective, because it is R^b-linear and $\partial_b(a_i a_i^*) = 1$. This relied on our assumption that the ideal in $\Bbbk$ generated by 2 and [2] contains the unit. After all, $\partial_b(\alpha_b) = 2$ and $\partial_b(\alpha_r) = -[2]$, so that this ideal comprises all the scalars in the image of ∂_b.

Now we give an alternate description of $\mathcal{SBSBim}_q$.

Definition 2.16. Let $\mathcal{SBSBim}_q$ denote the 2-category with objects and 1-morphisms as in Definition 2.9. The 2-morphisms will be generated by cups and caps, as well as boxes. A box appearing in a region labeled $I \subset S$ is labeled by a polynomial in R^I, and has degree equal to the degree of that polynomial (with the convention $\deg \alpha_b = 2$). Boxes add and multiply as polynomials do. In addition to the isotopy relation (2.3) we have the following relations.

$$
\boxed{\,\bigcirc\,} = \boxed{\alpha_b} \qquad \boxed{\,\bigcirc\,} = \boxed{\alpha_r} \,. \tag{2.8a}
$$

$$
\boxed{\,\boxed{f}\,} = \boxed{\,\boxed{f}\,} \quad \text{when } f \in R^b. \tag{2.8b}
$$

$$
\boxed{\,\boxed{f}\,} = \boxed{\partial_b(f)} \,. \tag{2.8c}
$$

$$
\boxed{\,\bigcap\,\bigcup\,} = \boxed{\,\boxed{\Delta_b}\,} \,. \tag{2.8d}
$$

In this last equation, Δ_b is meant to represent the action of $R \otimes_{R^b} R$ inside $\mathrm{End}(\emptyset b \emptyset)$, by placing $\Delta_{(1)}$ in the white region on the left, and $\Delta_{(2)}$ in the white region on the right. Equation (2.8b) guarantees that this action is well defined.

In [8] it is shown that this definition of $\mathcal{SBSBim}_q$ is equivalent to the one in Definition 2.9. It is also a reasonable exercise for the reader.

Now we define the singular Soergel 2-category.

Definition 2.17. Let $\mathcal{S}\mathbb{S}\mathrm{Bim}_q = \mathcal{S}\mathbb{S}\mathrm{Bim}_q(\mathfrak{sl}_2)$ denote the following full sub-2-category of **Cat**. The objects are proper subsets of S, where $I \subsetneq S$ is identified with the category of graded R^I-modules (using the notation $R^\emptyset = R$). The 1-morphisms are summands of (grading shifts of) iterated induction and restriction functors, between R^b- or R^r-modules and R-modules. The 2-morphisms are all natural transformations (i.e. homogeneous bimodule maps).

There is a 2-functor from $\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim}_q$ to $\mathcal{S}\mathbb{S}\mathrm{Bim}_q$ which is the identity on objects. It sends $b\emptyset$ to the restriction functor (with a grading shift by 1), and $\emptyset b$ to the induction functor. The cups and caps are sent to units and counits of adjunction.

Claim 2.18. *Suppose that $\Bbbk = \mathbb{Q}(q)$, or that $q = 1$ and $\Bbbk$ is a field of characteristic $\neq 2$. Then this 2-functor is well defined and fully faithful.*

See [8, Section 5.2.2 and Corollary 5.30] for more details.

2.5. Sketch of proof. Let us quickly sketch the proof of Theorem 2.15 given in [8, Section 4]. There are three main facets of the argument. One uses idempotent decompositions (or the algebraic theory of Frobenius extensions) to prove essential surjectivity up to grading shift. One uses considerations on the Grothendieck group to determine the graded dimension of morphism spaces. Finally, one uses a diagrammatic argument to show that $\mathcal{F}$ is full.

The fact that R is free over R^b says that the R^b-bimodule R should split into copies of the R^b-bimodule R^b. In $\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim}_q$, this amounts to the fact that $b\emptyset b$ is isomorphic to several copies of b (with grading shifts). One way to prove this splitting is to construct an idempotent decomposition of the identity of $b\emptyset b$. This is obtained by rotating (2.8d) by 90 degrees. When 2 is invertible, this can be realized more concretely by rotating (2.6a) by 90 degrees and dividing by 2.

Any 1-morphism in $m\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim}_q$ which alternates between red and blue, i.e. $r\emptyset b\emptyset r\emptyset \ldots \emptyset b$, is clearly in the image of $\mathcal{F}$. Any other 1-morphism must contain either $b\emptyset b$ or $r\emptyset r$. By the above paragraph, one can replace $b\emptyset b$ with b after taking summands, and similarly with $r\emptyset r$. Therefore, any object is a direct sum of alternating 1-morphisms, proving the essential surjectivity up to grading shift.

Using standard representation theory or combinatorics, one knows the dimension of the space of crossingless matchings between any two colored sequences $\underline{x}$ and $\underline{y}$ in Fund_q^Ω. Using the fact that $\mathcal{S}\mathbb{S}\mathrm{Bim}_q$ categorifies the Hecke algebroid with its standard trace, one knows the (graded) dimension of morphism spaces between any two objects in $m\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim}_q$. From this one can deduce that there are no morphisms of negative degree between alternating 1-morphisms, and that the maps in degree zero are the correct size for $\mathcal{F}$ to induce isomorphisms. (Note that the proof in [8] does not use precisely the same argument: it uses Theorem 2.15 to prove that $\mathcal{S}\mathbb{S}\mathrm{Bim}_q$ categorifies the Hecke algebra, not the other way around. Instead, one proves computationally that $\mathcal{F}$ is faithful.)

Thus we need only show that $\mathcal{F}$ is full onto maps of minimal degree. This is a purely diagrammatic argument, the details of which are not particularly relevant. First one proves that any closed diagram (i.e. a diagram in the endomorphism ring of an identity 1-morphism) reduces to a polynomial (in R or R^b or R^r, depending on the color of the boundary). Obviously any diagram with a non-trivial polynomial is not of minimal degree, so we can assume our diagram has no closed subdiagrams. Now we wish to show that blue regions are only adjacent (crossing white regions) to red regions, which would imply that one could deformation retract the white regions into a crossingless matching, and the diagram is in the image of $\mathcal{F}$. However, if there are two separate blue regions separated by a white region, then "fusing" them
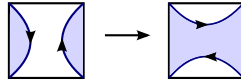
would yield a morphism of lower degree. (This argument is analogous to the one made in [8], which passes through a separate diagrammatic calculus to achieve the same result.)

3. The diagrammatic approach: $\mathfrak{sl}_3$

In this chapter, let $\Omega = \{\overline{0}, \overline{1}, \overline{2}\}$ denote the group $\mathbb{Z}/3\mathbb{Z}$. Let the set $S = \{b, r, y\}$ be an Ω-torsor, and assign the elements colors: blue to b, red to r, and yellow to y. The action of $+1 \in \Omega$ acts in alphabetical order to send b to r, r to y, and y to b. We call this *color rotation*.

We also assign names and colors to proper subsets $I \subsetneq S$. We abusively let r (resp. b, g) denote the singleton subset $\{r\} \subsetneq S$. Let $\emptyset$ denote the empty set, to which we assign the color white. To pairs we associate the natural compound color: purple p to $\{r, b\}$, green g to $\{b, y\}$, and orange o to $\{r, y\}$. Therefore, the action of $+1 \in \Omega$ sends p to o, o to g, and g to p (reverse alphabetical order, unfortunately).

For the rest of this chapter we fix a $\mathbb{Z}[q^{\pm 1}]$-algebra $\Bbbk$. In order to connect these results to representation theory we will set $\Bbbk = \mathbb{Q}(q)$, or $\Bbbk = \mathbb{C}$ with $q = 1$. However, the algebraic Satake equivalence is defined in more generality.

### 3.1. Webs. Background on this material can be found in [29] or [5].

Definition 3.1. Let $\mathrm{Fund}_q^\Omega = \mathrm{Fund}_q^\Omega(\mathfrak{sl}_3)$ (also known as the 2-category of colored $\mathfrak{sl}_3$ webs) be the $\Bbbk$-linear 2-category defined as follows, using Notation 2.1. It has objects $\{o, g, p\}$ (i.e. $I \subset S$ of size 2), and has generating 1-morphisms from any color to any different color.

Example 3.2. The 1-morphism $pogogp$: .

The 2-morphisms are generated by colored cups and caps, and by trivalent vertices. More precisely, there is a cap map $sts \to s$ and a cup map $s \to sts$ for any two

colors s, t. There is a trivalent vertex $ogpo \rightarrow r$, and another trivalent vertex $opgo \rightarrow r$, as pictured below.
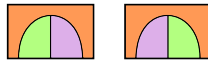
The relations are as follows; they hold for every valid coloring. Recall that $[2] = q + q^{-1}$ and $[3] = q^2 + 1 + q^{-2}$.

The Isotopy relations:

                                                    (3.1)

                                                    (3.2)

The Circle relation:

                                                    (3.3)

The Bigon relation:

                                                    (3.4)

The Square relation:

                                                    (3.5)

This ends the definition.

The existence of cups and caps, together with (3.1), is equivalent to the statement that the 1-morphisms *po* and *op* are biadjoint. The cyclicity relation (3.2) states that all 2-morphisms are cyclic with respect to these biadjunctions. Together, this allows one to unambiguously define the various trivalent vertices appearing above as rotations of the generating trivalent vertices. There is an action of Ω on Fund$^{\Omega}$ by color rotation.

Definition 3.3. A *non-elliptic web* is an isotopy class of colored trivalent graph with boundary embedded in the planar strip, which may have interior hexagons, octagons, et cetera, but may have no interior squares or bigons.

Example 3.4. A non-elliptic web in $\mathrm{Hom}(opgpgpopogo, opgpogopgo)$:
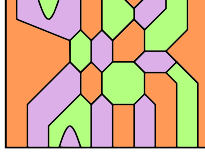
Claim 3.5. *Each morphism space* $\mathrm{Hom}(\underline{x}, \underline{y})$ *has a basis over* $\Bbbk$ *given by non-elliptic webs.*

Now we define another 2-category. Note that there is a bijection between Ω and the set of central characters of SL_3.

Definition 3.6. Let $\mathrm{Rep}_q^\Omega = \mathrm{Rep}_q^\Omega(\mathfrak{sl}_3)$ denote the following full sub-2-category of **Cat**. The objects are in bijection with Ω, where each $\xi \in \Omega$ is identified with the category of $U_q(\mathfrak{sl}_3)$ representations having central character Ω. The 1-morphisms are (functors arising from) tensor products with finite dimensional $U_q(\mathfrak{sl}_3)$ representations, and the 2-morphisms are all natural transformations (i.e. $U_q(\mathfrak{sl}_3)$-morphisms between tensor products).

Now we fix an identification of Ω with $\{o, g, p\}$. Say we identify o with the trivial character, g with the character of V_{ω_1}, and p with the character of V_{ω_2}. There is a 2-functor from Fund_q^Ω to Rep_q^Ω which on objects sends $\{o, g, p\}$ to Ω. The 1-morphisms go, pg, and op all correspond to tensoring with $V = V_{\omega_1}$, while their biadjoints correspond to tensoring with $V^* = V_{\omega_2}$. The cups and caps correspond to inclusion and projection between $V \otimes V^*$ and the trivial summand. The trivalent vertex $opgo \to o$ corresponds to projection from $V^{\otimes 3}$ to the trivial summand $\Lambda^3 V$. These projections and inclusions are unique up to scalar; one chooses the scalars so that (2.1) and (3.4) hold. For any such choice of scalars, it turns out that (3.3) and (3.5) also hold.

Claim 3.7. *Suppose that* $\Bbbk = \mathbb{Q}(q)$. *For any identification of S with Ω, this 2-functor is well defined and fully faithful.*

One again, the proofs of these claims are well known (also see [29]), and are worthwhile exercises. The result also holds when $\Bbbk = \mathbb{C}$, $q = 1$, and $U_q(\mathfrak{sl}_3)$ is replaced by $\mathfrak{sl}_3$. A similar claim holds when $\Bbbk = \mathbb{Z}[q, q^{-1}]$, for the correct integral form of $U_q(\mathfrak{sl}_3)$.

3.2. Polynomials. Our next goal is to define a 2-category $\mathcal{SBS}\mathrm{Bim}_q(\mathfrak{sl}_3)$, in similar fashion to $\mathcal{SBS}\mathrm{Bim}_q(\mathfrak{sl}_2)$. However, because polynomials are slightly more complex for $\mathfrak{sl}_3$ than $\mathfrak{sl}_2$, we choose to discuss them first in this section. Again, we need a mild assumption on $\Bbbk$ in order for the inclusion of an invariant subring to be a Frobenius extension. In this case, the assumption that 3 is invertible will suffice.

Let $R = \Bbbk[\alpha_r, \alpha_b, \alpha_y]$ be a polynomial ring in 3 variables, graded so that linear terms have degree $+2$. We define an action of three operators b, r, y on R as follows:

$$b(\alpha_b) = -\alpha_b, \qquad b(\alpha_r) = \alpha_r + \alpha_b, \qquad b(\alpha_y) = a_y + q^{-1}\alpha_b, \quad \text{(3.6a)}$$
$$r(\alpha_b) = \alpha_b + \alpha_r, \qquad r(\alpha_r) = -\alpha_r, \qquad r(\alpha_y) = \alpha_y + q\alpha_r, \quad \text{(3.6b)}$$
$$y(\alpha_b) = \alpha_b + q\alpha_y, \qquad y(\alpha_r) = \alpha_r + q^{-1}\alpha_y, \qquad y(\alpha_y) = -\alpha_y. \quad \text{(3.6c)}$$

In particular, each operator is an involution, and the reader can check that $(br)^3 = (by)^3 = (ry)^3 = 1$. This is a q-deformation of the reflection representation of the affine Weyl group of $\mathfrak{sl}_3$. Note that, unlike the case of $\mathfrak{sl}_2$, there is no color symmetry in this action, so the colors are not interchangeable.

Let $I \subsetneq S$, and let R^I denote the subring of polynomials invariant under all reflections in I. For $s \in S$ we let $\partial_s \colon R \to R^s$ be defined as before, via the formula

$$\partial_s(f) = \frac{f - s(f)}{\alpha_s}. \tag{3.7}$$

This entirely describes the action of the reflection on the linear terms. The values of $\partial_s(\alpha_t)$ for $s, t \in S$ can be encoded in the q-deformed affine $\mathfrak{sl}_3$ Cartan matrix.

$$\begin{pmatrix} 2 & -1 & -q^{-1} \\ -1 & 2 & -q \\ -q & -q^{-1} & 2 \end{pmatrix}. \tag{3.8}$$

The reader should confirm that ∂_s and ∂_t do not satisfy the braid relation! In particular, one has

$$\partial_b \partial_r \partial_b = \partial_r \partial_b \partial_r, \tag{3.9a}$$
$$q\partial_b \partial_y \partial_b = \partial_y \partial_b \partial_y, \tag{3.9b}$$
$$q^{-1}\partial_r \partial_y \partial_r = \partial_y \partial_r \partial_y. \tag{3.9c}$$

Either $\partial_b \partial_y \partial_b$ or $\partial_y \partial_b \partial_y$ will define a Frobenius structure map $R \to R^g$, but there is no particular reason a priori to choose one over the other. The choice we make will lead to some convenient formulas.

Definition 3.8. For each $I \subsetneq S$ we will define a set of *positive roots* Φ_I, in such a way that $\Phi_I \subset \Phi_J$ whenever $I \subset J$. We have $\Phi_\emptyset = \emptyset$, and $\Phi_s = \{\alpha_s\}$ for $s \in S$. For doubletons, we have

$$\Phi_p = \{\alpha_b, \alpha_r, b(\alpha_r) = \alpha_r + \alpha_b\}, \quad \Phi_g = \{\alpha_b, \alpha_y, b(\alpha_y) = \alpha_y + q^{-1}\alpha_b\},$$
$$\Phi_o = \{\alpha_r, \alpha_y, r(\alpha_y) = \alpha_y + q\alpha_r\}.$$

We call the union of these positive roots the *finite positive roots* of S. Let μ_J^I denote the product of the roots in Φ_J which are not in Φ_I.

For later reference, we let $c_{s,t}$ denote the coefficient of α_s in the third root of $\Phi_{s,t}$. Thus $c_{r,b} = c_{b,r} = c_{y,b} = c_{y,r} = 1$, while $c_{b,y} = q^{-1}$ and $c_{r,y} = q$.

Definition 3.9. For each $I \subsetneq J \subsetneq S$ we will define a map $\partial_J^I \colon R^I \to R^J$. These maps are *compatible* in that $\partial_K^J \circ \partial_J^I = \partial_K^I$ whenever $I \subsetneq J \subsetneq K \subsetneq S$. In particular, we need only define the map ∂_J^I when I and J differ by a single element; this will define all ∂_K^I uniquely by the comptability requirement, once we check that this is consistently defined. For each $s \in S$ we choose $\partial_s^\emptyset = \partial_s$ as defined above. Then we choose

$$\partial_p^b = \partial_b \partial_r, \qquad \partial_p^r = \partial_r \partial_b, \qquad \partial_g^b = q \partial_b \partial_y,$$
$$\partial_g^y = \partial_y \partial_b, \qquad \partial_o^r = q^{-1} \partial_r \partial_y, \qquad \partial_o^y = \partial_y \partial_r.$$

The compatibility relations follow from (3.9).

Note that $\partial_s \partial_t = c_{s,t} \partial_{s,t}^s$.

Claim 3.10. *The maps $\partial_J^I \colon R^I \to R^J$ equip the ring extension $R^J \subset R^I$ with the structure of a Frobenius extension. For any choice of dual bases, one has $\mu_J^I = \sum_i a_i a_i^*$. In particular, this implies that $\partial_J^I(\mu_J^I)$ is an integer, equal to the rank of the extension.*

This claim is not difficult to prove, by explicitly constructing dual bases. We invite the reader to check that whenever $s \in S$ and I is a doubleton containing s, one has $\partial_I^s(\mu_I^s) = 3$. When performing such computations, one should make use of the *twisted Leibniz rule*

$$\partial_s(fg) = \partial_s(f)g + s(f)\partial_s(g).$$

Once again, the fact that one has a Frobenius extension relies upon the fact that ∂_I^s is surjective, which was implied by our assumption that 3 is invertible.

For more information on choosing structure maps for collections of Frobenius extensions, see Sections §5.4 and §5.5.

Before moving on to the diagrammatic definition of singular Bott–Samelson bimodules, we state some remarkable identities which will play a role in the algebraic Satake equivalence, and which will generalize in some sense to $\mathfrak{sl}_n$.

Claim 3.11. *The following identities hold. They are defined in a color-symmetric way, but each case should be checked individually.*

- *Let $I, J \subset S$ be two distinct doubletons, with intersection s. Then*

$$\partial_I^s(\mu_J^s) = [3] = q^2 + 1 + q^{-2}. \tag{3.10}$$

- *Let $I = \{s, t\}$ be a doubleton with complement $u \in S$. Let $\mu_I^{s,t}$ denote the product of the roots in Φ_I which are not in Φ_s or Φ_t; this is none other than the third root in Φ_I, either $b(\alpha_r)$, $b(\alpha_y)$, or $r(\alpha_y)$. Then*

$$\partial_u(\mu_I^{s,t}) = -[2] = -(q + q^{-1}). \tag{3.11}$$

- *For any $s, t, u \in S$ distinct we have*

$$c_{s,u} c_{t,u} = 1. \tag{3.12}$$

Proof. These are straightforward exercises.   □

3.3. Singular Soergel diagrams. We now define $\mathcal{SBSBim}_q(\mathfrak{sl}_3)$ analogously to the $\mathfrak{sl}_2$ case in Definition 2.16.

Definition 3.12. Let $\mathcal{SBSBim}_q = \mathcal{SBSBim}_q(\mathfrak{sl}_3)$ be the graded $\Bbbk$-linear 2-category defined as follows, using Notation 2.1. The objects are proper subsets $I \subsetneq S$. There is a generating 1-morphism between two colors if and only if one contains the other and their sizes differ by one.

Example 3.13. The 1-morphism $b\emptyset roygb\emptyset$: .

In the drawing above, we placed an orientation on the dots consistent with the rule that the right-hand color contains the left-hand color. This orientation is redundant information, but helps to make 2-morphisms easier to describe. Similarly, we will color the strands in our pictures by the color which is added or subtracted by the 1-morphism.

The 2-morphisms are generated by oriented cups and caps (again, the orientation itself is redundant) and boxes, as well as crossings. Within a region colored I one can place a box labeled by a polynomial in R^I. There are crossing maps $\emptyset bp \to \emptyset rp$ as pictured below. More generally, there are crossing maps $\emptyset sI \to \emptyset tI$ for any doubleton $I = \{s, t\} \subset S$.
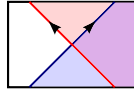
We place a grading on the 2-morphisms as follows. A clockwise cup between $\emptyset$ and $s \in S$ has degree $+1$, while a clockwise cup between s and $\{s, t\}$ has degree $+2$. Anticlockwise cups have the opposite degree. Upward-oriented crossings have degree 0. Note however that one can use cups and caps to draw a crossing with the arrows pointing either down or sideways. By computation, a downward-oriented crossing also has degree 0, while a sideways-oriented crossing has degree $+1$. Boxes have degree given by the degree of the polynomial.

The (homogeneous) relations given below are in terms of the Frobenius structures defined above. As a result, they will not be invariant under color rotation! For illustrative purposes, however, we choose specific colors to exemplify relations.

One has the isotopy relations as well as all the relations of (2.8), with the obvious modification

 (3.13)

We also have the analogous relations for the other Frobenius extensions.

$$\partial_p^b(f) \tag{3.14a}$$

$$\mu_p^b \tag{3.14b}$$

 when $f \in R^p$ (3.14c)

$$\Delta_p^b \tag{3.14d}$$

We have a number of relations analogous to the Reidemeister II move.

$$\tag{3.15a}$$

$$\partial\Delta_p \tag{3.15b}$$

$$\mu_p^{r,b} \tag{3.15c}$$

These relations require some explanation. The symbol $\partial\Delta_p$ in (3.15b) represents the element

$$\partial\Delta_p = \partial_b(\Delta_{p(1)}^r) \otimes \Delta_{p(2)}^r = \Delta_{p(1)}^b \otimes \partial_r(\Delta_{p(2)}^b) \in R^b \otimes_{R^p} R^r.$$

Meanwhile, the symbol $\mu_p^{r,b}$ is the product of the roots in Φ_p which are not in either Φ_b or Φ_r, which in this particular case is the linear term $b(\alpha_r)$.

We pause to note a consequence of the preceding relations, which is akin to (3.15b).

$$\Delta_{p(1)}^b \qquad \partial_r(f\Delta_{p(2)}^b) \tag{3.15d}$$

The relations above are actually completely general, for any compatible system of Frobenius extensions. See [13] for more details. The final relations below are specific to our particular circumstance.

$$\tag{3.16}$$

To compare these relations to our choice of Frobenius structure, the reader should note the following fact. Take the RHS of each equation, and apply a cap to each of the four walls of the diagram, to obtain diagrams with 3 bubbles each. The cubic polynomial obtained is precisely μ_I for the doubleton color I appearing on the LHS. Said another way, the coefficient of the diagram with colors s, t and more connected components colored s is $c_{s,t}$.

This ends the definition. The diagrams which arise in this fashion are called *singular Soergel diagrams*.

There is also a "boxless" definition analogous to Definition 2.9, though it is less convenient. In this alternate definition, the 2-morphisms are given purely in terms of oriented colored 1-manifolds, and boxes containing polynomials are defined in terms of these diagrams. One reason we find the definition above to be more convenient is that boxes should be thought of as "colorless" for the purpose of many proofs.

Each relation above used at most two colors, just as the relations for $\mathfrak{sl}_2$ (in the version of Definition 2.16) used at most one color. This follows the general principle that relations amongst Soergel bimodules are governed by the relations which appear for finite parabolic subgroups.

Let us pause to note one consequence of these relations.

Claim 3.14. *For any two colors, we have an equality exemplified as below.*

$$\tag{3.17}$$

The scalar $c_{r,b}$ was given in Definition 3.8.

Proof. We summarize the proof of [8, Claim 6.5]. By applying (3.15b) to the right half of the LHS, and then sliding polynomials and using (3.15a), one obtains the endomorphism of rpr given in Sweedler notation by $\partial_r \partial_b \Delta^r_{p(1)} \otimes \Delta^r_{p(2)} = c_{r,b} \partial^r_p \Delta^r_{p(1)} \otimes \Delta^r_{p(2)} \in R^r \otimes_{R^p} R^r$. For any basis of R^r over R^p, only a single term can survive the application of ∂^r_p for degree reasons, and this is the term which pairs against the unit in the dual basis. Thus $\partial^r_p \Delta^r_{p(1)} \otimes \Delta^r_{p(2)} = 1 \otimes 1$. □

Now we define the singular Soergel 2-category.

Definition 3.15. Let $\mathcal{SSBim}_q = \mathcal{SSBim}_q(\mathfrak{sl}_3)$ denote the following full sub-2-category of **Cat**. The objects are proper subsets of S, where $I \subsetneq S$ is identified with the category of graded R^I-modules (using the notation $R^\emptyset = R$). The 1-morphisms are summands of (grading shifts of) iterated induction and restriction functors, between R^I-modules and R^J-modules for $I \subset J$. The 2-morphisms are all natural transformations (i.e. homogeneous bimodule maps).

There is a 2-functor $\mathcal{G}$ from $\mathcal{SBSBim}_q$ to $\mathcal{SSBim}_q$ which is the identity on objects, sending downwards arrows to restriction functors (with grading shifts) and upwards arrows to induction functors. The cups and caps are sent to units and counits of adjunction. An upwards crossing is sent to the natural isomorphism between iterated induction functors. The grading shift on the restriction functor from R^I to R^J is $\ell(J) - \ell(I)$, where $\ell(p) = \ell(g) = \ell(o) = 3$, $\ell(r) = \ell(b) = \ell(y) = 1$, and $\ell(\emptyset) = 0$; this is the length of the longest element in the corresponding parabolic subgroup of the affine Weyl group.

Claim 3.16. *Suppose that* $\mathbb{k} = \mathbb{Q}(q)$*, or that* $q = 1$ *and* $\mathbb{k}$ *is a field of characteristic* $\neq 2, 3$*. Then* $\mathcal{G}$ *is well defined and fully faithful.*

The proof of this fact is a long algorithmic computation, and can be found in the appendix.

3.4. The equivalence.

Definition 3.17. Let $m\mathcal{SBSBim}_q$ denote the full sub-2-category of $\mathcal{SBSBim}_q$ defined as follows. The objects are $\{o, g, p\}$. The 1-morphisms are generated by IsJ whenever I, J are doubletons with $I \cap J = s$. The 2-morphisms are full. In other words, our diagrams may contain any color, but the right-hand and left-hand colors must be in $\{o, g, p\}$, and the bottom and top boundaries do not contain the color white.

Definition 3.18. Let $\mathcal{F}$ be the 2-functor $\mathrm{Fund}^\Omega_q \to m\mathcal{SBSBim}_q$ defined as follows. On objects, it acts as the identity. On 1-morphisms, it sends $go \mapsto gyo$, $og \mapsto oyg$ and so forth. On 2-morphisms, it acts as below.
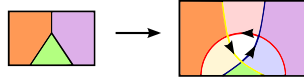
Claim 3.19. *The 2-functor $\mathcal{F}$ is well defined, and its image consists of degree 0 maps.*

Proof. The isotopy relations in Fund_q^{Ω} follow from the isotopy relations in $m\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim}_q$. We check the remaining relations for a specific coloration, but the computation for other colorations is almost identical.

Apply $\mathcal{F}$ to the LHS of the circle relation (3.3) to obtain a clockwise circle inside a anticlockwise circle. By (3.14b) and (3.14a), the result is multiplication by $\partial_p^r(\mu_o^r)$. By (3.10), this is [3].

Apply $\mathcal{F}$ to the LHS of the bigon relation (3.4).
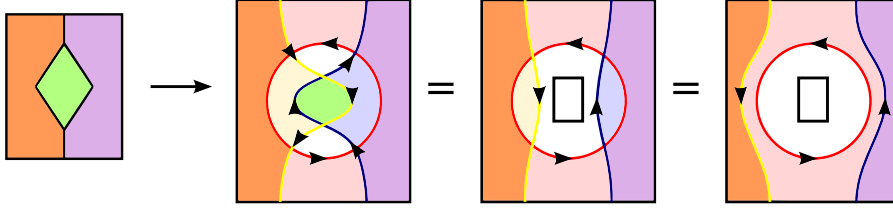
The first equality arises from (3.15c), and the second from (3.15a), while the polynomial in the box is $\mu_g^{b,y}$. Evaluating the red circle with (2.8c), one obtains $\partial_r(\mu_g^{b,y})$, which by (3.11) is $-[2]$ as desired.

Apply $\mathcal{F}$ to the LHS of the square relation (3.5).
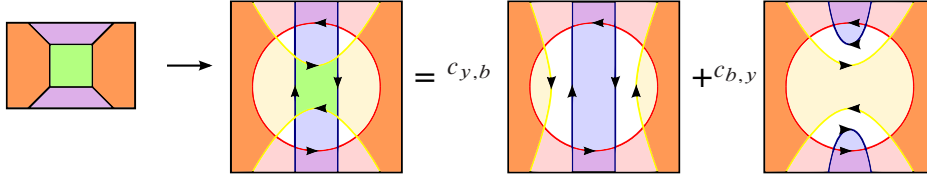
The equality arises from relation (3.16). By applying (3.15a) several times, this becomes
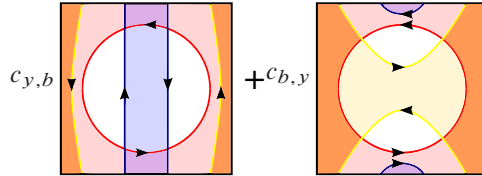
after which we apply (3.17) to each side to obtain
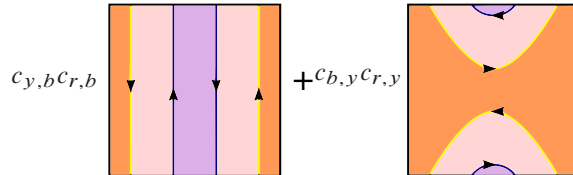
Now the result follows from (3.12).   □

Remark 3.20. Most people would agree that representation theory is the "easy" side of the geometric Satake equivalence, while perverse sheaves or Soergel bimodules are the "hard" side, at least in terms of the requisite background. Computationally, however, we like to argue that $\mathcal{SBSBim}_q$ is in some sense "easier" to work with than Fund_q, because the relations are more local. For example, the generating trivalent vertex is sent by $\mathcal{F}$ to a composition of several simpler maps, and the bigon and square relations each follow (after applying $\mathcal{F}$) from a sequence of relations on subdiagrams.

Theorem 3.21. *Suppose that* 3 *is invertible in* $\Bbbk$. *Then the* 2*-functor* $\mathcal{F}$ *is a degree-zero equivalence.*

Proof. This proof follows a similar diagrammatic argument to the $\mathfrak{sl}_2$ case, but is quite a bit more complicated. It can be found in the appendix. $\square$

4. Representations and $\mathfrak{sl}_n$-webs

The goal of this chapter is to give some general introduction to the 2-category of representations of a complex semisimple Lie algebra $\mathfrak{g}$. The focus will be on the algebraic, generators and relations approach. We feel as though the literature is lacking an introduction of this form, though another similarly-focused introduction can be found in [5].

When we eventually apply this technology to geometric Satake, we will be applying it to the case of the Langlands dual $\mathfrak{g}^\vee$. However, in the interest of making this chapter more readable, we stick with $\mathfrak{g}$ until §4.3.

**4.1. Some algebraic philosophy.** Let $\mathrm{Rep} = \mathrm{Rep}(\mathfrak{g})$ denote the monoidal category of finite dimensional representations of $\mathfrak{g}$, and let $\mathrm{Rep}_q = \mathrm{Rep}(U_q(\mathfrak{g}))$ denote its q-deformation.

The category Rep is semisimple, which implies that for any two irreducibles V_λ and V_μ one has $\mathrm{Hom}(V_\lambda, V_\mu) = \delta_{\lambda\mu}\mathbb{C}$. In other words, describing morphisms between arbitrary objects (given a decomposition into irreducibles) is a trivial task. A semisimple category over $\mathbb{C}$ is determined up to equivalence by its Grothendieck group, reinforcing the fact that the morphisms do not play an essential role.

However, Rep has a monoidal structure, and a semisimple monoidal category is not determined by its Grothendieck ring (see, for instance, [14, Problem 1.42.8 and following]). As such, we should study its morphisms in more detail.

Goal 4.1. *Describe the morphism spaces* $\mathrm{Hom}(V_{\lambda_1} \otimes \cdots \otimes V_{\lambda_d}, V_{\mu_1} \otimes \cdots \otimes V_{\mu_e})$ *for any sequences of irreducible representations, along with the structures of vertical (i.e. usual) and horizontal (i.e. monoidal) composition.*

This appears to be an incredibly difficult problem, and it remains open for almost all $\mathfrak{g}$. A more tractable approach is to study $\mathrm{Fund} = \mathrm{Fund}_\mathfrak{g}$, the full monoidal

subcategory whose objects are tensor products $V_{\omega_{i_1}} \otimes \cdots \otimes V_{\omega_{i_d}}$ of fundamental representations. Every irreducible is a summand of such a tensor product, so that the *idempotent closure* or *Karoubi envelope* **Kar**(Fund) is isomorphic to Rep. An introduction to Karoubi envelopes can be found in [2], or on Wikipedia. Here are three main goals in the algebraic study of Fund and Rep, with a discussion of the known results.

Goal 4.2. *Describe the monoidal category Fund by generators and relations.*

The language of description tends to be planar diagrammatics. Of course, the presentation will depend on a choice of generating morphisms (just as our category Fund came from a choice of generating objects), but choices can be made so that the relations take on a nice form.

- When $\mathfrak{g} = \mathfrak{sl}_2$, Fund is isomorphic to the Temperley–Lieb category [42], which was given a diagrammatic interpretation by Kauffman [23]. The diagrams involved are called crossingless matchings.
- When $\mathfrak{g}$ has rank 2, a presentation for Fund was found by Kuperberg [29]. His term for the diagrams, *webs*, has become standard, so that crossingless matchings can also be called $\mathfrak{sl}_2$-webs.
- When $\mathfrak{g} = \mathfrak{sl}_n$, a conjectural presentation for Fund due to Morrison [37] (and Kim [27] for $n = 4$) was recently proven by Cautis–Kamnitzer–Morrison [5], in terms of $\mathfrak{sl}_n$-*webs*.
- Outside of these cases, a presentation for Fund is unknown.

In this paper we rely only on the description of Fund($\mathfrak{sl}_n$) by $\mathfrak{sl}_n$-webs, as well as an analogous description of a nice subcategory of Soergel bimodules. However, to motivate the introduction of Fund, we should at least explain how one would complete this approach into a solution to the original problem, Goal 4.1.

Let $\lambda = \sum a_i \omega_i$ be a decomposition of a dominant integral weight λ as a sum of fundamental weights ω_i, so that $a_i \in \mathbb{N}$. Let $V = V_{\omega_{i_1}} \otimes \cdots \otimes V_{\omega_{i_d}}$ be some tensor product where each fundamental representation V_{ω_i} appears a_i times. Then V_λ is a summand of V with multiplicity one, so that there is some canonical idempotent $e_V \in \text{End}(V)$ which projects to V_λ. In the literature, such idempotents are often called *clasps*.

Goal 4.3. *Using the algebraic description of Fund from the previous goal, give a (possibly recursive) formula for e_V.*

- When $\mathfrak{g} = \mathfrak{sl}_2$, the clasps are known as Jones–Wenzl projectors [22, 43], and several recursive formulas [4, 43] and a closed formula [36] are known.
- When $\mathfrak{g}$ has rank 2, a recursive formula for clasps was found by Kim [27].
- The author [7] has recently conjectured a recursive formula for clasps when $\mathfrak{g} = \mathfrak{sl}_n$, proven to hold for $n \leq 4$.
- Outside of these cases, nothing is known.

Goal 4.3 is too naive, so we replace it by Goal 4.4 below (though it seems unlikely that one could solve the former without the latter). Consider the finite set $\mathbb{X}_\lambda$ consisting of all tensor products of fundamental representations whose fundamental weights add up to λ (they are all isomorphic). For any two $V, W \in \mathbb{X}_\lambda$, there is a 1-dimensional space of maps $V \to W$ which factor through the common summand V_λ.

Goal 4.4. *Find a family φ_λ, consisting of maps $\varphi_{V,W} \colon V \to W$ for each $V, W \in \mathbb{X}_\lambda$, such that $\varphi_{V,W}$ factors through V_λ, and such that $\varphi_{W,X} \circ \varphi_{V,W} = \varphi_{V,X}$ for all $V, W, X \in \mathbb{X}_\lambda$.*

Note that $\varphi_{V,V}$ is none other than the clasp e_V. Together, the family φ_λ gives a canonical identification of V_λ as a common summand of each object in the family $\mathbb{X}_\lambda$.

Given a presentation of a monoidal category and an intertwining family φ (such as φ_λ), it is a trivial operation to produce a presentation of the partial Karoubi envelope, where the image of φ is formally added as a new object. In the literature (e.g. [25]), such a planar calculus which describes morphisms in a partial Karoubi envelope is called a *thick calculus*, or perhaps a *thicker calculus* (with the thickest calculus being the entire Karoubi envelope). Adding φ_λ for each dominant weight λ, we obtain a solution to Goal 4.1.

This need not be the end of the story, as the thick calculus may itself have a number of interesting formulas and consequences, and this may simplify the presentation. For example, the $3j$ and $6j$ symbols and theta networks in $\mathfrak{sl}_2$ representation theory are computations within the thick calculus. For more about these topics, see [16] and the references it contains.

One additional feature of the known diagrammatic presentations is that they are easily adapted to the study of quantum group representations. A q-deformation of the $\mathfrak{sl}_n$-web algebras discussed above will yield a description of Fund_q, tensor products of fundamental representations of $U_q(\mathfrak{sl}_n)$. When q is specialized to a root of unity (with $U_q(\mathfrak{sl}_n)$ denoting the Lusztig form), representations of $U_q(\mathfrak{g})$ are no longer semisimple; $\mathbf{Kar}(\mathrm{Fund}_q)$ will no longer be equivalent to Rep_q, being equivalent to the category of tilting modules instead. By understanding which denominators in the formulas for clasps will vanish at a root of unity, one can address the behavior of these tilting modules in an algebraic manner. Of course, this is already known using abstract representation theory, but knowing the clasps explicitly will still yield measurable gains. In the broader context of Soergel bimodules, computing the denominators in analogous idempotents is of great interest in modular representation theory.

4.2. Monoidal-graded categories.

Definition 4.5. (c.f. [15]) Let H be a finite abelian group. A monoidal category $\mathcal{C}$ is *H-graded-monoidal* if there are subcategories $\mathcal{C}_\xi$ for $\xi \in H$, for which $\mathcal{C} = \oplus_{\xi \in H} \mathcal{C}_\xi$, and for which $\mathcal{C}_\xi \otimes \mathcal{C}_\nu \subset \mathcal{C}_{\xi\nu}$.

Definition 4.6. The *fundamental group* $\pi(\mathfrak{g})$ is defined as $\Lambda_{\mathrm{wt}}/\Lambda_{\mathrm{rt}}$, the quotient of the weight lattice by the root lattice. We also denote it Ω.

The category $\mathrm{Rep} = \mathrm{Rep}(\mathfrak{g})$ is Ω-graded-monoidal (as is Rep_q). For $\xi \in \Omega$, let Rep_ξ denote those representations whose irreducible constituents all have highest weights in the coset ξ. Clearly the tensor product obeys the group law. From this monoidal grading one can reconstruct the category of G representations for any corresponding algebraic group.

Example 4.7. Representations of $\mathfrak{sl}_2$ or SL_2 can be split into even and odd representations. These can be distinguished by the action of the subgroup $\{\pm 1\} \subset SL_2$ on these representations. The subcategory of even representations is equivalent to the category of $\mathbb{P}SL_2$-representations.

Remark 4.8. The adverb H-**graded** modifies the adjective **monoidal**. One should not confuse this with other uses of the term **graded**. For instance, a category is *graded* or $\mathbb{Z}$-*graded* if (by conflating maps of all "degrees") its Hom spaces can be enriched in graded vector spaces. Therefore, every (homogeneous) **morphism** has an associated degree in $\mathbb{Z}$. On the other hand, every (indecomposable) **object** in an H-graded-monoidal category has an associated character in H.

Construction 4.9. It is a trivial operation to take an H-graded-monoidal category and replace it with a 2-category $\mathfrak{C}$ which encodes (almost) the same data. Let S be a right H-torsor. The objects of $\mathfrak{C}$ will be identified with $s \in S$, and for $\xi \in H$ the morphism category from s to $s\xi$ is $\mathcal{C}_\xi$. A 2-category where the objects form an H-torsor, equipped with a corresponding H action which identifies $\mathrm{Hom}(s, t) = \mathrm{Hom}(s\xi, t\xi)$, we will call an H-2-*category*, for lack of better terminology.

Remark 4.10. The data of an H-2-category $\mathfrak{C}$ together with a choice of $s \in S$ (corresponding to $1 \in H$) and the data of an H-graded-monoidal category $\mathcal{C}$ are equivalent. Philosophically speaking, the difference between them is analogous to the difference between an algebra and its regular representation. Because $\mathcal{C}$ is monoidal, it acts on itself by tensor product. Under the splitting $\mathcal{C} = \oplus_{\xi \in H} \mathcal{C}_\xi$, the action of $M \otimes (\cdot)$ for any object M can be expressed as an $H \times H$ matrix. Each entry of this matrix is a corresponding 1-morphism in $\mathfrak{C}$.

Thus from Rep (resp. Rep_q) we can construct a Ω-2-category, which we denote Rep^Ω (resp. Rep_q^Ω). It is simple to pass from a diagrammatic description of $\mathrm{Rep}(\mathfrak{g})$ (say, by $\mathfrak{sl}_n$-webs) to a diagrammatic description of Rep^Ω. One simply colors each planar region by an element of Ω, and imposes the rule that tensoring with V_{ω_i} will shift the color by ω_i. This is illustrated by the $\mathfrak{sl}_2$ and $\mathfrak{sl}_3$ cases in previous chapters.

The reason we switch from graded monoidal categories to 2-categories is that it makes it easier to compare Rep_q^Ω with $\mathcal{SS}\mathrm{Bim}_q$, which is more naturally a 2-category. When $q = 1$, the subcategory $m\mathcal{S}\mathbb{BS}\mathrm{Bim}_q$ which is the target of the algebraic Satake equivalence is equivalent to a graded Ω-2-category, and thus arises from some

$\mathbb{Z}$-graded Ω-graded-monoidal category, though this is perhaps an unnatural way of viewing it.

Remark 4.11. The monoidal category Rep has a symmetric structure. There is no notion of a symmetric structure on a 2-category, but the corresponding structure on an Ω-2-category is not hard to formulate.

4.3. The fundamental group. We return to the setting of geometric Satake, which cares about representations of $\mathfrak{g}^\vee$. We provide several important facts about the fundamental group $\Omega = \pi(\mathfrak{g}^\vee)$. It is a finite abelian group, so that its Pontrjagin dual or character group Ω^* is also a finite abelian group, non-canonically isomorphic to Ω.

Let Γ denote the Dynkin diagram of $\mathfrak{g}$, and $\Gamma^\vee$ the Langlands dual Dynkin diagram. Write G_{adj} or G_{sc} for the adjoint or simply-connected algebraic groups for Γ. Let $\mathcal{K} = \mathbb{C}((t))$ and $\mathcal{O} = \mathbb{C}[[t]]$.

Claim 4.12. *The following finite abelian groups are all canonically isomorphic, and will all be denoted* Ω.

- $\pi(\mathfrak{g}^\vee)$.

- $\pi(\mathfrak{g})^*$.

- $Z(G_{sc})$, *the center of* G_{sc}.

- $\pi_1(G_{adj})$.

- $\pi_0(G_{adj}(\mathcal{K}))$, *the component group of* $G_{adj}(\mathcal{K})$.

Let $\widetilde{\Gamma}$ denote the affine Dynkin diagram of Γ, and let 0 denote the affine vertex (i.e. choose an embedding $\Gamma \hookrightarrow \widetilde{\Gamma}$). Call a vertex $v \in \widetilde{\Gamma}$ *removable* if $\widetilde{\Gamma} \setminus v$ is isomorphic to Γ, and let $\widetilde{\Theta}$ denote the set of removable vertices. They form a single orbit under the automorphism group of $\widetilde{\Gamma}$. Let $\Theta = \widetilde{\Theta} \setminus 0 \subset \Gamma$. Identifying the vertices of $\Gamma^\vee$ and Γ, we can view Θ as a subset of $\Gamma^\vee$. It is known that the fundamental representations of $\mathfrak{g}^\vee$ associated to vertices in Θ are precisely the *miniscule* fundamental representations, meaning that their weights form a single Weyl group orbit.

Claim 4.13. *The set* $\widetilde{\Theta}$ *is in canonical bijection with* Ω. *That is, it is an* Ω-torsor *with a distinguished element* 0.

Essentially, the statement is that the miniscule fundamental weights, together with the zero weight, enumerate a list of representatives in $\Lambda_{\text{wt}}^\vee / \Lambda_{\text{rt}}^\vee = \Omega$.

Remark 4.14. Here, Langlands duality is essential. The miniscule fundamental weight in type B is precisely the removable vertex from affine type C, and vice versa.

These facts are quite standard, see for instance [46, Chapter 2].

5. Singular Soergel bimodules

The goal of this chapter is to provide background material on algebraic Soergel theory. For the connections with geometry, see §6. This background is not essential to the main theorems.

The exotic realizations of affine $\mathfrak{sl}_n$ for $n > 2$ are examples of "odd-unbalanced" realizations, as defined in [8], while the ordinary realization of affine $\mathfrak{sl}_n$ is odd-balanced. Being odd-unbalanced adds an additional layer of technicality which readers uninterested in the q-deformation of geometric Satake can ignore. Because of this, it may be better for readers to get their first introduction to singular Soergel bimodules elsewhere, such as in [45]. This paper is the first to use odd-unbalanced realizations in an essential way, which is why we devote a great deal of effort to explaining these technicalities.

5.1. Overview. Here is a big picture to keep in mind while reading the nitty-gritty.

Given a (reflection faithful) representation $\mathfrak{h}$ of a Coxeter group W, one can consider the polynomial ring $R = \mathcal{O}(\mathfrak{h})$. This is a graded ring whose linear terms $\mathfrak{h}^*$ will be placed in degree 2, equipped with a natural action of W. The letters I, J will designate subsets of the simple reflections S of W. For any parabolic subgroup $W_I \subset W$ there is a subring of invariants $R^I \subset R$. Under mild assumptions, the ring extension $R^s \subset R$ is graded Frobenius, meaning that induction and restriction functors are biadjoint (up to a shift).

Definition 5.1. Let $\mathbb{BSBim} = \mathbb{BSBim}_q$, the monoidal category of *Bott–Samelson bimodules*, be the full subcategory of (R, R)-bimodules whose 1-morphisms are generated by

$$B_s = R \otimes_{R^s} R(1)$$

for each simple reflection $s \in S$, the composition of restriction with induction and a shift. The grading shift places $1 \otimes 1$ in degree -1.

Definition 5.2. The monoidal category $\mathbb{SBim}$ of *Soergel bimodules* is the Karoubi envelope of $\mathbb{BSBim}$.

Thus to define $\mathbb{SBim}$ one only needs the data of a reflection representation $\mathfrak{h}$ satisfying some conditions. Under some additional conditions, Soergel [41] has proven that $\mathbb{SBim}$ has a number of desirable properties: it categorifies the Hecke algebra, and the Hecke algebra is equipped with a pairing which encodes the graded dimension of morphism spaces in $\mathbb{SBim}$.

However, our goal is not just to define Soergel bimodules but to find an algebraic presentation for them. One must first choose a set of generating morphisms. As it turns out, this arises from a choice of Frobenius structure on the extensions $R^s \subset R$ for $s \in S$, which boils down to an explicit choice of simple roots in $\mathfrak{h}^*$. A reflection representation paired with a choice of simple roots is (essentially) what we call a *simple realization* below, the word *simple* indicating that is corresponds to a choice

of simple roots. A simple realization can roughly be encoded in a Cartan matrix. From a simple realization one can define a diagrammatic category by generators and relations which is equivalent to $\mathbb{BS}$Bim. Moreover, for a slightly more general notion of a simple realization (allowing for non-faithful representations and more) one can still define this diagrammatic category and prove the categorification results of Soergel, even when it is no longer equivalent to $\mathbb{BS}$Bim. This is shown in [12].

Now return to the original setting of a (reflection faithful) representation $\mathfrak{h}$. Under certain assumptions, whenever $W_I \subset W_J$ are two finite parabolic subgroups, the ring extension $R^J \subset R^I$ will be graded Frobenius (this is an upgraded version of Chevalley's Theorem [21, Chapter 3]). Let $\ell(I)$ denote the length of the longest element of W_I.

Definition 5.3. Let $\mathcal{S}\mathbb{BS}$Bim $= \mathcal{S}\mathbb{BS}$Bim$_{\text{alg}}$, the 2-category of *singular Bott–Samelson bimodules*, be defined as follows. The objects are the rings R^I for finite parabolic subgroups W_I. The morphism category $\text{Hom}(R^I, R^J)$ will be a full subcategory of (R^J, R^I)-bimodules. The 1-morphisms in $\mathcal{S}\mathbb{BS}$Bim are generated by the induction bimodule ${}_{R^I}R^I_{R^J}$ and the restriction bimodule ${}_{R^J}R^I_{R^I}(\ell(J)-\ell(I))$ for the extensions $R^J \subset R^I$ whenever $I \subset J$.

Definition 5.4. The 2-category $\mathcal{S}\mathbb{S}$Bim of *singular Soergel bimodules* is the Karoubi envelope of $\mathcal{S}\mathbb{BS}$Bim.

Thus to define $\mathcal{S}\mathbb{S}$Bim one only needs the data of a representation $\mathfrak{h}$, satisfying some conditions. Under additional conditions, Williamson [45] has proven that $\mathcal{S}\mathbb{S}$Bim has a number of desirable properties: it categorifies the Hecke algebroid, and the Hecke algebroid is equipped with a pairing which determines the graded rank of 2-morphism spaces.

To find an algebraic presentation for $\mathcal{S}\mathbb{BS}$Bim one must choose a set of generating 2-morphisms. Again, this arises from a choice of Frobenius structures on every extension $R^J \subset R^I$ together with some consistency condition on these structures. We call this data a *Frobenius realization* (Definition 3.9). One way to specify a Frobenius realization is to make an explicit choice of positive roots in $\mathfrak{h}^*$, which we call a *root realization* (Definition 3.8). From a Frobenius realization one can define a diagrammatic category by generators and relations which is equivalent to $\mathcal{S}\mathbb{BS}$Bim. As above, one can generalize the notion of a Frobenius or root realization, and one should be able to define 2-categories which possess the categorification properties proven by Williamson, even without being equivalent to $\mathcal{S}\mathbb{BS}$Bim. This will eventually be proven in type A in [10]. It was proven for dihedral groups in [8], and will be proven for affine $\mathfrak{sl}_3$ in the appendix.

Let α_s denote the simple root attached to $s \in S$. In familiar notions of root systems (as in, say, Humphreys' book [21]), a choice of simple roots determines a choice of positive roots. One definition is that the positive roots are the elements of the set $\{s_1 s_2 \ldots s_{d-1}(\alpha_{s_d})\}$ for each reduced expression $\underline{w} = s_1 \ldots s_d$. This set has redundancies: for example, when $m_{st} = 3$, one has $s(\alpha_t) = t(\alpha_s)$. However,

for more general realizations, the previously redundant descriptions of the same positive root may differ by a scalar! For instance, using the notation from §3.2, one has $m_{by} = 3$ but $y(\alpha_b) = qb(\alpha_y)$. This kind of ambiguity can occur whenever the Cartan matrix of the realization is odd-unbalanced, in which case the choice of positive roots is a piece of additional data. In a precisely analogous computation, the Demazure operators attached to simple roots will be used to define the trace maps for all Frobenius extensions $R^J \subset R^I$, but they only satisfy the braid relations when the realization is odd-balanced.

There is not a great deal of complication involved — just keeping track of some scalars which measure the failure of roots to be well defined! The notion of a root realization is developed as a natural way to keep track of these scalars. Remember also that the 2-categories themselves only depend on the representation $\mathfrak{h}$, while the presentation depended on the root realization.

The sections to follow will explore these concepts in more detail, with constant reference to the exotic realization of affine $\mathfrak{sl}_n$. The one remaining subtlety is that, unlike the familiar affine Cartan matrix, the exotic affine Cartan matrix is non-degenerate. This may lead to possible confusion when comparing the exotic and familiar realizations, which we do our best to forestall. This motivates the discussion of dual realizations below.

5.2. Simple realizations. Fix a Coxeter group W with a set of simple reflections S, and let m_{st} be the order of st, for $s, t \in S$. Let $\Bbbk$ be a commutative ring. We simplify the discussion somewhat by assuming that $\Bbbk$ has no 2-torsion. Typical choices for $\Bbbk$ are $\mathbb{C}$, $\mathbb{Q}(q)$, and $\mathbb{Z}[\zeta_n]$ for a primitive n-th root of unity ζ_n.

Definition 5.5. A *(simple) realization* of (W, S) over $\Bbbk$ is a free, finite rank $\Bbbk$-module $\mathfrak{h}$, together with subsets $\Delta^\vee = \{\alpha_s^\vee \mid s \in S\} \subset \mathfrak{h}$ and $\Delta = \{\alpha_s \mid s \in S\} \subset \mathfrak{h}^* = \mathrm{Hom}_\Bbbk(\mathfrak{h}, \Bbbk)$, called *simple coroots* and *simple roots* respectively, satisfying:

(1) $\langle \alpha_s^\vee, \alpha_s \rangle = 2$ for all $s \in S$;

(2) the assignment

$$s(v) \overset{\text{def}}{=} v - \langle v, \alpha_s \rangle \alpha_s^\vee$$

for all $v \in \mathfrak{h}$ yields a representation of W.

We will often refer to $\mathfrak{h}$ as a realization, however the choice of simple roots and coroots is always implicit.

Note that $\mathfrak{h}^*$ is equipped with a contragredient action of W, for which

$$s(f) = f - \langle \alpha_s^\vee, f \rangle \alpha_s.$$

Definition 5.6. Given a realization $(\mathfrak{h}, \mathfrak{h}^*, \Delta^\vee, \Delta)$, the *dual realization* is $(\mathfrak{h}^*, \mathfrak{h}, \Delta, \Delta^\vee)$. This makes sense, since $\mathfrak{h}^{**} \cong \mathfrak{h}$ canonically.

Definition 5.7. The *Cartan matrix* associated to a realization $\mathfrak{h}$ is the $S \times S$ matrix A valued in $\Bbbk$, with entries $a_{s,t} = \langle a_s^\vee, \alpha_t \rangle$.

Note that we did not assume that either Δ or $\Delta^\vee$ is a basis, or is even linearly independent. When $\Delta^\vee$ is a basis for $\mathfrak{h}$, the realization is determined by its Cartan matrix. However, "enlarging" the representation beyond the span of Δ may affect the image of α_s or $\alpha_s^\vee$ in $\Bbbk$.

Example 5.8. Let W be a Weyl group equipped with a root system. Then the familiar reflection representation (with symmetric Cartan matrix, as in [21]) is a self-dual realization of W over $\mathbb{R}$.

Example 5.9. Let W be an affine Weyl group, with $n + 1$ simple reflections. It has a familiar realization over $\mathbb{R}$ of rank $n + 1$ where Δ forms a basis, while $\Delta^\vee$ is linearly dependent. It has a dual realization of rank $n + 1$ where $\Delta^\vee$ forms a basis, and Δ is linearly dependent. It has a self-dual realization of rank $n + 2$ where both Δ and $\Delta^\vee$ are linearly independent, but neither is a basis.

Claim 5.10. *Suppose that the Cartan matrix has invertible (not just non-zero) determinant in $\Bbbk$. Then Δ and $\Delta^\vee$ are each linearly independent. As W representations one has $\mathfrak{h} \cong \Bbbk \cdot \Delta^\vee \oplus \mathbb{T}$ for some trivial representation $\mathbb{T}$ in the kernel of Δ. Similarly, one has $\mathfrak{h}^* \cong \Bbbk \cdot \Delta \oplus \mathbb{T}^*$.*

Proof. The fact that Δ and $\Delta^\vee$ are linearly independent is obvious. Let v be a vector in $\mathfrak{h}$ which is not in $\Bbbk \cdot \Delta^\vee$. Using linear algebra, some vector $x \in \Bbbk \cdot \Delta^\vee$ satisfies $\alpha_s(x) = \alpha_s(v)$ for all $s \in S$. Then $v - x$ is in the kernel of Δ, so that W acts trivially on $v - x$. In this fashion, one can replace any basis for $\mathfrak{h}$ extending $\Delta^\vee$ by a basis where the remaining elements are in the kernel of Δ. $\qquad\square$

It is obvious how one equips the base change $\mathfrak{h} \otimes_\Bbbk \Bbbk'$ with the structure of a realization over $\Bbbk'$. If $\Delta^\vee$ spans a free module of rank $|S|$ in $\mathfrak{h}$, then it does the same in $\mathfrak{h} \otimes_\Bbbk \Bbbk'$. However, the same can not be said for Δ. For any Cartan matrix over a domain $\Bbbk$ with non-zero non-invertible determinant D, base change to $\Bbbk/D$ will lower the rank of the span of Δ. Therefore, base change does not commute with duality.

5.3. The exotic realization. In this section we let W be the affine Weyl group in type $\widetilde{A_n}$, and let $\Bbbk = \mathbb{Z}[q, q^{-1}]$. The finite Weyl group W_{fin} will be embedded inside W, generated by the first n simple reflections.

Definition 5.11. We define the *exotic affine Cartan matrix of $\mathfrak{sl}_{n+1}$*. For $n \geq 2$ it has a form which is exemplified by the case $n = 4$:

$$
A = \begin{pmatrix}
2 & -1 & 0 & 0 & -q^{-1} \\
-1 & 2 & -1 & 0 & 0 \\
0 & -1 & 2 & -1 & 0 \\
0 & 0 & -1 & 2 & -q \\
-q & 0 & 0 & -q^{-1} & 2
\end{pmatrix}.
$$

Exotic affine $\mathfrak{sl}_2$ on the other hand is given by a symmetric Cartan matrix:

$$A = \begin{pmatrix} 2 & -(q+q^{-1}) \\ -(q+q^{-1}) & 2 \end{pmatrix}.$$

We call any realization over $\mathbb{Z}[q^{\pm 1}]$ with this matrix an *exotic realization* of W.

Obviously, setting $q = 1$ yields the usual affine Cartan matrix for $\widetilde{A_n}$.

Exercise 5.12. Whenever $a_{st}a_{ts} = 1$, one has $sts = tst$. Therefore, this matrix gives a realization of W.

Thus whenever $n \geq 2$, specialization of q will not change the order of $s_i s_j$ for any simple reflections, so that the realization is dihedrally faithful (as in [8]).

Exercise 5.13. The exotic Cartan matrix has determinant $2 - q^2 - q^{-2}$, which is non-zero and non-invertible. Therefore, in any base change of the exotic realization, $\Delta^\vee$ will be linearly independent.

Exercise 5.14. For the exotic realization of rank $n + 1$ where $\Delta^\vee$ is a basis, base change to $q^2 = 1$ will make Δ linearly dependent. This is the realization for which the action of W on $\mathfrak{h}^*$, after specialization to $q = 1$, is the reflection representation of W defined in [21]. In particular, an exotic realization is faithful. There are also exotic realizations for which base change to $q = 1$ will yield the other realizations mentioned in Example 5.9.

Number the simple roots so that $\Delta = \Delta_{\text{fin}} \cup \{\alpha_0\}$, where $\Delta_{\text{fin}} = \{\alpha_1, \ldots, \alpha_n\}$ has the usual finite Cartan matrix and α_0 is the affine simple root. Let s_i denote the reflection corresponding to α_i. Let $\beta = \sum_{i=1}^n \alpha_i$ and $\beta^\vee = \sum_{i=1}^n \alpha_i^\vee$, the highest finite root and coroot. Let t denote reflection across β, which can be written as $t = s_1 s_2 \cdots s_{n-1} s_n s_{n-1} \cdots s_2 s_1$. Let $[n]$ denote the n-th quantum number in $\mathbb{Z}[q^{\pm 1}]$.

Exercise 5.15. The subgroup W' generated by s_0 and t inside W is the infinite dihedral group. Let $\Delta' = \{\beta, \alpha_0\}$ and $(\Delta^\vee)' = \{\beta^\vee, \alpha_0^\vee\}$. Then $(\mathfrak{h}, \mathfrak{h}^*, \Delta', (\Delta^\vee)')$ gives a realization of W', having as its Cartan matrix the exotic affine Cartan matrix of $\mathfrak{sl}_2$.

In [8] the exotic realization of $\mathfrak{sl}_2$ was studied in detail. It was shown that the realization is faithful unless q^2 is a root of unity. If q^2 is a primitive m-th root of unity then the kernel of the action is generated by $(uv)^m$ (for simple reflections u, v).

Proposition 5.16. *After base change, an exotic realization of $\mathfrak{sl}_{n+1}$ is still faithful unless q^2 is a primitive m-th root of unity, in which case the kernel K is generated by $(s_0 t)^m$.*

Presumably there is an elegant proof of this fact, but we provide a brute force proof.

Proof. We assume that q^2 is a primitive m-th root of unity in the rest of the proof; it is an exercise to modify this into a proof of faithfulness when q^2 is not a root of

unity. Our first step is to examine $H = K \cap \Lambda_{\mathrm{rt}}$, those translations within the kernel. Clearly H is closed under the action of W_{fin}. We wish to show that $H = m\Lambda_{\mathrm{rt}}$.

The element $(s_0 t)^m$ is translation by the m-th multiple of the highest root. Since this translation is in H, the entirety of $m\Lambda_{\mathrm{rt}}$ is in H. Moreover, $(s_0 t)^k \in H$ if and only if k is a multiple of m (by the dihedral result above), meaning that a k-th multiple of a root is in H if and only if k is a multiple of m.

If $\lambda \in H$ then $s_i(\lambda) - \lambda = \langle \lambda, \alpha_i \rangle \alpha_i \in H$, which is a multiple of a root. Therefore $\langle \lambda, \alpha_i \rangle$ is a multiple of m for all i, meaning that $\lambda \in m\Lambda_{\mathrm{wt}}$. Thus $m\Lambda_{\mathrm{rt}} \subset H \subset m\Lambda_{\mathrm{wt}}$. Since $m\Lambda_{\mathrm{wt}}/m\Lambda_{\mathrm{rt}} \cong \mathbb{Z}/(n+1)\mathbb{Z}$, there are not many choices for H: one has $H = H_k$ for some k dividing $n+1$, where H_k is generated by $m\Lambda_{\mathrm{rt}}$ and $km\Lambda_{\mathrm{wt}}$. Moreover, $H \subset \Lambda_{\mathrm{rt}}$, which places further restrictions. We seek to show that $k = n+1$, in which case $km\Lambda_{\mathrm{wt}}$ is contained inside $m\Lambda_{\mathrm{rt}}$.

Now consider the element $s_1 s_2 \ldots s_0 = x$. A simple computation shows that x preserves the subspace spanned by $\{\alpha_2, \ldots, \alpha_n\} \cup \{q^{-1}\alpha_1 + \alpha_0 = s_1(\alpha_n)\}$, and that it acts by a rescaled permutation matrix. From this it is easy to compute that the order of x is precisely mn. Note that x^n is translation by an $(n+1)$-st multiple of a fundamental weight which, because $\Omega = \mathbb{Z}/(n+1)\mathbb{Z}$, is in Λ_{rt}. Let ω be the fundamental weight with $(n+1)\omega = x^n$.

Let d be the greatest common divisor of $n+1$ and m, so that $n+1 = da$ and $m = db$. The fact that $km\Lambda_{\mathrm{wt}} \subset \Lambda_{\mathrm{rt}}$ implies that n divides km, meaning that $k = al$ for some l dividing d. However, K contains $km(\omega) = bl(n+1)\omega = (x^n)^{bl}$. By previous computation, this is impossible unless $l = d$ and $k = n+1$. Thus we have finally proven that $H = m\Lambda_{\mathrm{rt}}$.

Our next step is to examine J, the image of K under the quotient map $W \to W_{\mathrm{fin}}$ which kills all translations. Clearly J is a normal subgroup of S_{n+1}, which means that it is trivial or it contains the alternating subgroup A_{n+1} (or $n = 1$). If we show that J is trivial then we have proven the proposition. Thus we aim for a contradiction when $s_1 s_2 \in J$ (or $s_1 \in J$ for $n = 1$). We sketch the rest of the proof.

Let λ denote translation by $\lambda \in \Lambda_{\mathrm{rt}}$, and suppose that $w\lambda \in K$ for $w \in W_{\mathrm{fin}}$. By applying H we may freely shift λ by an element of $m\Lambda_{\mathrm{rt}}$. Clearly $\lambda \notin m\Lambda_{\mathrm{rt}}$, because W_{fin} acts faithfully. By a fairly straightforward argument using convolution and multiplication, one can show that $w(\lambda) - \lambda \in K$, which implies by the result above that $w(\lambda) - \lambda \in m\Lambda_{\mathrm{rt}}$. If this is the case for $w = s_1$ or $w = s_1 s_2$, it is a similar argument to show that $s_i(\lambda) - \lambda \in m\Lambda_{\mathrm{rt}}$ for all $i \leq n$, or in other words, that $\lambda \in m\Lambda_{\mathrm{wt}}$. Letting d be the greatest common divisor of $n+1$ and m, and a, b as above, this implies that $\lambda \in b\Lambda_{\mathrm{rt}} \cap m\Lambda_{\mathrm{wt}}$. After some manipulation we may assume that $\lambda = (x^n)^{bl}$ for some l dividing d.

Now the result follows from computing $s_1 x^{nk}$ and $s_1 s_2 x^{nk}$ on the subbasis $\{\alpha_2, \ldots, \alpha_n\} \cup \{q^{-1}\alpha_1 + \alpha_0 = s_1(\alpha_n)\}$, and observing that these operators are never trivial. $\square$

We conclude this section with a discussion of where the exotic representations of W can be found in the literature, and some history. This was explained to me by Lusztig.

There is another "deformation of the Cartan matrix" which may be more familiar to readers: one with the usual (integral) non-diagonal entries, but replacing the diagonal entries 2 with $v + v^{-1}$. (This is a different formal parameter v, not the same as q!) Warning: this is not a Cartan matrix, in the sense we have defined above. This deformation, originally found by Killing, was used by Kilmoyer in his thesis [26] (see also [6]) in finite type to define a representation of the Hecke algebra (with parameter v) which deforms the usual reflection representation. Analogously, the v-deformed q-exotic Cartan matrix provides a one-parameter family of affine Hecke algebra representations.

Kazhdan and Lusztig provide a definition of W-graphs [24], giving v-deformations of other W representations. In [33], Lusztig gives a number of W-graphs for affine Weyl groups, using a special left cell; in affine types A and C, these W-graphs can be observed to have periodicity properties. In [31], it is explicitly described how to take a periodic W-graph and obtain a one-parameter family of Hecke representations. In type A, the corresponding family of affine Hecke representations, after passing to $v = 1$, is precisely the exotic family of representations of W.

However, our parametrization of this family by the variable q appears to be new. Accordingly, study of what happens when q is specialized to a root of unity also appears to be new.

Remark 5.17. There is also a one-parameter family of reflection representations in affine type C. One might pray that a q-parametrization might exist which gives a quantum Soergel Satake equivalence.

5.4. Frobenius extensions.

Definition 5.18. A *graded (commutative) Frobenius extension of degree d* is an extension of graded commutative rings $A \subset B$ equipped with a *trace map* $\partial \colon B \to A$ which is A-bilinear and has degree $-d$. There must exist bases $\{b_i\}$ and $\{b_j^*\}$ of B as a finite rank free A-module, satisfying $\partial(b_i b_j^*) = \delta_{i,j}$.

The choice of dual bases is not part of the structure of a Frobenius extension, but the choice of trace map is. For any ring extension there are two canonical bimodule maps: inclusion $i \colon A \to B(d)$ of A-bimodules, having degree $-d$, and multiplication $m \colon B \otimes_A B(d) \to B$ of B-bimodules, having degree $+d$. For a Frobenius extension there are two additional bimodule maps: trace $\partial \colon B(d) \to A$ of A-bimodules, having degree $-d$, and comultiplication $\Delta \colon B \to B \otimes_A B(d)$ of B-bimodules, having degree $+d$. The comultiplication map satisfies $\Delta(1) = \Delta_A^B = \sum b_i \otimes b_i^*$, this element being independent of the choice of dual bases.

Note that for any invertible scalar λ, $\lambda \partial$ will also define a trace map, producing a different Frobenius extension structure. One way to "measure" the Frobenius

structure is by the invariant $\mu_A^B = m(\Delta(1)) \in B$, the *product-coproduct element*, which has degree $2d$. Rescaling ∂ by λ will rescale one of the bases in the dual pair by λ^{-1}, and thus will rescale μ by λ^{-1}. The pair (∂, μ) will always satisfy $\partial(\mu) = k$, where $k \in \mathbb{N}$ is the rank of the extension. Thus if ∂ and μ are known up to a scalar, specifying either one will suffice to pin down the Frobenius extension structure.

Let us quickly note that for any Frobenius extension the trace map ∂ must be surjective.

Now we define the notion of a Frobenius hypercube, as in [13].

Definition 5.19. Let S be a finite set, and let ℓ be a function which assigns an integer to each subset of S. A *Frobenius hypercube* is data of a (commutative graded) ring R^I for each $I \subset S$, and whenever $I \subset J \subset S$, a Frobenius extension $R^J \subset R^I$ of degree $\ell(J) - \ell(I)$. The trace maps are required to satisfy $\partial_I^J \partial_J^K = \partial_I^K$ whenever $K \subset J \subset I$.

A *partial Frobenius hypercube* is the same definition, except restricted to subsets $I \subset S$ lying within a certain ideal (under the inclusion partial order). For example, one might only consider proper subsets of S.

Our goal will be to define a partial Frobenius hypercube attached to a realization, though this will require slightly more data than the realization itself provides.

5.5. Polynomials and invariants.

Assumption 5.20. The maps $\alpha_s : \mathfrak{h} \to \mathbb{k}$ and $\alpha_s^\vee : \mathfrak{h}^* \to \mathbb{k}$ are surjective for each $s \in S$. This assumption is called *Demazure surjectivity*.

If the row and column corresponding to s in the Cartan matrix each contain an invertible element of $\mathbb{k}$, then Demazure surjectivity follows immediately. Thus Demazure surjectivity holds for any exotic realization of $\mathfrak{sl}_{n+1}$ for $n \geq 2$. Demazure surjectivity for exotic affine $\mathfrak{sl}_2$ can be guaranteed by assuming that 2 and [2] generate the unit ideal in $\mathbb{k}$ (and this is required when Δ and $\Delta^\vee$ are bases).

We let $R = \text{Sym}(\mathfrak{h}^*)$. When the simple roots form a basis for $\mathfrak{h}^*$, this is just a polynomial ring on the variables α_s, $s \in S$. We give R an even $\mathbb{Z}$-grading, so that $\deg(\mathfrak{h}^*) = 2$. There is an action of W on R by ring homomorphisms, extending the action on the linear terms.

Let R^s denote the invariants of R under $s \in S$. There exists an R^s-linear *Demazure operator* $\partial_s : R \to R^s$ of degree -2, defined on linear terms by $\partial_s = \alpha_s^\vee : \mathfrak{h}^* \to \mathbb{k}$, and satisfying the twisted Leibniz rule $\partial_s(fg) = \partial_s(f)g + s(f)\partial_s(g)$. It can be defined by (3.7).

The Demazure operator ∂_s equips the ring extension $R^s \subset R$ with the structure of a graded Frobenius extension (this relies on Demazure surjectivity). The invariant μ_s of this extension structure is equal to α_s. Thus while the rings R and R^s only depend on the action of W on $\mathfrak{h}^*$, the Frobenius structure depends on the choice of simple roots and coroots, and one should think of a simple realization as being the choice of Frobenius structures on $R^s \subset R$ for all $s \in S$.

We say that I is *finitary* when the parabolic subgroup $W_I \subset W$ is finite. Let $R^I \subset R$ denote the invariants under W_I. When I is finitary, the Chevalley Theorem states that R is free over R^I of finite rank (whereas when W_I is infinite, R^I has smaller transcendence degree, and the extension has infinite rank). Under some additional assumptions, $R^I \subset R$ can be given the structure of a Frobenius extension of degree $\ell(I)$, where $\ell(I)$ is the length of the longest element $w_I \in W_I$.

Claim 5.21. *The operators ∂_s satisfy the braid relations up to an invertible scalar. For any reduced expression $\underline{w}_I = s_1 s_2 \ldots s_{\ell(I)}$ for w_I, one has a composition of Demazure operators $\partial_{\underline{w}_I} = \partial_{s_1} \cdots \partial_{s_d}$ of degree $-\ell(I)$, and this depends on the reduced expression only up to invertible scalar. Under some additional assumptions on $\Bbbk$, any such map $\partial_{\underline{w}_I}$ will equip $R^I \subset R$ with the structure of a Frobenius extension. Similarly, when $J \subset I$ are both finitary and some additional assumptions hold, any reduced expression for $w_I w_J^{-1}$ will define an iterated Demazure operator giving a Frobenius trace $R^J \to R^I$.*

Remark 5.22. The statement that a map is a Frobenius trace is none other than the existence of dual bases, which requires the map to be surjective, and requires a certain determinant to be invertible. What precisely this determinant is for the various traces $R^J \to R^I$ does not appear to be in the literature, so that we can not make the required assumptions as explicit as Assumption 5.20. For the case of $\mathfrak{sl}_3$, the assumptions can be guaranteed by assuming that 3 is invertible (and this is required when Δ and $\Delta^\vee$ are bases).

Exercise 5.23. For the Frobenius trace ∂_g^y defined in §3.2, find dual bases for R^y over R^g under the assumption that 3 is invertible.

Definition 5.24. A *Frobenius realization* is a simple realization equipped with a choice of Frobenius structure ∂_I^J on $R^I \subset R^J$ for each finitary pair $J \subset I \subset S$. (In particular, we assume that such a Frobenius structure exists.) The Frobenius trace $R \to R^s$ must simply be ∂_s, as determined by the simple realization. We require that $\partial_I^J \partial_J^K = \partial_I^K$ whenever $K \subset J \subset I$. The Frobenius trace $\partial_I : R \to R^I$ must be proportional to $\partial_{\underline{w}_I}$ for some reduced expression of w_I.

This provides a partial Frobenius hypercube. As the Frobenius trace maps ∂_I^J are defined up to an invertible scalar, we can specify the structure by choosing the invariant μ_I^J. We do this by choosing a set of positive roots.

Definition 5.25. A *root realization* is a simple realization with the additional choice of a set of *finite positive roots* Φ. It must include the simple roots, and contain exactly one vector collinear with each element of $W_I \cdot \Delta_I$ for I finitary. In other words, whenever $\underline{w} = s_1 \ldots s_d$ is a reduced expression within a finite parabolic subgroup W_I, then Φ contains one element proportional to $s_1 s_2 \ldots s_{d-1}(\alpha_{s_d})$. Such an element is a *positive root for W_I* and lives in the subset $\Phi_I \subset \Phi$.

Definition 5.26. Assume that $\partial_{\underline{w}_I}$ is a Frobenius trace for some reduced expression of w_I, for each finitary I. Then any root realization will determine a unique Frobenius

realization for which the invariant μ_I of the extension $R^I \subset R$ is equal to the product of the positive roots for W_I. The invariant μ_I^J for $R^I \subset R^J$ is equal to the product of the positive roots for W_I which are not roots for W_J. We call this the *associated Frobenius realization*.

Finally, let us note when these scalar ambiguities do not arise. The following is proven in [8].

Definition 5.27. A simple realization is *odd-balanced* if, whenever m_{st} is odd for $s, t \in S$, then $a_{st} = a_{ts}$ satisfies the algebraic conditions to be equal to $-(\zeta + \zeta^{-1})$ for a primitive $2m_{st}$-th root of unity ζ.

Claim 5.28. *Whenever a realization is odd-balanced, there is a canonical choice of positive roots, and the Demazure operators satisfy the braid relations.*

5.6. Singular Bott–Samelson bimodules and diagrams. Fix a Frobenius realization of (W, S). One can now define $\mathcal{SBSBim}_{\mathrm{alg}}$ and $\mathcal{SSBim}$ precisely as in §5.1. One can also define a diagrammatic category $\mathcal{SBSBim}_{\mathrm{diag}}$ as in the previous sections (for $\mathfrak{sl}_2$ and $\mathfrak{sl}_3$).

Singular Soergel bimodules were defined by Williamson [45]. When the realization is *reflection faithful* (see [45] or the introduction), the category $\mathcal{SSBim}$ is well behaved. It categorifies the Hecke algebroid, an "idempotented" version of the Hecke algebra (and called the Schur algebroid in [45]). The size of morphism spaces between singular Soergel bimodules is governed by the standard trace on the Hecke algebroid; this is known as the *Soergel–Williamson Hom formula*. When $\Bbbk = \mathbb{C}$ and the realization is "geometric", the indecomposable objects in $\mathcal{SSBim}$ are (roughly speaking) the equivariant hypercohomologies of simple perverse sheaves on a Kac–Moody group, equivariant under various parabolic subgroups, and thus they descend to the Kazhdan–Lusztig basis of the Hecke algebroid.

The situation is analogous to the discussion of Rep and Fund in the previous chapter. $\mathcal{SSBim}$ is the genuine object of interest, while $\mathcal{SBSBim}$ is the combinatorial replacement. Finding a description of $\mathcal{SBSBim}$ by generators and relations has been elusive in general, and no progress has been made outside of rank 2 and type A. Rank 2 was accomplished in [8]. Williamson and I have a conjectural presentation of $\mathcal{SBSBim}$ in type A which we hope will shortly be available. Thus the presentation $\mathcal{SBSBim}_{\mathrm{diag}}$ for affine $\mathfrak{sl}_3$ is joint with Williamson; the proof of its correctness given in the appendix is original, but similar in style to arguments we attempted together.

When the realization is not reflection faithful, the category $\mathcal{SSBim}$ is not well behaved, and will not categorify the Hecke algebroid. An example will be when $\Bbbk = \mathbb{C}$ and $q \neq \pm 1$ is a root of unity, where the realization ceases to be faithful. However, the diagrammatic category $\mathcal{SBSBim}_{\mathrm{diag}}$ is still well behaved, and its morphism spaces always satisfy the Soergel–Williamson Hom formula. The presentation only encodes maps which occur generically; more morphisms may exist

between algebraic Soergel bimodules when q is a root of unity, but these morphisms are not in our diagrammatic presentation.

5.7. Maximally singular Bott–Samelson bimodules. Let us quickly mention the particular sub-2-category of $\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim}$ that we will use in this paper. Fix a Dynkin diagram Γ. Let (W, S) be the affine Weyl group of Γ, so that $S = \widetilde{\Gamma}$, and fix a Frobenius realization. Let $\Omega \cong \widetilde{\Theta} \subset S$ be as in §4.3. Maximal finitary subgroups have the form $I = S \setminus s$ for $s \in S$, and when $s \in \widetilde{\Theta}$ the parabolic subgroup W_I is isomorphic to the finite Weyl group.

Definition 5.29. Let $m\mathcal{S}\mathbb{S}\mathrm{Bim}$, the 2-category of *maximally singular Soergel bimodules*, denote the full sub-2-category of $\mathcal{S}\mathbb{S}\mathrm{Bim}$ whose objects are subsets of the form $I = S \setminus s$, for $s \in \widetilde{\Theta}$.

We can now state the Soergel Satake theorem.

Theorem 5.30 (Soergel Satake). *There is a 2-functor $\mathcal{F} \colon Rep^{\Omega}(\mathfrak{g}^{\vee}) \to m\mathcal{S}\mathbb{S}Bim(\widetilde{\Gamma})$, which is a degree 0 equivalence.*

Henceforth in this chapter we assume (W, S) is the affine Weyl group in type A. Then $m\mathcal{S}\mathbb{S}\mathrm{Bim}$ has an even better combinatorial replacement than $\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim} \cap m\mathcal{S}\mathbb{S}\mathrm{Bim}$.

Definition 5.31. Let $m\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim}$, the 2-category of *maximally singular Bott–Samelson bimodules*, denote the full sub-2-category of $\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim}$ defined as follows.

- The objects are subsets of the form $I = S \setminus s$, for $s \in \widetilde{\Theta}$.

- The 1-morphisms are monoidally generated by bimodules of the form ${}_{R^I} R_{R^J}^{I \cap J}(\ell(J) - \ell(I \cap J))$, where $I = S \setminus s$ and $J = S \setminus t$ for $s \neq t \in \widetilde{\Theta}$.

- The 2-morphisms are all graded bimodule maps.

Claim 5.32. *The Karoubi envelope of $m\mathcal{S}\mathbb{B}\mathbb{S}Bim$ is $m\mathcal{S}\mathbb{S}Bim$. In other words, every summand of a singular Bott–Samelson bimodule between maximally singular parabolic subsets is a summand of a maximally singular Bott–Samelson bimodule.*

Remark 5.33. This corresponds, under algebraic Satake, to the fact that all irreducible representations of $\mathfrak{sl}_n$ are summands of tensor products of miniscule representations. For general affine Weyl groups, more complicated bimodules are required to produce the non-miniscule fundamental representations.

Indecomposable singular Soergel bimodules are classified numerically, that is, by their behavior in the Grothendieck group. The proof of this claim is also numerical, showing that maximally singular Bott–Samelsons also have a certain upper-triangularity property in the Grothendieck group. For the classification of indecomposable bimodules, see [45].

Remark 5.34. One might also call the objects of $m\mathcal{S}\mathbb{B}\mathbb{S}\mathrm{Bim}$ by the name *reverse Bott–Samelson bimodules*. This is because, in their original context, Bott–Samelson

bimodules are R-bimodules obtained by restricting from R to R^s and inducing back to R. In other words, we begin at the minimal finitary parabolic subset $\emptyset$, move to a subminimal one $\{s\}$, and come back to the minimal one. For reverse Bott–Samelson bimodules we begin at a maximal finitary parabolic subset $S \setminus s$, move to a submaximal one $S \setminus \{s, t\}$, and return to a (different) maximal one $S \setminus t$.

We can now state the algebraic Satake theorem. We now reintroduce the subscript q to indicate that we are taking an exotic affine $\mathfrak{sl}_n$ realization.

Theorem 5.35 (Algebraic Satake). *There is a 2-functor $\mathcal{F}\colon Fund_q^\Omega \to m\mathcal{S}\mathbb{B}\mathbb{S}Bim_q$, which is a degree 0 equivalence.*

6. Reformulating geometric Satake

The reformulation from geometric Satake to Soergel Satake takes the form of three "replacements":

- Replacing monoidal categories with 2-categories,

- Replacing the affine Grassmannian with a partial flag variety for the affine Kac–Moody group,

- Replacing perverse sheaves with their global sections (i.e. their equivariant hypercohomology).

In the sections to come we will present this flow of ideas for a general complex semisimple lie algebra $\mathfrak{g}$. Nothing used is new or should be unfamiliar to the experts.

The transformation from Soergel Satake to algebraic Satake consists of replacing each side of the equivalence with a combinatorial additive subcategory. The technology for this procedure only exists currently in type A. This final replacement is discussed in greater detail in sections §4.1 and §5.7, but we also mention it briefly here.

6.1. Geometric Satake. We begin by recalling the results which are packaged under the name "geometric Satake." We refer the reader to [18] for better and more thorough introduction to this topic. We assume for the moment that all categories are $\mathbb{C}$-linear. For a group H acting on a space X, we let $\mathcal{P}_H(X)$ denote the category of H-equivariant perverse sheaves on X. When two groups H and K act on X, on the left and right respectively, we write $\mathcal{P}_{(H,K)}(X)$ for the category of biequivariant perverse sheaves.

Fix $T \subset G$ a maximal torus in a simple algebraic group, with Lie algebras $\mathfrak{h} \subset \mathfrak{g}$. Let $G^\vee$ (resp. $\mathfrak{g}^\vee$) denote the Langlands dual group (resp. Lie algebra). Let $\mathcal{K} = \mathbb{C}((t))$ and $\mathcal{O} = \mathbb{C}[[t]]$. As it is usually stated, geometric Satake is an equivalence between $G^\vee(\mathbb{C})$-rep and the category of $G(\mathcal{O})$-equivariant perverse sheaves on the affine Grassmannian $Gr_G = G(\mathcal{K})/G(\mathcal{O})$.

This equivalence intertwines the (non-equivariant) hypercohomology functor with the forgetful functor: the hypercohomology of an irreducible perverse sheaf is isomorphic, as a vector space, to the underlying vector space of the corresponding irreducible $G^\vee$ representation. In this context, the hypercohomology and forgetful functors are called *fiber functors*. The cohomological grading on hypercohomology transforms to the weight grading on the representation (with respect to some naturally-defined regular semisimple element). More interesting features arise when one considers the T-equivariant hypercohomology: instead of a vector space, one gets a graded module over the T-equivariant cohomology of a point, which is the polynomial ring $\mathbb{C}[\mathfrak{h}^*]$ (graded with $\deg \mathfrak{h}^* = 2$). Taking the quotient by the augmentation ideal (i.e. killing positive degree polynomials), one obtains the usual hypercohomology, though other specializations also hold interest. The usual tensor product on representations is intertwined with a convolution functor on perverse sheaves, though the geometric proof that convolution is symmetric monoidal is quite technical.

Remark 6.1. This long remark continues the discussion from the introduction (Remark 1.1) about the validity of discussing geometric Satake without the symmetric structure.

Given two semisimple monoidal categories with the same classification of indecomposables and the same decomposition of tensor products (i.e. the same Grothendieck ring), when are they equivalent? The answer is encoded in the associativity isomorphism. If these monoidal categories are also symmetric, then checking equivalence of symmetric monoidal categories also requires checking the commutation isomorphism. See [14, Problem 1.42.8 and following] for more details.

The numerical underpinnings of geometric Satake were shown by Lusztig in [32]. The characters of the irreducible perverse sheaves are encoded in Kazhdan–Lusztig polynomials for the affine Hecke algebroid. According to geometric Satake, these characters (evaluated at $v = 1$) should agree with the weight multiplicities in the corresponding irreducible $G^\vee$-representation. This was shown by Lusztig, who also demonstrated the link with orbit closures in the affine Grassmannian. Moreover, these results already encode a great deal about the monoidal structure on perverse sheaves. When one convolves two irreducible perverse sheaves, the BBD Decomposition Theorem [3] implies that the result splits into irreducibles according to its character, just as the splitting of a tensor product of G-representations is determined by its weight multiplicities. Thus perverse sheaves form a semisimple monoidal category with the same Grothendieck ring as $G^\vee(\mathbb{C})$-representations. The fiber functor (implicit also in Lusztig's work) also commutes with the convolution structure, making it easy to check the associativity isomorphism (note: this argument relies on semisimplicity). This is done explicitly in Ginzburg's preprint [18]. With Lusztig's results and the Decomposition Theorem in hand, the only further content in Geometric Satake is the comparison of commutativity isomorphisms. Without this comparison one has only

an equivalence of (vanilla) monoidal categories; this is what we accomplish in this paper.

An elementary proof of the commutativity isomorphism was attempted by Ginzburg in [18], but is not believed to be correct. A correct (and non-elementary) proof was given by Mirkovic and Vilonen [35], following ideas of Drinfeld. These results work with a "global approach," using the fusion product and the Beilinson–Drinfeld Grassmannian. This is the style also used by Gaitsgory in his quantum geometric Satake equivalence [17]. The author does not know how the global approach connects to the Soergel picture.

We do rely on Lusztig's numerical results, though we do not use the Decomposition theorem. Instead, we compare two monoidal categories by looking inside at strictly-monoidal additive subcategories (where the associativity isomorphism is by definition the identity map). To show two such categories are equivalent one has no convenient tricks (to the author's knowledge), but must compare the morphism algebras in full, as we did above. In essence, we have converted the questions of decomposition and associativity into a question about morphisms in this subcategory. The proof we use does not use the Decomposition Theorem or semisimplicity in any way, relying on Soergel's categorification results but not on the Soergel conjecture. In other words, we never need to study indecomposable singular Soergel bimodules and their numerical properties, only those bimodules which correspond to tensor products of fundamental representations. This is one advantage of working directly with generators and relations.

A proof independent of the Decomposition Theorem is required for the quantum algebraic Satake equivalence (and not just because there is no corresponding geometry)! After all, when q is a root of unity, the analog of the Soergel conjecture will fail, and the indecomposable objects will have the "wrong" characters, just as for representations of the (Lusztig form of the) quantum group at a root of unity. Nevertheless, tensor products of fundamental representations (and their Soergel analogs) continue to have the same numerics.

6.2. From monoidal categories to 2-categories. Our first step will be to apply a sequence of tautologies in order to reformulate geometric Satake as an equivalence of (strict) abelian 2-categories. In §4.2 it is explained how to begin with representations of $\mathfrak{g}^\vee$ and obtain Rep^Ω, an Ω-2-category. This will be one side of our equivalence. By conventional geometric Satake, the geometric category paired with representations of $\mathfrak{g}^\vee$ would be perverse sheaves on $Gr_{G_{\text{adj}}}$ in adjoint type. So we begin by extracting a 2-category from G_{adj}, after which we will explain how the same data is encoded within the affine Grassmannian for any G.

Suppose that $G = G_{\text{adj}}$ has adjoint type, so that $\pi_0(G(\mathcal{K}))$ is the group Ω defined in §4.3. As $G(\mathcal{O})$ is connected, this is also the component group of Gr_G. Let Gr_0 denote the *nulcomponent*, i.e. the component associated to the identity in Ω. Since Ω is abelian, conjugation by any element will preserve the nulcomponent $G(\mathcal{K})_0$. Any

irreducible perverse sheaf is supported on a single component, and perverse sheaves on different components admit no morphisms or extensions. Convolution will act on the "component support" of a perverse sheaf precisely via multiplication in Ω, so that $\mathcal{P}_{G(\mathcal{O})}(Gr_G)$ already forms an Ω-graded monoidal category, from which one can construct an Ω-2-category.

We choose to think of this Ω-2-category in a different way, altering equivariant structure rather than working on different components. Let us pick an equivariant perverse sheaf $\mathcal{F}$ supported on component $\xi \in \Omega$. Let $x \in G(\mathcal{K})$ be an element in the component ξ. Then $x^*\mathcal{F}$ is supported on the nulcomponent Gr_0, and the $G(\mathcal{O})$-equivariant structure on $\mathcal{F}$ produces a natural $(x^{-1}G(\mathcal{O})x)$-equivariant structure on $x^*\mathcal{F}$. The group $(x^{-1}G(\mathcal{O})x)$ depends on ξ, not on the choice of x, so we call this group G_ξ. Therefore, instead of considering $\mathcal{P}_{G(\mathcal{O})}(Gr_G)$, we can consider $\bigoplus \mathcal{P}_{G_\xi}(Gr_0)$ for all $\xi \in \Omega$.

Example 6.2. Suppose that $\mathfrak{g} = \mathfrak{sl}_4$, so that $G_{\text{adj}} = \mathbb{P}GL_4$, $G_{\text{sc}} = SL_4$, and $\Omega = \mathbb{Z}/4\mathbb{Z}$. The four subgroups $G_\xi \subset G_{\text{adj}}(\mathcal{K})$ are as follows:

$$
G_{\overline{0}} = \begin{pmatrix} \mathcal{O} & \mathcal{O} & \mathcal{O} & \mathcal{O} \\ \mathcal{O} & \mathcal{O} & \mathcal{O} & \mathcal{O} \\ \mathcal{O} & \mathcal{O} & \mathcal{O} & \mathcal{O} \\ \mathcal{O} & \mathcal{O} & \mathcal{O} & \mathcal{O} \end{pmatrix}, \qquad
G_{\overline{1}} = \begin{pmatrix} \mathcal{O} & t^{-1}\mathcal{O} & t^{-1}\mathcal{O} & t^{-1}\mathcal{O} \\ t\mathcal{O} & \mathcal{O} & \mathcal{O} & \mathcal{O} \\ t\mathcal{O} & \mathcal{O} & \mathcal{O} & \mathcal{O} \\ t\mathcal{O} & \mathcal{O} & \mathcal{O} & \mathcal{O} \end{pmatrix},
$$

$$
G_{\overline{2}} = \begin{pmatrix} \mathcal{O} & \mathcal{O} & t^{-1}\mathcal{O} & t^{-1}\mathcal{O} \\ \mathcal{O} & \mathcal{O} & t^{-1}\mathcal{O} & t^{-1}\mathcal{O} \\ t\mathcal{O} & t\mathcal{O} & \mathcal{O} & \mathcal{O} \\ t\mathcal{O} & t\mathcal{O} & \mathcal{O} & \mathcal{O} \end{pmatrix}, \qquad
G_{\overline{3}} = \begin{pmatrix} \mathcal{O} & \mathcal{O} & \mathcal{O} & t^{-1}\mathcal{O} \\ \mathcal{O} & \mathcal{O} & \mathcal{O} & t^{-1}\mathcal{O} \\ \mathcal{O} & \mathcal{O} & \mathcal{O} & t^{-1}\mathcal{O} \\ t\mathcal{O} & t\mathcal{O} & t\mathcal{O} & \mathcal{O} \end{pmatrix}.
$$

These groups are obtained from $G_{\overline{0}}$ by conjugation by a diagonal matrix with entries $(t, \ldots, t, 1, \ldots, 1)$. This diagonal matrix is always an element of $\mathbb{P}GL_4(\mathcal{K})$, though not usually an element of $SL_4(\mathcal{K})$. Note that the intersection $\cap_\xi G_\xi$ is an Iwahori subgroup, a fact which is special to type A.

Perverse sheaves which are $G(\mathcal{O})$-equivariant on $G(\mathcal{K})/G(\mathcal{O})$ have a natural monoidal structure from convolution. The collection of G_ξ-equivariant perverse sheaves on the nulcomponent of $G(\mathcal{K})/G_{\overline{0}}$ does not have an obvious monoidal structure (it is a module over the previous monoidal category). However, by considering the collection of all G_ξ-equivariant perverse sheaves on the nulcomponent of $G(\mathcal{K})/G_\eta$ for $\xi, \eta \in \Omega$, one obtains a 2-category where convolution again makes sense. We denote this 2-category by **P**.

For mental simplification, we choose to work with $\mathcal{P}_{(G_\xi, G_\eta)}(G(\mathcal{K})_0)$, i.e. bi-equivariant sheaves on the nulcomponent of $G(\mathcal{K})$, instead of $\mathcal{P}_{G_\xi}(G(\mathcal{K})_0/G_\eta)$, i.e. equivariant sheaves on the nulcomponent of the quotient. Because G_ξ orbits on $G(\mathcal{K})$ are infinite-dimensional, it is difficult to comprehend what is meant by equivariant perverse sheaves on $G(\mathcal{K})$. However, in any sense in which such things are defined, there will be an equivalence between these two categories. We choose the former

mostly for the notational symmetry, and because it makes the convolution structure "seem" more natural.

Now let us explain how the same 2-category **P** can be extracted from any Lie group G lifting $\mathfrak{g}$. The natural map $G \to G_{\mathrm{adj}}$ is a Galois cover with kernel Z, the center of G. It induces a map $G(\mathcal{O}) \to G_{\mathrm{adj}}(\mathcal{O})$ whose kernel is also Z, and a map $G(\mathcal{K}) \to G_{\mathrm{adj}}(\mathcal{K})$ whose kernel, restricted to the nulcomponent, is also Z. Thus the induced map $Gr_G \to Gr_{G_{\mathrm{adj}}}$ is an isomorphism on the respective nulcomponents. Thus the space on which we consider our sheaves does not depend on the choice of group.

The group Ω acts on $G(\mathcal{K})_0$ by group automorphisms, for any G. One can still think of this action as "conjugation" by an element in $G_{\mathrm{adj}}(\mathcal{K})$, but this automorphism may be an outer automorphism for $G(\mathcal{K})$. Thus one still has groups $G_\xi = \xi(G(\mathcal{O}))$ for each $\xi \in \Omega$. These groups will also be the preimages of the corresponding groups $G_\xi \subset G_{\mathrm{adj}}(\mathcal{K})_0$, and their orbits on $G(\mathcal{K})$ will be the preimages of orbits on $G_{\mathrm{adj}}(\mathcal{K})$.

Example 6.3. Continuing example 6.2 above, one can define $G_{\overline{k}}$ exactly as above for each $0 \le k \le 3$, with the assumption that one only considers matrices of that form lying within G. Conjugation by the diagonal matrix with entries $(t, \ldots, t, 1, \ldots, 1)$ is a (potentially outer) automorphism of $G(\mathcal{K})$.

Now compare the equivariant categories $\mathcal{P}_{G_\xi}(G(\mathcal{K})_0/G_\eta)$. As the group changes, the space $G(\mathcal{K})_0/G_\eta$ is unchanged, but the group which acts G_ξ will be altered by some central kernel Z. Clearly Z acts trivially, being contained in both G_ξ and G_η. However, Z is in the nulcomponent of any stabilizer, meaning that it will also act trivially on any equivariant structure. Thus the equivariant category is also independent of the choice of G.

Taking the sum over all $\mathcal{P}_{G_\xi}(G(\mathcal{K})_0/G_\eta)$, one obtains a 2-category which is equivalent to **P**. This 2-category is larger than merely $\mathcal{P}_{G(\mathcal{O})}(Gr_G)$; those G_ξ which are genuinely conjugate to $G(\mathcal{O})$ will encode perverse sheaves on other connected components, while the ones which are only outer-conjugate will allow for the extra data.

6.3. From affine Grassmannians to Kac–Moody groups.

Consider the Kac–Moody group G_{aff} associated to the affine Dynkin diagram $\widetilde{\Gamma}$, as in [28]. This is an analogue of $G(\mathcal{K})$. The essential difference is that G_{aff} has a torus whose dimension is bigger by two: it has a copy of $\mathbb{C}^*$ acting on $G(\mathcal{K})$ by loop rotation (sending $t \mapsto \lambda t$), and a copy of $\mathbb{C}^*$ arising from a central extension.

The set of removable vertices $\widetilde{\Theta}$ was defined in §4.3; recall that it is a Ω-torsor. For any subset of $\widetilde{\Gamma}$ one has a corresponding parabolic subgroup of G_{aff}, so for each $s \in \widetilde{\Theta}$ one has a parabolic subgroup P_s corresponding to $\widetilde{\Gamma} \setminus s$. Because each of these parabolic subgroups also contains the affine torus, the extra dimensions cancel in the quotient space G_{aff}/P_s, which is isomorphic to $Gr_0^\vee$ for any $s \in \widetilde{\Theta}$. However, the structure of P_t-equivariance on G_{aff}/P_s is stronger than the structure

of $G_\xi^\vee$-equivariance, because of the larger torus. The central extension $\mathbb{C}^*$, being central, will act trivially on G_{aff}/P_s, and being connected, will therefore act trivially on any equivariant structure. However, the loop rotation $\mathbb{C}^*$ will add new structure.

Thus our new 2-category of interest is the collection of P_t-equivariant perverse sheaves on G_{aff}/P_s for various $s, t \in \widetilde{\Theta}$, or alternatively, $\mathcal{P}_{(P_t, P_s)}(G_{\mathrm{aff}})$. This is a souped-up version of the original 2-category, having an extra "dimension" of equivariant structure.

Note that in type A, because $\widetilde{\Theta} = \widetilde{\Gamma}$, the intersection of all the P_s will be precisely the Borel subgroup. This is analogous to the situation in Example 6.2, where the intersection of the G_ξ was an Iwahori subgroup.

6.4. From perverse sheaves to bimodules.
We now describe some results of Soergel and Williamson, in order to motivate the theorem of Härterich we shall use.

Recall that the T-equivariant cohomology of a point $H_T(\star)$ is naturally isomorphic to $R_{\mathrm{fin}} = \mathbb{C}[\mathfrak{h}_{\mathrm{fin}}^*]$, a polynomial ring generated by dual Cartan subalgebra. We temporarily denote this ring A, not to be confused with the ring R we have used elsewhere in this paper. For T-spaces with a cellular filtration, passing from the T-equivariant cohomology $H_T(X)$ to the ordinary cohomology $H(X)$ is obtained by tensoring over $H_T(\star)$ with the one-dimensional module $A/A_+ = \mathbb{C}$, by the localization theorem. There is a natural action of the finite Weyl group W_{fin} on A, and the G-equivariant cohomology of a point is $A^{W_{\mathrm{fin}}}$. For the parabolic subgroup P_I associated to $I \subset S_{\mathrm{fin}}$, the P_I-equivariant cohomology of a point is $A^I = A^{W_I}$.

Therefore, given a (P_I, P_J)-equivariant perverse sheaf on G, its equivariant hypercohomology (i.e. equivariant pushforward to a point) will naturally be a graded (A^I, A^J)-bimodule.

Theorem 6.4. *The functor of equivariant hypercohomology from $\mathcal{P}_{(P_I, P_J)}(G)$ to graded (A^I, A^J)-bimodules is fully faithful on semisimple objects, for any $I, J \subset S$. (Note that this functor sends higher extensions to bimodule maps of nonzero degree, as the rings A^I are graded homologically.) The images of the semisimple objects are known as* singular Soergel bimodules *for W_{fin}.*

Remark 6.5. Unfortunately, this precise version of the theorem does not appear in the literature, so we will briefly discuss the history and sketch the proof. Soergel first studied the case $P_I = P_J = B$. A proof using ordinary (rather than equivariant) hypercohomology is in [39], while its equivariant analog is in [40, Proposition 3.4.4]. The general idea of these proofs is the same. First, one shows that pullbacks and pushforwards between partial flag varieties (or induction/restriction of equivariance) is sent by hypercohomology to induction and restriction of bimodules. This implies that the images of semisimple objects are Soergel bimodules. Next, one uses a parity-vanishing argument to deduce that the spectral sequence computing morphisms between hypercohomologies actually degenerates on the first page, implying that hypercohomology is faithful. Finally, one uses the Soergel Hom formula to

deduce that the dimensions of morphism spaces agree. For the general case a similar argument will apply, using Williamson's generalization of the Soergel Hom formula [45]; though not in the literature, this proof is known to experts.

The theorem was upgraded to affine Kac–Moody groups by Soergel's student Härterich [20]. Now one works with the Kac–Moody group G_{aff} and its parabolic subgroups P_I for $I \subset S$.

Theorem 6.6. *The functor of equivariant hypercohomology from $\mathcal{P}_{(P_I, P_J)}(G_{\mathrm{aff}})$ to graded $(H_{P_I}(\star), H_{P_J}(\star))$-bimodules is fully faithful on semisimple objects, for any $I, J \subset S$ with W_I, W_J finite. The images of the semisimple objects are known as singular Soergel bimodules for W_{aff}.*

Let T_{aff} denote the torus of G_{aff}, let T' denote the finite torus extended by loop rotation, and let S denote the set of affine reflections. Then $H_{T'}(\star)$ is naturally isomorphic to the polynomial ring $R = \mathbb{C}[\mathfrak{h}_{\mathrm{aff}}^*]$. This is the ring R we have used throughout this paper, defined in Example 5.9, or in §5.3 at $q = 1$. On the other hand $R_{\mathrm{aff}} = H_{T_{\mathrm{aff}}}(\star)$ is a larger ring, corresponding to the self-dual $n + 2$-dimensional realization also mentioned in Example 5.9. The subrings $H_{P_I}(\star)$ are isomorphic to R_{aff}^I, so that $H_{P_s}(\star)$ is isomorphic to $R_{\mathrm{aff}}^{W_{\mathrm{fin}}}$ for any $s \in \widetilde{\Theta}$.

The difference between R_{aff} and R is a polynomial ring in one variable, on which W acts trivially: namely, $H_{\mathbb{C}^*}(\star)$ for the central extension. All computations done in this paper work for R_{aff} as well as for R. Conversely, the theorem above will apply equally well to bimodules over the subrings of $H_{P_I}(\star)$ which ignore the central extension. Henceforth, we will work with singular Soergel bimodules for R rather than for R_{aff}.

The singular Soergel bimodules defined in §5.1 agree with the singular Soergel bimodules in this theorem. Therefore, the (semisimple part of the) 2-category of interest, $\mathcal{P}_{(P_t, P_s)}(G_{\mathrm{aff}})$, is equivalent to $m\mathcal{S}\mathbb{S}\mathrm{Bim}$, defined in §5.7. In fact, every perverse sheaf in this 2-category happens to be semisimple (warning: higher extensions do exist).

Here we see the payoff of having replaced the affine Grassmannian with the Kac–Moody group. The $G(\mathcal{O})$-equivariant cohomology of the point is the same as the G-equivariant cohomology, which is $R_{\mathrm{fin}}^{W_{\mathrm{fin}}}$. Meanwhile, the P_s-equivariant cohomology of the point is $R_{\mathrm{aff}}^{W_{\mathrm{fin}}}$, a larger ring. One could not repeat the constructions of the previous chapters with R_{fin}, because only one copy of W_{fin} acts on it, while for R_{aff} or R there are a number of copies of W_{fin} inside W_{aff} which can act.

6.5. Combinatorial subcategories.
When investigating an additive (monoidal) category, it may be difficult to compute all the morphisms between indecomposable objects. However, there may be a class of objects (closed under the tensor product) for which the computation of morphisms is tractable, and for which every indecomposable object appears as a summand. We call the corresponding full subcategory a *combinatorial replacement* for the original additive category.

A discussion of this philosophy and the known results for the special case of Rep($\mathfrak{g}$) or Rep$^\Omega$($\mathfrak{g}$) can be found in §4.1. In that section, it is discussed that tensor products of fundamental representations form a combinatorial replacement in type *A*, where the morphisms between such tensor products can be described by $\mathfrak{sl}_n$-webs. Unfortunately, morphisms between tensor products of fundamental representations remain undiscovered outside of type *A* and rank 2. The description of webs for $\mathfrak{sl}_2$ and $\mathfrak{sl}_3$ is given explicitly in §2.1 and §3.1 respectively.

The same philosophy applies to perverse sheaves, and in fact was a strong motivation for Soergel's original definition of Soergel bimodules. Soergel originally worked with the finite flag variety, or $\mathcal{P}_{(B,B)}(G)$. Finding irreducible perverse sheaves is difficult, but the decomposition theorem implies that a given irreducible perverse sheaf is a summand of the pushforward of the constant sheaf from a proper resolution of singularities. Schubert varieties in the finite flag variety have "combinatorial" resolutions known as *Bott–Samelson resolutions*, and the corresponding pushforwards are known as *Bott–Samelson sheaves*. Continuing to pushforward these sheaves to a point, one obtains the Bott–Samelson bimodules defined in Remark 5.34 or in [41]. Bott–Samelson bimodules are easy to compute with, and thus form an excellent combinatorial replacement for perverse sheaves. As the philosophy in §4.1 points out, returning from an algebraic description of Bott–Samelson bimodules to a similar description of Soergel bimodules is very difficult, and amounts to a computation of certain idempotents.

Williamson [45] continued to treat the case of partial flag varieties (or rather, equivariance under various parabolic subgroups), and defined singular Soergel bimodules. Singular Bott–Samelson bimodules, defined in §5.1 in terms of induction and restriction bimodules, form a combinatorial replacement in this context. For the specific case of $\mathcal{P}_{(P_s,P_t)}(G_{\mathrm{aff}})$ in type *A*, it turns out that there is also a further replacement which works, discussed in §5.7.

A. Diagrammatic proofs for $\mathfrak{sl}_3$

In this appendix, our primary goal is to prove the following theorem, which uses terminology from §5, and recalls the discussion of §5.6.

Theorem A.1. *Consider a Frobenius realization of affine $\mathfrak{sl}_3$ with exotic Cartan matrix. Let $\mathcal{SBS}Bim_{diag}$ denote the 2-category defined in §3.3, and let $\mathcal{SBS}Bim_{alg}$ denote the 2-category defined in §5.1. Then $\mathcal{SBS}Bim_{diag}$ categorifies the Hecke algebroid, and satisfies the Soergel–Williamson Hom formula. There is a faithful, essentially surjective 2-functor $\mathcal{G}\colon \mathcal{SBS}Bim_{diag} \to \mathcal{SBS}Bim_{alg}$ (defined at the end of §3.3). If the realization is reflection faithful, then $\mathcal{G}$ is an equivalence.*

Recall that when the realization is reflection faithful, Williamson's results imply that $\mathcal{SBS}Bim_{alg}$ will also categorify the Hecke algebroid, and satisfy the Soergel–

Williamson Hom formula. Therefore, faithfulness of $\mathcal{G}$ will imply fullness. When the realization is not even faithful (e.g. when q is a root of unity in $\Bbbk$), $\mathcal{SBSB}im_{alg}$ will fail to categorify the Hecke algebroid, and $\mathcal{G}$ will not be full. Theorem A.1 clearly implies Claim 3.16, and the reflection faithful condition explains the restriction on $\Bbbk$ given in that claim.

Our second goal is to prove Theorem 3.21, which states that the functor $\mathcal{F}\colon \mathrm{Fund}_q^\Omega \to m\mathcal{SBSB}im_q$ is a degree zero equivalence. This is independent of the choice of realization (e.g. it still holds when q is a root of unity).

Both results were proven for $\mathfrak{sl}_2$ in [8]. After setting up and proving some diagrammatic lemmata, the main arc of the proof is completely analogous to the $\mathfrak{sl}_2$ case. Like all diagrammatic proofs, it is extremely delicate.

Henceforth, to simplify notation, $\mathcal{D}$ will denote the diagrammatic 2-category $\mathcal{SBSB}im_{diag}$, and $\mathcal{C}$ will denote the algebraic category $\mathcal{SBSB}im_{alg}$. A *diagram* refers to a singular Soergel diagram.

A.1. Preliminaries. We recall and amplify our conventions from §3. Let W be the affine Weyl group of $\mathfrak{sl}_3$. Let $S = \{r, b, y\}$ denote the set of *primary colors*, which can be combined to form the *secondary colors* $\{g, o, p\}$. The color *brown*, which combines all three primary colors, is forbidden as a region label in a diagram. For a subset $K \subset S$, a diagram is said to only have the primary colors in K if every region label is contained in K.

Definition A.2. Let $K \subsetneq S$ be a subset of the primary colors. Let $\mathcal{D}(K)$ denote the 2-category defined as in Definition 3.12 except that

- 2-morphisms may only have primary colors in K (this also restricts the objects and 1-morphisms),
- one only imposes the relations with primary colors in K.

Let $\mathcal{C}(K)$ denote the algebraic subcategory of singular Bott–Samelson bimodules defined as in Definition 3.15, except that one only allows sets I, J contained in K.

Note that both categories use the same base ring R, defined in §3.2, regardless of the subset K (we consider polynomials to be colorless).

There is a natural 2-functor $\iota_K\colon \mathcal{D}(K) \to \mathcal{D}$, which we will eventually show is fully faithful. At the moment, it is not even clear that ι_K is faithful, because $\mathcal{D}$ has more relations. We prove this very soon - that relations between diagrams with extraneous primary colors do not affect diagrams without those primary colors. It is far less obvious that ι_K should be full - that diagrams without certain primary colors on the boundary should be in the span of diagrams without those primary colors anywhere.

The Dihedral Cathedral [8] is an in-depth study of Soergel bimodules for dihedral groups. Its appendix treats odd-unbalanced realizations such as the exotic realization of affine $\mathfrak{sl}_3$. We will quote this paper for many "two-color" results in this proof, such as the following proposition.

Proposition A.3. *Let $K \subsetneq S$. Then the analogous 2-functor $\mathcal{G}_K$ is well defined, full, and faithful from $\mathcal{D}(K)$ to $\mathcal{C}(K)$.*

Proof. Observe that any exotic realization of affine $\mathfrak{sl}_3$ restricts to a reflection faithful representation of any dihedral parabolic subgroup. Now the result follows from the Cathedral. □

All the relations which define $\mathcal{D}$ involve at most two primary colors. Therefore, Proposition A.3 implies that these relations hold algebraically in $\mathcal{C}$, and the 2-functor $\mathcal{G}$ is well defined. Moreover, $\mathcal{G}_K = \mathcal{G} \circ \iota_K$ is faithful, which implies that ι_K must be faithful.

The proof that ι_K is full will comprise the bulk of this appendix, after which certain categorification considerations will complete the proof of Theorem A.1. It amounts to a "color-removal algorithm."

We say that a diagram *reduces* to another class of diagrams if it can be rewritten, using the relations, as a linear combination of diagrams in this class. A region is *external* if it touches the boundary, and *internal* otherwise. A diagram has *featureless boundary* if it has a single exterior region. A *loop* is a closed 1-manifold in a diagram (which may intersect strands of other colors), and may be oriented either clockwise or anticlockwise.

Proposition A.4. *Suppose that $K \subsetneq S$. A diagram whose external regions only contain primary colors in K will reduce to diagrams where every region only contains primary colors in K.*

This is a restatement of the fullness of ι_K. We now summarize the argument.

Suppose that $K = \{b, y\}$. The region colors we wish to remove, namely r, o, p, we will call "reddish." Recall that we have also colored the strands in our diagrams, so that a red strand separates reddish regions from non-reddish regions. Any diagram without external reddish regions but with some internal reddish region must have a clockwise red loop. Color-removal is equivalent to the statement that diagrams without clockwise loops will span all diagrams in $\mathcal{D}$.

A diagram without any crossings or closed 1-manifolds is called a *crossingless matching (with boxes)*. It is built entirely of cups, caps, and boxes. A crossingless matching with featureless boundary must be a box. One nice feature of crossingless matchings on simply-connected domains is that they necessarily have a cup/cap on the boundary (i.e. two adjacent points connected by a strand). However, this does not hold true of crossingless matchings on non-simply-connected domains, such as a configuration of radii on an annulus.

Lemma A.5. *Suppose that every region in a diagram (on a simply-connected domain) is reddish. Then the diagram reduces to a crossingless matching.*

Proof. As the primary color red is never removed there can be no red strands, only blue and yellow strands. These can never cross, lest there exist a brown region. So we

need only show that any closed loop can be removed. Suppose there is a closed blue or yellow loop. Choosing an innermost loop (i.e. no loops in its interior; this makes sense since loops can not cross), we replace it and its interior with a box using (3.14a) or (3.14b). Repeating this process, we remove all closed components. □

To remove a clockwise red loop, one step of the algorithm will be to simplify its interior. A naive approach would go as follows. Suppose one could remove all closed loops from the interior of the red loop. Then this interior is a diagram where every region is reddish, and by the above lemma this interior is a crossingless matching and must have a cup/cap on the boundary. Applying (3.15) one should be able to pull this cup/cap out of the yellow loop, reducing the number of strands which intersect the red loop, and allowing an inductive argument. (Boxes are easy to deal with and do not complicate the argument overmuch, thanks to (3.15d) and similar relations.)

Although this style of argument is useful, the supposition is too naive. A complicating feature of the interior is that it may contain numerous anticlockwise red loops, and anticlockwise loops cannot be removed in general! The resulting reddish domain is no longer simply connected, and we can not proceed as above. Instead, the interior of an anticlockwise loop can be simplified, and this is where we must begin a more sophisticated argument. For this purpose, we prove some lemmata about Bott–Samelson objects.
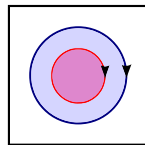
A.2. Bott–Samelson objects.

Definition A.6. A *BS object* in $\mathcal{SBS}\mathrm{Bim}$ is an object of the form $\emptyset s_1 \emptyset s_2 \emptyset \cdots \emptyset s_d \emptyset$. We say that a diagram on the disk has *BS boundary* if the boundary (read around the circle) is a BS object.

The following map has an alternating two-color BS boundary, and will be denoted v_k.

$$\tag{A.1}$$

Pictured is the case $k = 6$. When $k = 0$, it looks like below.

In v_0, switching the roles of blue and red would make no difference, as (3.14b), (2.8b), and (3.13) imply that v_0 is equal to the polynomial μ_p. We continue to denote the map v_k, regardless of which two primary colors are used.

The diagram v_3 is special, which relates to the fact that the order of rb in W is 3. In [8, Section 6.1.2] it was shown that every v_k for $k \geq 4$ can be rewritten as a composition of the maps v_3, together with cups and caps. It was also shown that every v_k for $k < 3$ can be rewritten in terms of crossingless matchings. This can be seen directly: for $k = 2$ this is (3.16), for $k = 1$ this is (3.15c), and for $k = 0$ this follows as above.

The following claim is one of the results in the Cathedral.

Claim A.7. *Any diagram with at most two primary colors and BS boundary reduces to a composition of cups, caps, boxes, and the diagram v_3. If the BS boundary is empty (i.e. featureless and white), then the diagram reduces to a box. If the BS boundary is non-empty then after reduction the diagram will either have a cap/cup on the boundary, or will be equal to v_k for some $k \geq 3$ (with boxes allowed in any white regions).*
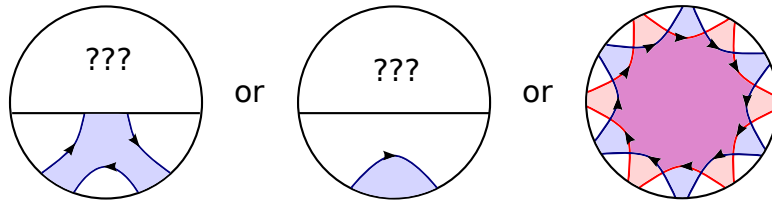
In particular, if the BS boundary does not alternate between the two colors, but instead has repetition $\emptyset b \emptyset b \emptyset$, then there is guaranteed to be a cup/cap on the boundary as in the first option above. One way to see this is to apply (2.8d) to the intermediate $b \emptyset b$.

Lemma A.8. *Consider the neighborhood of an anticlockwise red loop, with no reddish regions inside. It can be reduced to a diagram where the red loop is absent (i.e. every region is reddish), or to a diagram as below for $k \geq 3$.*

$$(A.2)$$
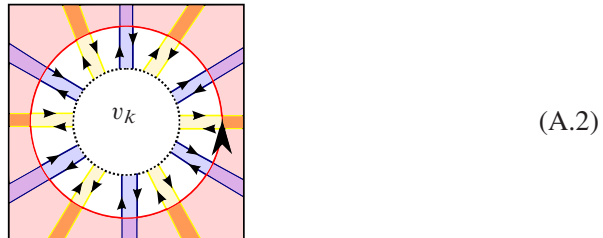
Proof. Consider the subdiagram which is the interior of the red loop. The boundary of this subdiagram can not contain the color green, lest the exterior of the loop contain a forbidden brown region. Therefore, the boundary of the interior is either featureless or is a non-empty BS object. Either way, we may apply Claim A.7 to the interior. When the boundary is featureless, the interior reduces to a box, and then we may apply (2.8c) or (3.14a) to remove the loop entirely. Otherwise, either the

interior is v_k for $k \geq 3$ as desired, or there is a cup/cap on the boundary. In the latter case, we apply (3.15) to pull the cup/cap out of the red loop, reducing the size of the Bott–Samelson boundary by 1. By induction, this particular anticlockwise loop either disappears or has the desired form.   □

Note that the diagram (A.2) above is an example of a 2-morphism with an anticlockwise loop which can not be reduced to diagrams without loops. These kinds of diagrams can appear within the image of $\mathcal{F}$ applied to non-elliptic webs (c.f. Claim 3.5). The non-elliptic condition is equivalent to the condition $k \geq 3$.

A.3. Loop and color removal. We work one color at a time.

Lemma A.9. *Any diagram reduces to a diagram without clockwise red loops. Any diagram with featureless reddish boundary reduces to a box.*

Proof. Consider the neighborhood of a red loop. Either no other strands intersect it (i.e. it is featureless), or the strands which intersect it form a non-empty yellow-and-blue BS object.


Suppose that this BS object has repetition, i.e. $\emptyset b \emptyset b \emptyset$. Then applying (2.8d) and (3.15), we can reduce the number of intersecting strands.
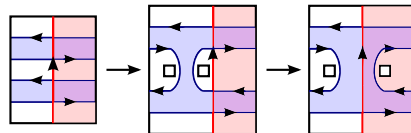
Ergo, up to placing boxes in various regions, we can assume that the intersection is either empty or an alternating BS object. Like every step of this algorithm, this procedure is local, and will not interfere with other operations.

Consider the case of an innermost clockwise red loop, and let D denote the diagram in its interior. Using the above paragraph, we may assume either that the boundary of D is featureless, or that it consists of an alternating yellow-and-blue BS object with red added; i.e. it is something like $rprorpror$ read cyclically around the disk. In the latter case, we assume that there are j copies of p and j copies of o, for

some $j > 0$. Meanwhile, D has some number $n \geq 0$ of anticlockwise red loops in its interior, so that its reddish zone appears as below (only red is shown).
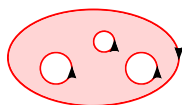
The arguments below will also apply to any diagram D with featureless reddish boundary, regardless of whether or not it was the interior of a clockwise loop.

We now prove by induction that any such diagram D reduces to a diagram where $n = 0$, and that if the diagram has featureless boundary then it reduces to a box. Moreover, if D is the interior of a red loop and $n = 0$, then it reduces to a diagram where $j = 0$, i.e. to the featureless case. This will finish the proof of the Lemma: when $n = j = 0$ and D is the interior of a red loop, D reduces to a box which can be slid out using (2.8b) or (3.14c), and then the loop itself can be eliminated using (3.13) or (3.14b).

We wish to argue that the diagram in the reddish zone can be reduced to a crossingless matching on the n-punctured disk (i.e. it has no crossings or closed blue or yellow loops). The lack of crossings is clear, and any contractible blue or yellow loop can be removed as in the argument of Lemma A.5. Suppose that $n > 0$ and there is a blue or yellow non-contractible loop within the reddish region, and choose an innermost one.
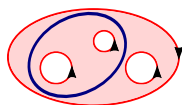

Its interior is another diagram D' with featureless reddish boundary, where the reddish zone is a crossingless matching. Applying the arguments below to D' will replace D' with a box. Then we can repeat with any other non-contractible loops in a local fashion. Thus we assume that the reddish zone is a crossingless matching.

Consider the base case when $n = 0$. Then the reddish zone is simply-connected. If its boundary is featureless then the reddish zone just contains a box. If not, then any crossingless matching on a simply-connected region has a cup/cap on the boundary. Using (3.15) will lower j. To be very precise, this will remove one index from the alternating BS object on the boundary, producing a repetition, which can then be removed as in the first paragraph; the overall effect is to lower j by one. (This phenomenon will repeat often enough that we leave it unstated — whenever we remove a single index from the BS boundary, we say that it lowers j by one, leaving unstated the fact that one must then remove the repeated color.) This handles the base case.

Now suppose that $n > 1$. We can assume by Lemma A.8 that each anticlockwise red loop has a neighborhood which looks like (A.2) for some $k \geq 3$. We call this a

puncture of size k. We now proceed by induction on the size k of any given puncture, or the number j of the exterior. We do all our computations within the reddish zone, treating the interiors of the anticlockwise red loops (the punctures) as black boxes.

Consider what any red (not reddish!) region will look like. Instead of reading a list of region colors around a boundary, as we have done in previous arguments, now consider the list of strand colors around the boundary of the region. They must alternate between red and blue/yellow. If the color blue appears twice (even with yellow in between) applying (3.14d) will chop the red region into smaller red regions.
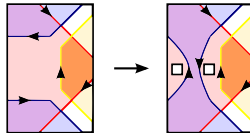

If the red region is a bigon, then one can apply (3.15) as usual to pull the cup/cap either out of the exterior red loop, or into one of the punctures, and use induction. Thus we can assume each red region is a square, having one blue and one yellow strand as parallel walls.
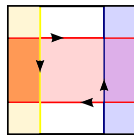

The non-reddish parts of this diagram all occur either inside a puncture or outside the exterior clockwise loop. Thus we can think of red regions as "edges" which connect punctures or the exterior loop; we will use such a graph below.

Consider what any purple (resp. orange) region will look like. Its boundary alternates between red and blue strands, and it could be a bigon, square, hexagon, octagon, etc. We now argue that the diagram will simplify unless the purple region has at least 6 sides. If it is a bigon, i.e. if a cup/cap is attached to one of the punctures or the exterior, then applying (3.15) will reduce k or j by one. If it is a square, connecting two punctures or a puncture with the exterior, then one has a local picture like
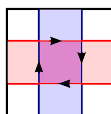

where the two red strands belong to different red loops. Applying (3.16), there are two terms which remain.
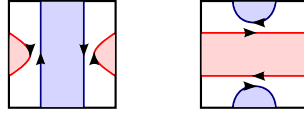
In the second term, the value of k of each puncture (or j for the exterior) is reduced by one. In the first term, the two punctures are fused into one puncture, or the puncture is fused with the exterior loop, and either case reduces n by one.

Finally, if the purple region is a square connecting the exterior with itself or connecting any puncture with itself, then one has a local picture like
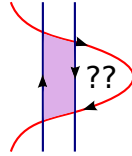
The ??? region can be quite complicated, having punctures, additional strands crossing the red strand, etc. However, once again, applying (3.16) will yield two terms. In the second term, the value of either j or some k will decrease. In the first term, either the original exterior loop is split into two smaller loops, or a new clockwise loop is created within a puncture (containing the ??? region). Either way, the clockwise loop containing the ??? region will be strictly smaller (by one of our inductive criteria) than the original loop, and thus will disappear. The remainder will also have had either j or some k decreased.

Thus induction applies unless every red region is a square and every purple and orange region is a hexagon or larger. To finish the proof, we make a planar Euler characteristic argument to prove that this never occurs.

Let n_k be the number of punctures with boundary of size k, as above. Let p_k (resp. o_k) denote the number of purple (resp. orange) regions with $2k$ sides, and r the number of red regions. Contract every puncture into a vertex, and if $j > 0$ then contract the outer loop into a vertex, so that one obtains a graph embedded in S^2. Now, each puncture of size k yields a $4k$-valent vertex, with $2k$ red regions, k purple regions, and k orange regions adjoining. After this contraction, the purple regions enumerating p_k have k adjoining edges, and the red regions are bigons. Therefore, by counting purple (resp. orange) regions paired with an adjoining vertex, we obtain

$$\sum_k k p_k = j + \sum_k k n_k = \sum_k o_k. \tag{A.3}$$

Counting red regions paired with a vertex, we obtain

$$2r = \sum_k 2k n_k + 2j, \tag{A.4}$$

so that r is also equal to the quantity in (A.3).

Each edge appears adjacent to a single red region, so the number of edges is $2r$. The number of vertices is $1 + \sum n_k$ if $j > 0$, and $\sum n_k$ if $j = 0$. The number of regions is $r + \sum p_k + \sum o_k$. Thus (when $j > 0$) we have:

$$2 = V - E + R = (1 - j) + \sum(1 - k)n_k + \sum p_k + \sum o_k.$$

When $j = 0$, the $(1-j)$ term is ignored. Now, by assumption each $k \geq 3$. Therefore

$$\sum p_k + \sum o_k \leq \frac{1}{3} \sum (kp_k + ko_k) = \frac{2}{3}(j + \sum k n_k),$$

so that

$$2 \leq (1 - \frac{j}{3}) + \sum(1 - \frac{k}{3})n_k.$$

However, $(1 - \frac{k}{3})$ is nonpositive, and the term $(1 - \frac{j}{3})$ is at most one (and disappears when $j = 0$). This is a contradiction, and concludes the proof of Lemma A.9. □

Corollary A.10. *Let $K \subsetneq S$. Then ι_K is full.*

Proof. Suppose that K does not contain the primary color red. Then any 2-morphism between 1-morphisms in $\mathcal{D}(K)$ will have no reddish exterior regions. By removing all clockwise red loops, we thereby remove the primary color red entirely. The resulting 2-morphism is in the image of ι_K. □

The lemma essentially gave an algorithm to remove any given innermost clockwise loop from a diagram. However, the procedure did screw up the topology of the differently-colored strands which intersected that loop, and thus can create (blue or yellow) loops as well as destroy them. There does not seem to be an easy way to deduce that one can remove all clockwise loops of different colors at once, directly from the lemma. Instead, we use a circuitous route.

A.4. Arguments from categorification.

Lemma A.11. *2-morphisms in $\mathcal{D}$ between BS objects are spanned by diagrams where the only instances of purple (resp. orange, green) have neighborhoods which look like v_k for $k \geq 3$.*

Proof. Consider a purple region in such a diagram, which must be internal, and its interior D. Then D is a diagram with featureless reddish boundary, so by Lemma A.9 it reduces to a box. In this fashion, we can assume every purple region is simply connected, and has no yellowish regions in its neighborhood. This neighborhood is now a morphism in $\mathcal{D}(r, b)$, and the result follows from the Cathedral. □

Recall that each v_k for $k > 3$ can be rewritten using v_3.

BS objects form a monoidal category, which is supposed to categorify the Hecke algebra **H** of W. In [12], a diagrammatic category $\mathcal{D}_{BS}$ (called $\mathcal{D}$ in that paper) is described by generators and relations (for an arbitrary Coxeter group), which is meant to encode the morphisms between BS bimodules. In particular, there is a 2-functor $\mathcal{I}$ from $\mathcal{D}_{BS}$ to $\mathcal{D}$, defined on dihedral parabolic subgroups in [8], whose 1-morphism image consists of the BS objects. The 2-morphism image consists of those morphisms between BS objects generated by cups and caps, as well as the map v_3.

Proposition A.12. *The natural 2-functor $\mathcal{I}$ from $\mathcal{D}_{BS}$ to $\mathcal{D}$ is fully faithful. The 2-functor $\mathcal{G}$ is faithful from objects in the image of $\mathcal{I}$.*

Proof. Lemma A.11 indicates that $\mathcal{I}$ is full. The composition $\mathcal{G} \circ \mathcal{I}$ was shown to be faithful in [12], so that $\mathcal{I}$ must also be faithful, and $\mathcal{G}$ must be faithful from the image. □

The wonderful fact about Soergel bimodules and their diagrammatic categories is that the behavior of morphism spaces is governed by the BS objects.

Corollary A.13. *The 2-functor $\mathcal{G}$ is faithful. 2-morphism spaces in $\mathcal{D}$ satisfy the Soergel–Williamson Hom Formula [45, Theorem 7.2.2], and $\mathcal{D}$ categorifies the Hecke algebroid.*

Proof. This is a quick summary of an argument made almost verbatim in the Cathedral. Recall the definitions of the *Hecke algebra* **H** (resp. *Hecke algebroid* $\mathfrak{H}$) from [45] or from [8, Section 2.4]. Recall the definition of a *potential categorification* of the Hecke algebra (resp. algebroid) from [8, Sections 2.3 and 2.4.3]. Proposition A.3) implies that $\mathcal{D}$ categorifies all the relations in $\mathfrak{H}$ coming from dihedral parabolic subgroups. As all the relations in $\mathfrak{H}$ come from finite parabolic subgroups, we see that $\mathcal{D}$ is a potential categorification of $\mathfrak{H}$, and therefore induces a trace map on $\mathfrak{H}$. Any trace map on $\mathfrak{H}$ is determined by its evaluation on the subalgebroid **H**, or in other words, the sizes of arbitrary morphism spaces in $\mathcal{D}$ (and their behavior under $\mathcal{G}$) are determined by the morphism spaces between BS objects. Thus Proposition A.12 implies that $\mathcal{G}$ is faithful everywhere, as desired (c.f. [8, Sections 2.4.3 and 5.4.2]). □

Theorem A.1 follows immediately. So does the fact that all clockwise loops can be removed simultaneously. One may also assume (up to reduction) that every region in a diagram is simply-connected.

A.5. Proof of Theorem 3.21. Recall the definition of $m\mathcal{S}\mathbb{B}\mathbb{S}\text{Bim}$ and the functor $\mathcal{F}: \text{Fund}^{\Omega} \to m\mathcal{S}\mathbb{B}\mathbb{S}\text{Bim} \subset \mathcal{D}$ from §3.4. Our goal is to show that $\mathcal{F}$ is a degree zero equivalence. Much of what needs to be shown is already implied by Theorem A.1.

It is an exercise in the Hecke algebroid and the basic properties of $\mathcal{S}\mathbb{S}$Bim to show that $\mathcal{F}$ is essentially surjective up to grading shift.

We wish to show that $\mathcal{F}$ is fully faithful in degree zero. However, morphism spaces in $\mathcal{D}$ have a graded dimension governed by the Soergel–Williamson Hom Formula. The fact that morphisms between 1-morphisms in the image of $\mathcal{F}$ are positively graded, and that their degree 0 morphisms have the correct size, is a purely numerological one. It was first observed by Lusztig in the seminal paper [32] from which the idea of Geometric Satake was birthed. Thus it remains to show that $\mathcal{F}$ is full, which is a purely diagrammatic argument.

A diagram representing a morphism in $m\mathcal{S}\mathbb{B}\mathbb{S}$Bim has no exterior white regions. Consider a given white region. Using Proposition A.4 we can remove any color on the interior of the white region. This leaves a box, and nothing more. The region is delineated by k primary-colored strands, with no two of the same color being adjacent. If $k = 1$ the region can be removed using (2.8c). If $k = 2$ the region can be removed using (3.15b). If $k = 3$ then a neighborhood of this triangular white region is precisely $\mathcal{F}$ applied to the trivalent vertex
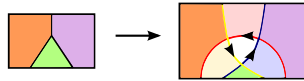
though possibly with a box in the white region. If $k > 4$ then there must be a repetition of colors. Using a now-familiar argument, one applies (2.8d) to this repetition, which splits the white region into two smaller regions. Therefore, we may assume that every white region is a triangle.

Remove a neighborhood of each white region, and consider the remaining diagram on the punctured disk. The remaining diagram has no crossings, lest there exist a brown region. We wish to show that it is in the image of $\mathcal{F}$ applied to a diagram on the punctured disk without any trivalent vertices. In other words, we must show that every red region has exactly one purple and one orange region adjacent to it (and something similar for blue and yellow regions). Note that, because of the behavior of the diagram at the boundary and near each puncture, the regions surrounding the red region must alternate between orange and purple.

Consider any red region. Exactly as in the proof of Lemma A.9, if it abuts two purple regions then one may fuse them with (3.14d), splitting the red region into two smaller regions. This can be repeated until each red region has exactly one purple and one orange region adjacent to it, as desired.

Therefore, every morphism between objects in $m\mathcal{S}\mathbb{B}\mathbb{S}$Bim is in the span of diagrams in the image of $\mathcal{F}$, but with boxes in various regions. Only when there are no boxes will the diagram have degree zero. This concludes the proof of fullness. In fact, we have proven a slightly stronger statement.

Proposition A.14. *Morphism spaces (of arbitrary degree) in $m\mathcal{S}\mathbb{B}\mathbb{S}$Bim between objects in the image of $\mathcal{F}$ are generated by diagrams in the image of $\mathcal{F}$, under the operation of placing boxes in various regions.*

References

[1] H. H. Andersen and D. Tubbenhauer, Diagram categories for $\mathbf{U}_q$ tilting modules at roots of unity, *preprint*, 2014. arXiv:1409.2799

[2] D. Bar-Natan and S. Morrison, The Karoubi envelope and Lee's degeneration of Khovanov homology, *Algebr. Geom. Topol.*, **6** (2006), 1459–1469. Zbl 1130.57012 MR 2253455

[3] A. A. Beĭlinson, J. Bernstein, and P. Deligne, Faisceaux pervers, in *Analysis and topology on singular spaces, I (Luminy, 1981)*, 5–171, Astérisque, 100, Soc. Math. France, Paris, 1982. Zbl 0536.14011 MR 0751966

[4] J. Bernstein, I. Frenkel, and M. Khovanov, A categorification of the Temperley–Lieb algebra and Schur quotients of $U(\mathfrak{sl}_2)$ via projective and Zuckerman functors, *Selecta Math. (N.S.)*, **5** (1999), no. 2, 199–241. Zbl 0981.17001 MR 1714141

[5] S. Cautis, J. Kamnitzer, and S. Morrison, Webs and quantum skew Howe duality, *Math. Ann.*, **360** (2014), no. 1-2, 351–390. Zbl 06357947 MR 3263166

[6] C. W. Curtis, N. Iwahori, and R. Kilmoyer, Hecke algebras and characters of parabolic type of finite groups with (B, N)-pairs, *Publications Mathématiques de l'IHÉS*, **40** (1971), no. 1, 81–116. Zbl 0254.20004 MR 0347996

[7] B. Elias, Light ladders and clasp conjectures, *preprint*, 2015. arXiv:1510.06840

[8] B. Elias, The two-color Soergel calculus, *Compos. Math.*, **152** (2016), no. 2, 327–398. Zbl 06563391 MR 3462556

[9] B. Elias and N. Libedinsky, Soergel bimodules for universal Coxeter groups, to appear in *Transactions of the American Mathematical Society*. arXiv:1401.2467

[10] B. Elias and G. Williamson, Singular Soergel bimodules in type A, *in preparation*.

[11] B. Elias and G. Williamson, The Hodge theory of Soergel bimodules, *Ann. of Math. (2)*, **180** (2014), no. 3, 1089–1136. Zbl 1326.20005 MR 3245013

[12] B. Elias and G. Williamson, Soergel calculus, *Represent. Theory*, **20** (2016), 295–374. MR 3555156

[13] B. Elias, G. Williamson, and N. Snyder, On cubes of Frobenius extensions, *preprint*. arXiv:1308.5994

[14] P. Etingof, S. Gelaki, D. Nikshych, and V. Ostrik, *Tensor categories*, lecture notes. Available at: http://www-math.mit.edu/~etingof/tenscat.pdf

[15] P. Etingof, D. Nikshych, and V. Ostrik, On fusion categories, *Ann. of Math. (2)*, **162** (2005), no. 2, 581–642. Zbl 1125.16025 MR 2183279

[16] I. Frenkel, C. Stroppel, and J. Sussan, Categorifying fractional Euler characteristics, Jones–Wenzl projector and $3j$-symbols, *Quantum Topol.*, **3** (2012), no. 2, 181–253. Zbl 1256.17006 MR 2901970

[17] D. Gaitsgory, Twisted Whittaker model and factorizable sheaves, *Selecta Mathematica*, **13** (2008), no. 4, 617–659. Zbl 1160.17009 MR 2403306

[18] V. Ginzburg, Perverse sheaves on a loop group and Langlands' duality, *preprint*, 1995. arXiv:alg-geom/9511007

[19] F. M. Goodman and H. Wenzl, Ideals in the Temperley–Lieb category, *preprint*, 2002. arXiv:math/0206301

[20] M. Härterich, *Kazhdan–Lusztig–Basen, unzerlegbare Bimoduln und die Topologie der Fahnenmannigfaltigkeit einer Kac–Moody–Gruppe*, Ph.D. thesis, Albert-Ludwigs-Universität, Freiburg, 1999.

[21] J. E. Humphreys, *Reflection groups and Coxeter groups*, Cambridge Studies in Advanced Mathematics, 29, Cambridge University Press, Cambridge, 1990. Zbl 0725.20028 MR 1066460

[22] V. F. R. Jones, The annular structure of subfactors, in *Essays on geometry and related topics, Vol. 1, 2*, 401–463, Monogr. Enseign. Math., 38, Enseignement Math., Geneva, 2001. Zbl 1019.46036 MR 1929335

[23] L. H. Kauffman, State models and the Jones polynomial, *Topology*, **26** (1987), no. 3, 395–407. Zbl 0622.57004 MR 0899057

[24] D. Kazhdan and G. Lusztig, Representations of Coxeter groups and Hecke algebras, *Invent. Math.*, **53** (1979), no. 2, 165–184. Zbl 0499.20035 MR 0560412

[25] M. Khovanov, A. D. Lauda, M. Mackaay, and M. Stošić, Extended graphical calculus for categorified quantum sl(2), *Mem. Amer. Math. Soc.*, **219** (2012), no. 1029, vi+87pp. Zbl 1292.17013 MR 2963085

[26] R Kilmoyer, *Some irreducible representations of a finite group with BN-pair*, Ph.D. thesis, Massachusetts Institute of Technology, 1969. MR 2940017

[27] D. Kim, *Graphical calculus on representations of quantum Lie algebras*, Ph.D. thesis, University of California, Davis, ProQuest LLC, Ann Arbor, MI, 2003. MR 2704398

[28] S. Kumar, *Kac-Moody groups, their flag varieties and representation theory*, Progress in Mathematics, 204, Birkhäuser Boston Inc., Boston, MA, 2002. Zbl 1026.17030 MR 1923198

[29] G. Kuperberg, Spiders for rank 2 Lie algebras, *Comm. Math. Phys.*, **180** (1996), no. 1, 109–151. Zbl 0870.17005 MR 1403861

[30] A. D. Lauda, An introduction to diagrammatic algebra and categorified quantum sl_2, *Bull. Inst. Math. Acad. Sin. (N.S.)*, **7** (2012), no. 2, 165–270. Zbl 1280.81073 MR 3024893

[31] G. Lusztig, Periodic W-graphs, *Represent. Theory*, **1** (1997), 207–279 (electronic). Zbl 0895.20031 MR 1464171

[32] G. Lusztig, Singularities, character formulas, and a q-analog of weight multiplicities, *Astérisque*, **101** (1983), no. 102, 208–229. Zbl 0561.22013 MR 0737932

[33] G. Lusztig, Some examples of square integrable representations of semisimple *p*-adic groups, *Trans. Amer. Math. Soc.*, **277** (1983), no. 2, 623–653. Zbl 0526.22015 MR 0694380

[34] G. Lusztig, *Hecke algebras with unequal parameters*, CRM Monograph Series, 18, American Mathematical Society, Providence, RI, 2003. Zbl 1051.20003 MR 1974442; *Updated preprint*, 2014. arXiv:math/0208154

[35] I. Mirković and K. Vilonen, Geometric Langlands duality and representations of algebraic groups over commutative rings, *Ann. of Math. (2)*, **166** (2007), no. 1, 95–143. Zbl 1138.22013 MR 2342692

[36] S. Morrison, A formula for the Jones–Wenzl projections, *preprint*, 2002. arXiv:1503.00384

[37] S. Morrison, A diagrammatic category for the representation theory of $U_q(\mathfrak{sl}_n)$, Ph. D. Thesis, University of California, Berkeley, 2007. MR 2710589

[38] S. Riche and G. Williamson, Tilting modules and the *p*-canonical basis, *preprint*, 2016. arXiv:1512.08296

[39] W. Soergel, Kategorie $\mathcal{O}$, perverse Garben und Moduln über den Koinvarianten zur Weylgruppe, *J. Amer. Math. Soc.*, **3** (1990), no. 2, 421–445. Zbl 0747.17008 MR 1029692

[40] W. Soergel, Langlands' philosophy and Koszul duality, in *Algebra. Representation theory (Constanta, 2000)*, 379–414 NATO Sci. Ser. II Math. Phys. Chem., 28, Kluwer Acad. Publ., Dordrecht, 2001. Zbl 1001.22010 MR 1858045

[41] W. Soergel, Kazhdan–Lusztig–Polynome und unzerlegbare Bimoduln über Polynomringen, *J. Inst. Math. Jussieu*, **6** (2007), no. 3, 501–525. Zbl 1192.20004 MR 2329762

[42] H. N. V. Temperley and E. H. Lieb, Relations between the "percolation" and "colouring" problem and other graph-theoretical problems associated with regular planar lattices: some exact results for the "percolation" problem, *Proc. Roy. Soc. London Ser. A*, **322** (1971), no. 1549, 251–280. Zbl 0211.56703 MR 0498284

[43] H. Wenzl, On sequences of projections, *C. R. Math. Rep. Acad. Sci. Canada*, **9** (1987), no. 1, 5–9. Zbl 0622.47019 MR 0873400

[44] B. W. Westbury, The representation theory of the Temperley–Lieb algebras, *Math. Z.*, **219** (1995), no. 4, 539–565. Zbl 0840.16008 MR 1343661

[45] G. Williamson, Singular Soergel bimodules, *Int. Math. Res. Not. IMRN*, (2011), no. 20, 4555–4632. Zbl 1236.18009 MR 2844932

[46] S. Wu, Miniscule representations, Gauss sum and modular invariance, *preprint*, 2008. arXiv:0802.2038

B. Elias, Mathematics Department, Fenton Hall, University of Oregon, Eugene, OR 97403-1222, USA

E-mail: belias@uoregon.edu