Gelfand-Tsetlin modules in the Coulomb context

Ben Webster

Abstract. This paper gives a new perspective on the theory of principal Galois orders as developed by Futorny, Ovsienko, Hartwig and others. Every principal Galois order can be written as eFe for any idempotent e in an algebra F, which we call a flag Galois order; and in most important cases we can assume that these algebras are Morita equivalent. These algebras have the property that the completed algebra controlling the fiber over a maximal ideal has the same form as a subalgebra in a skew group ring, which gives a new perspective to a number of result about these algebras.

We also discuss how this approach relates to the study of Coulomb branches in the sense of Braverman-Finkelberg-Nakajima, which are particularly beautiful examples of principal Galois orders. These include most of the interesting examples of principal Galois orders, such as $U(\mathfrak{gl}_n)$. In this case, all the objects discussed have a geometric interpretation which endows the category of Gelfand-Tsetlin modules with a graded lift and allows us to interpret the classes of simple Gelfand-Tsetlin modules in terms of dual canonical bases for the Grothendieck group. In particular, we classify Gelfand-Tsetlin modules over $U(\mathfrak{gl}_n)$ and relate their characters to a generalization of Leclerc's shuffle expansion for dual canonical basis vectors.

1. Introduction

Let Λ be a Noetherian commutative ring, and $\widehat{W}$ a monoid acting faithfully on Λ; let $L = \mathrm{Frac}(\Lambda)$ be the fraction field of Λ. Assume that $\widehat{W}$ is the semi-direct product of a finite subgroup W and a submonoid $\mathcal{M}$ and that $\#W$ is invertible in Λ. For simplicity, we assume throughout the introduction that $\mathcal{M}$ has finite stabilizers in its action on $\mathrm{MaxSpec}(\Lambda)$.

A **principal Galois order** (Def. 2.1) is an subalgebra of invariants of the skew group ring $(L\#\mathcal{M})^G$ equipped with (amongst other structure) an inclusion of $\Gamma = \Lambda^W$ as a subalgebra (usually called the Gelfand-Tsetlin subalgebra) and a faithful action on Γ.

We call a finitely generated module **Gelfand-Tsetlin** if it is locally finite under the action of Γ, and thus decomposes as a direct sum of generalized weight spaces. An important motivating question for a great deal of work in recent years has been the question:

Question A. *Given a principal Galois order U, classify the simple Gelfand-Tsetlin modules and describe the dimensions of their generalized weight spaces for the different maximal ideals of Γ.*

Work of Drozd-Futorny-Ovsienko [DFO94, Th. 18] shows that the "fiber" over a maximal ideal m_γ of Γ is controlled by a pro-finite length algebra $\widehat{U}_\gamma$, which naturally

Date: March 27, 2019.

acts on the corresponding generalized weight space for any U-module and whose simple modules are the γ-generalized weight spaces of the different simple Gelfand-Tsetlin modules where this generalized weight space is non-zero. Thus, we can rephrase Question A as the question of how to understand these algebras in specific special cases.

One perspective shift we want to strongly emphasize is that taking invariants for a group action is a very bad idea, and that we should instead consider subalgebras F in the skew group ring of the semi-direct product $L\#(W \ltimes \mathcal{M})$, which we call **principal flag orders** (Def. 2.2). So, now the algebra Γ is replaced by the smash product $\Lambda\#W$, which in particular contains W. If we let $e \in \mathbb{Z}[\frac{1}{\#W}][W]$ be the symmetrization idempotent, then for any principal flag order F, the centralizer $U = eFe$ is a principal Galois order, and every principal Galois order appears this way.

One can easily check that Λ will be a Harish-Chandra subalgbra (in the sense of [DFO94, §1.3]) and so we can apply the results of that paper in this situation as well. Thus, for any maximal ideal $\mathfrak{m}_\lambda \subset \Lambda$, we have an algebra $\widehat{F}_\lambda$ which controls the $\mathfrak{m}_\lambda$-weight spaces for different modules. Let $\widehat{W}_\lambda \subset \widehat{W}$ be the stabilizer of $\lambda \in \operatorname{MaxSpec}(\Lambda)$ and $\widehat{\Lambda}$ the completion of Λ with respect to this maximal ideal.

Theorem B. *The algebra $\widehat{F}_\lambda$ is a principal flag order for the ring $\widehat{\Lambda}$ and the group $\widehat{W}_\lambda$, that is, it is a subalgebra of the skew group ring $\widehat{K}\#\widehat{W}_\lambda$ such that $\widehat{F}_\lambda \otimes_{\widehat{\Lambda}} \widehat{K} \cong \widehat{K}\#\widehat{W}_\lambda$, with an induced action on $\widehat{\Lambda}$.*

The difference between $\widehat{F}_\lambda$ and $\widehat{U}_\gamma$ for $\mathsf{m}_\gamma = \mathfrak{m}_\lambda \cap \Gamma$ is controlled by the stabilizer W_λ of λ in W. We have that $\widehat{U}_\gamma = e_\lambda \widehat{F}_\lambda e_\lambda$ for the symmetrizing idempotent e_λ in $\mathbb{Z}[\frac{1}{\#W}][W_\lambda]$. Thus, generically these algebras will simply be the same.

In particular, by [FO10, Th. 4.1(4)], the center of $\widehat{F}_\lambda$ is the invariants $\widehat{\Lambda}_\lambda = \widehat{\Lambda}^{\widehat{W}_\lambda}$ and any simple module over $\widehat{F}_\lambda$ will factor through the quotient $F_\lambda^{(1)}$ by the unique maximal ideal of the center. Thus, this gives a canonical way choosing a finite dimensional quotient of $\widehat{F}_\lambda$ through which all simples factor.

Note, the situation will be simpler if we work in the context of [FGRZ18], where we assume that:

($\star$) The algebra Λ is the symmetric algebra on a vector space V, the group W is a complex reflection group acting on V, $\mathcal{M}$ is a subgroup of translations, and F is free as a left Λ-module.

In this case, we can always choose F so that U and F are Morita equivalent via the bimodules eF and Fe, and the dimension of $F_\lambda^{(1)}$ is easy to calculate: it is just $(\#\widehat{W}_\lambda)^2$. Furthermore, the quotient by the maximal ideal $\mathfrak{m}_\lambda$ has dimension $\#\widehat{W}_\lambda$, and every simple module as a quotient. In particular, the sum of the dimensions of the λ-generalized weight space for all simple Gelfand-Tsetlin-modules is $\leq \#\widehat{W}_\lambda$.

If we consider how the results apply to $\widehat{U}_\gamma$, then they are almost unchanged, except that we replace the order of the group $\widehat{W}_\lambda$ with the number of cosets $S(\gamma) = \frac{\#\widehat{W}_\lambda}{\#W_\lambda}$ for any maximal ideal $\mathfrak{m}_\lambda$ lying over m_γ in Λ; this is the same statistic called $S(\mathsf{m}_\gamma, \mathsf{m}_\gamma)$ in [FO14]. With the assumptions ($\star$), the algebra $U_\gamma^{(1)}$ is $S(\gamma)^2$-dimensional, and the sum of the dimensions of the γ-generalized weight space for all simple Gelfand-Tsetlin-modules is $\leq S(\gamma)$. This seems to be implicit in the results of [FO14] (in particular, Cor. 6.1) but this perspective makes the result manifest.

1.1. **Coulomb branches.** These results however are fairly abstract and give no indication of how to actually compute the algebras $U_\gamma^{(1)}$ and understand their representation theory. However, the most interesting examples of principal Galois orders actually arise from a geometric construction: the **Coulomb branches** of Braverman, Finkelberg and Nakajima [Nak, BFNb]. These include the primary motivating example, the orthogonal Gelfand-Zetlin[1] algebras of Mazorchuk [Maz99] (including $U(\mathfrak{gl}_n)$), and a number of examples that seem to have escaped the notice of experts, such as the spherical Cherednik algebras of the groups $G(\ell, 1, n)$ and hypertoric enveloping algebras.

The Coulomb branch is an algebra constructed from the data of a gauge group G and matter representation N. For example:

- In the case where G is abelian and N arbitrary, the Coulomb branch is a hypertoric enveloping algebra as defined in [BLPW12]; the isomorphism of this with a Coulomb branch (defined at a "physical level of rigor") is proven in [BDGH, §6.6.2]; it was confirmed this matches the BFN definition of the Coulomb branch in [BFNb, §4(vii)].
- In the case where $G = GL_n$ and $N = \mathfrak{gl}_n \oplus (\mathbb{C}^n)^{\oplus \ell}$, the Coulomb branch is a spherical Cherednik algebra of the group $G(\ell, 1, n)$ by [KN]. We'll confirm in forthcoming work with LePage that the spherical Cherednik algebra for $G(\ell, p, n)$ is also a principal Galois order.
- In the case where

$$G = GL_{v_1} \times \cdots \times GL_{v_{n-1}} \tag{1.1a}$$

$$N = M_{v_n, v_{n-1}}(\mathbb{C}) \oplus M_{v_{n-1}, v_{n-2}}(\mathbb{C}) \oplus \cdots \oplus M_{v_2, v_1}(\mathbb{C}), \tag{1.1b}$$

  the Coulomb branch is an orthogonal Gelfand-Zetlin algebra associated to the dimension vector $(v_1, \ldots, v_n)$ as shown in [Wee, §3.5]. In particular, $U(\mathfrak{gl}_n)$ arises from $(1, 2, 3, \ldots, n)$.

In this case, the algebras $U_\gamma^{(1)}$ also have a geometric interpretation in terms of convolution in homology:

Theorem C. *The Coulomb branch for any group G and representation N is a principal Galois order with $\Lambda = \mathrm{Sym}^\bullet(\mathfrak{t})[h]$, the symmetric algebra on the Cartan of G with an extra loop parameter h and $\widehat{W}$ the affine Weyl group of G acting naturally on this space.*

[1]As any savvy observer knows, there is no universally agreed-upon spelling of Гельфанд-Цетлин in the Latin alphabet; in fact it's not even spelled consistently in Russian, since some authors write Цейтлин, a different transliteration of the same Yiddish name. We will write "Tsetlin" as this is the spelling that will elicit the most correct pronunciation from an English-speaker. However, since "OGZ" is well-established as an acronym, we will not change the spelling of these algebras.

For each maximal ideal m_γ of Γ, there is a Levi subgroup $G_\gamma \subset G$, with parabolic P_γ and a P_γ-submodule $N_\gamma^- \subset N$ such that

$$U_\gamma^{(1)} \cong H_*^{BM}(\{(gP_\gamma, g'P_\gamma, n) \in G_\gamma/P_\gamma \times G_\gamma/P_\gamma \times N \mid n \in gN_\gamma^- \cap g'N_\gamma^-\})$$

(1.2)

$$U_{\mathscr{S}}^{(1)} \cong \bigoplus_{\gamma,\gamma' \in \mathscr{S}} H_*^{BM}(\{(gP_\gamma, g'P_{\gamma'}, n) \in G_\gamma/P_\gamma \times G_{\gamma'}/P_{\gamma'} \times N \mid n \in gN_\gamma^- \cap g'N_{\gamma'}^-\})$$

for any set $\mathscr{S}$ contained in a single $\widehat{W}$-orbit, where the right hand side is endowed with the usual convolution multiplication (as in [CG97, (2.7.9)]*).*

This is a Steinberg algebra in the sense of Sauter [Sau]. One notable point to consider is that this algebra is naturally graded. Thus, for any choice of (G, N) and $\widehat{W}$-orbit $\mathscr{S}$, this give a graded lift $\widetilde{\Gamma\text{Ц}}(\mathscr{S})$ of the category of Gelfand-Tsetlin modules supported on this orbit. It's a consequence of the Decomposition theorem that the classes of simple modules form a dual canonical basis of the Grothendieck group $K_0(\widetilde{\Gamma\text{Ц}}(\mathscr{S}))$.

Algebras in this style have appeared numerous places in the literature. In particular, in the case of (1.1a–1.1b), the algebras that appear are already well-known: they are very closely related to the Stendhal algebras $\tilde{T}_{\mathbf{v}}\tilde{T}$ as defined in [Web17, Def. 4.5] corresponding to the Lie algebra $\mathfrak{sl}_n$, with its Dynkin diagram identified as usual with the set $\{1, \dots, n-1\}$. These algebras correspond to a list of highest weights, which we will take to be v_n copies of the $n-1$st fundamental weight ω_{n-1}; the dimension vector $(v_1, \dots, v_{n-1})$ determining the number times each Dynkin node appears as a label on a black strand. The author has proven in [Webc, Cor. 4.9] that the ring $\tilde{T}$ is an equivariant Steinberg algebra for the space appearing in (1.2). These are algebras closely related to KLR algebras [KL09], but instead of categorifying the universal enveloping algebra $U(\mathfrak{n})$ of the strictly lower triangular matrices in $\mathfrak{gl}_n$, by [Web17, Prop. 4.39], they category the tensor product of $U(\mathfrak{n})$ with the v_nth tensor power of the defining representation of $\mathfrak{gl}_n$. In particular, the classes of simples modules over this algebra match the dual canonical basis in this space (which is proven in the course of the proof of [Web15, Th. 8.7]).

The center of the algebra $\tilde{T}_{\mathbf{v}}\tilde{T}$ is a copy of

$$\Gamma = H^*(BGL_{v_1} \times \cdots \times BGL_{v_{n-1}}) \cong \bigotimes_{i=1}^{n-1} \mathbb{C}[y_{i,1}, \dots, y_{i,v_i}]^{S_{v_i}}.$$

Quotienting out by the unique graded maximal ideal in this ring gives a quotient $\tilde{T}'$; this quotient is, of course, the non-equivariant convolution algebra that appears in (1.2). That is:

Corollary D. *For $\mathscr{S}$ the set of integral elements of* $\operatorname{MaxSpec}(\Gamma)$*, the algebra $U_{\mathscr{S}}^{(1)}$ is Morita equivalent to the algebra $\tilde{T}'$.*

This gives a new way of interpreting the results of [KTW$^+$, §6]; in particular, Corollary D is effectively equivalent to Theorem 6.4 of *loc. cit.* In particular, this gives us a criterion in terms of which weight spaces are not zero that classifies the different simple

Gelfand-Tsetlin modules with integral weights for an orthogonal Gelfand-Zetlin algebra (Theorem 5.9).

Acknowledgements

A great number of people deserve acknowledgement in the creation of this paper: my collaborators Joel Kamnitzer, Alex Weekes and Oded Yacobi, since this project grew out of our joint work; Walter Mazorchuk, who first suggested to me that a connection existed between our previous work and the study of Gelfand-Tsetlin modules, and who, along with Elizaveta Vishnyakova, Slava Futorny, Dimitar Grantcharov and Pablo Zadunaisky, suggested many references and improvements; and Hiraku Nakajima who, amongst many other things, pointed out to me the argument used in the proof of Theorem 4.4.

The author was supported by NSERC through a Discovery Grant. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Research, Innovation and Science.

2. Generalities on Galois orders

Following the notation of [Har], let Λ be a noetherian integrally closed domain, L its fraction field. Note that this implies Hartwig's condition (A3), and we lose no generality in assuming this by [Har, Lem. 2.1]. Let W be a finite group[2] acting faithfully on Λ and $\Gamma = \Lambda^W, K = L^W$. Let $\mathcal{M}$ be a submonoid of $\mathrm{Aut}(\Lambda)$ which is normalized by W, and let $\widehat{W} = \mathcal{M} \ltimes W$, which we also assume acts faithfully (this implies Hartwig's (A1) and (A2)). Let $\mathcal{L}$ be the smash product $L\#\mathcal{M}$, $\mathcal{F} = \mathcal{L}\#W$, and $\mathcal{K} = \mathcal{L}^W$. Note that L is a $\mathcal{L}$ module in the obvious way, and thus K is a $\mathcal{K}$-module.

The more general notion of Galois orders was introduced by Futorny and Ovsienko [FO10], but we will only be interested in a special class of these considered in Hartwig in [Har], which makes these properties easy to check.

Definition 2.1 ([Har, Def 2.22 & 2.24])**.** *The* **standard order** *is the subalgebra*

$$\mathcal{K}_\Gamma = \{X \in \mathcal{K} \mid X(\Gamma) = \Gamma\}.$$

A subalgebra $A \subset \mathcal{K}_\Gamma$ is a **principal Galois order** *if $KA = \mathcal{K}$.*

It is a well-known principle in the analysis of quotient singularities that taking the smash product of an algebra with group acting on it is a much better behaved operation than taking invariants. Similarly in the world of Galois orders, there is a larger algebra that considerably simplifies the analysis of these algebras.

Definition 2.2. *The* **standard flag order** *is the subalgebra*

$$\mathcal{F}_\Lambda = \{X \in \mathcal{F} \mid X(\Lambda) = \Lambda\}.$$

A subalgebra $F \subset \mathcal{F}_\Lambda$ is called a **principal flag order** *if $KF = \mathcal{F}$ and $W \subset F$.*

[2]Note that this is a departure from [Har], where this group is denoted G. We will be most interested in the case where W is the Weyl group of a semi-simple Lie group acting on the Cartan, so we prefer to save G for the name of this group.

It's an easy check, via the same proofs, that the analogues of [Har, Prop. 2.5, 2.14 & Thm 2.21] hold here: that is F is a Galois order inside $\mathcal{F}$ with Λ maximal commutative; in order to match the notation of [FO10], we must take $G = \{1\}$ and $\mathcal{M} = W \ltimes \mathcal{M}$.

Let $e = \frac{1}{\#W} \sum_{w \in W} w \in \mathcal{F}_\Lambda$. Note that $\mathcal{K} \subset \mathcal{F}$ via the obvious inclusion, and that given $k \in \mathcal{K}$, the element $eke \in \mathcal{F}$ acts on Γ by the same operator as k. Thus, $k \mapsto eke$ is an algebra isomorphism $\mathcal{K} \cong e\mathcal{F}e$.

Lemma 2.3. *The isomorphism above induces an isomorphism $\mathcal{K}_\Gamma \cong e\mathcal{F}_\Lambda e$.*

Proof. If $a \in \mathcal{F}_\Lambda$, then $eae\Gamma = ea\Gamma \subset e\Lambda = \Gamma$, so $eae \in e\mathcal{K}_\Gamma e$. On the other hand, $e\mathcal{K}_\Gamma e$ acts trivially on the elements of Λ that transform by any non-trivial irrep, and sends Λ to Λ, so indeed, this lies in $e\mathcal{F}_\Lambda e$. □

Thus, we have that for any flag order F, the centralizer algebra $U = eFe$ is a principal Galois order. As usual with the centralizer algebra of an idempotent:

Lemma 2.4. *The category of U-modules is a quotient of the category of F-module via the functor $M \mapsto eM$; that is, this functor is exact and has right and left adjoints $N \mapsto Fe \otimes_U N$ and $N \mapsto \mathrm{Hom}_U(eF, N)$ that split the quotient functor.*

Furthermore, every principal Galois order appears this way. Consider the smash product $\Lambda \# W \subset \mathrm{End}_{\Lambda^W}(\Lambda)$, and let D be a subalgebra satisfying $\Lambda\#W \subset D \subset \mathrm{End}_\Gamma(\Lambda) \subset L\#W$. Note that in this case, $eDe = \Gamma$, since this is true when $D = \Lambda\#W$ or $D = \mathrm{End}_\Gamma(\Lambda)$. Let $F_D = De \otimes_\Gamma U \otimes_\Gamma eD$ endowed with the obvious product structure (using the map $eD \otimes_D De \to \Gamma$).

Lemma 2.5. *For any principal Galois order U, and any D as above, the obvious algebra map $F_D \to \mathcal{F}_\Lambda$ makes F_D into a principal Galois order such that $U = eF_De$.*

Proof. First, note that since $D \subset L\#W$, can identify eD and De with Λ-submodules of $L \cong e(L\#W) = (L\#W)e$. Since the natural map $(L\#W)e \otimes_K \mathcal{K} \otimes_K e(L\#W) \to \mathcal{F}$ is an isomorphism, this shows that F_D injects into $\mathcal{F}$, and this is clearly an algebra map. Thus, we will use the same symbol to denote the image.

First, note that F_D is a principal flag order, since $KF_D \supset \mathcal{K}\Lambda W = \mathcal{L}W = \mathcal{F}$ and by assumption F_D contains the smash product $\Lambda\#W$. Furthermore,

$$eF_De = eDe \otimes_\Gamma U \otimes_\Gamma eDe = \Gamma \otimes_\Gamma U \otimes_\Gamma \Gamma = U$$

so we have all the desired properties. □

2.1. **Gelfand-Tsetlin modules.** Now, fix a flag Galois order $F \subset \mathcal{F}_\Lambda$. We wish to understand the representation theory of this algebra. Consider the weight functors

$$\mathcal{W}_\lambda(M) = \{m \in M \mid \mathfrak{m}_\lambda^N m = 0 \text{ for some } N \gg 0\}$$

for $\lambda \in \mathrm{MaxSpec}(\Lambda)$. The reader might reasonably be concerned about the fact that this is a *generalized* eigenspace; in this paper, we will always want to consider these, and thus will omit "generalized" before instances of "weight."

Definition 2.6. *We call a finitely generated F-module M a* **weight module** *or* **Gelfand-Tsetlin module** *if $M = \bigoplus_{\lambda \in \mathrm{MaxSpec}(\Lambda)} \mathcal{W}_\lambda(M)$.*

Remark 2.7. *One subtlety here is that we have not assumed that $\mathcal{W}_\lambda(M)$ is finite dimensional. We'll see below that this holds automatically if the stabilizer of λ in $\widehat{W}$ is finite.*

Since many readers will be more interested in the Galois order $U = eFe$, let us compare the weight spaces of a module M with those of the U-module eM. Of course, in U, we only have an action of Γ. Let $\gamma \in \mathrm{MaxSpec}(\Gamma)$ be the image of λ under the obvious map and

$$\mathcal{W}_\gamma(N) = \{m \in eM \mid \mathsf{m}_\gamma^N m = 0 \,\forall N \gg 0\}.$$

Lemma 2.8. *If M is a Gelfand-Tsetlin F-module, then eM is a Gelfand-Tsetlin U-module with*

$$\mathcal{W}_\gamma(eM) \cong e_\lambda \mathcal{W}_\lambda(M).$$

Proof. Let $\mathsf{m}_\gamma = \Gamma \cap \mathfrak{m}_\lambda$; by standard commutative algebra, the other maximal ideals lying over m_γ are those in the orbit $W \cdot \lambda$. Thus, we have that

$$\mathcal{W}_\gamma(eM) = e \cdot \left(\bigoplus_{\lambda' \in W\lambda} Wei_{\lambda'}(M) \right).$$

This space $\bigoplus_{\lambda' \in W\lambda} Wei_{\lambda'}(M)$ has a W-action induced by the inclusion $W \subset F$, and is isomorphic to the induced representation $\mathrm{Ind}_{W_\lambda}^W \mathcal{W}_\lambda(M)$ since it is a sum of subspaces which it permutes like the cosets of this subgroup. Thus, its invariants are canonically isomorphic to the invariants for W_λ on $\mathcal{W}_\lambda(M)$. □

2.2. **The fiber for a flag order.**

Definition 2.9. *Fix an integer N. The* **universal Gelfand-Tsetlin module** *of weight λ and length N is the quotient $F/F\mathfrak{m}_\lambda^N$.*

This is indeed a Gelfand-Tsetlin-module by [Har, Lem. 3.2]. Obviously, this represents the functor of taking generalized weight vectors killed by ${\mathfrak{m}_\lambda}^N$:

$$\mathrm{Hom}_F(F/F{\mathfrak{m}_\lambda}^N, M) = \{m \in M \mid {\mathfrak{m}_\lambda}^N m = 0\}.$$

In particular, every simple Gelfand-Tsetlin-module with $\mathcal{W}_\lambda(S) \neq 0$ is a quotient of $F/F\mathfrak{m}_\lambda$, since it must have a vector killed by $\mathfrak{m}_\lambda$. Taking inverse limit $\varprojlim F/F\mathfrak{m}_\lambda^N$, we obtain an universal (topological) Gelfand-Tsetlin module of arbitrary length. Consider the algebra

$$\widehat{F}_\lambda = \varprojlim F/\big(F\mathfrak{m}_\lambda^N + \mathfrak{m}_\lambda^N F\big)$$

As noted in [DFO94, Th. 18], this algebra controls the λ weight spaces of all modules, and in particular simple modules.

Let $\widehat{W}_\lambda$ be the subgroup of $\widehat{W} = W \ltimes \mathcal{M}$ which fixes λ. For the remainder of this section, we assume that $\widehat{W}_\lambda$ is finite. This implies that Λ is finitely generated over $\Lambda_\lambda = \Lambda^{\widehat{W}_\lambda}$.

Definition 2.10. *Let F_λ be the intersection $F \cap K \cdot \widehat{W}_\lambda \subset \mathcal{K} = K\widehat{W}$ with the K-span of $\widehat{W}_\lambda$. Since F_λ is the intersection of two subalgebras, it is itself a subalgebra.*

This has an obvious left and right module structure over Λ but Λ is not central.

Lemma 2.11. *The image of F_λ spans $F/\big(F\mathfrak{m}_\lambda^N + \mathfrak{m}_\lambda^N F\big)$ for all N.*

Proof. This is essentially just a restatement of the proof of [FO14, Lemma 5.3]. The quotient $F/\big(F\mathfrak{m}_\lambda^N + \mathfrak{m}_\lambda^N F\big)$ is finitely generated as a Λ-Λ-bimodule, and thus generated by the images of finitely many elements $f_1, \dots, f_n$ of F. Thus, there is some finite set T given by the union of the supports in $\widehat{W}$ of these elements. We induct on the number of elements of T that don't lie in $\widehat{W}_\lambda$.

If t is such an element, then there is some polynomial $p \in \mathfrak{m}_\lambda^N$ which does not vanish at $p(t^{-1}\cdot\lambda)$ for any $t \in T$; that is, p^t is a unit mod $\mathfrak{m}_\lambda^N$. Thus, $p^t \otimes 1 - 1 \otimes p$ acts invertibly on the quotient $F/\big(F\mathfrak{m}_\lambda^N + \mathfrak{m}_\lambda^N F\big)$, so the elements $p^t f_k - f_k p$ are still generators, but their support now lies in $T \setminus \{t\}$ by [FO10, Lem. 5.2]. Applied inductively, this achieves the result. □

Lemma 2.12. *The ring F_λ is finitely generated as a left module and as a right module over Λ and satisfies $F_\lambda\Lambda = \Lambda F_\lambda = K \cdot \widehat{W}_\lambda$. In fact, F_λ is a Galois order for the group $\mathcal{M} = \widehat{W}_\lambda$ and commutative ring Λ, using the notation of* [FO10].

This shows in particular that Λ is **big at** λ in the terminology of [DFO94].

Proof. The finite generation is an immediate consequence of the fact that F is an order. Similarly, that F_λ has the order property, i.e. its intersection with any finite dimensional K-subspace for the left/right action of $K\cdot\widehat{W}_\lambda$ is finitely generated for the left/right action of Λ is an immediate consequence of the same property for F.

Thus, it only remains to show that $F_\lambda\Lambda = \Lambda F_\lambda = K \cdot \widehat{W}_\lambda$. Since $\mathcal{F} = \Lambda F$, for any $w \in \widehat{W}_\lambda$, we have $w = \sum k_i f_i$ for $k_i \in K$, and $f_i \in F$. As in the proof of 2.11 above, we can assume that the f_i's have support in some set T, and if $t \in T$ but not in $\widehat{W}_\lambda$, then we have a polynomial p as before, vanishing at λ, but not at $t^{-1}\cdot\lambda$. Note that we have $w = \frac{1}{p^{wt}-p^w}(p^{wt}w - wp)$, with the $p^t - p$ being non-zero in K since it does not vanish at λ. Substituting in our formula for w, we have

$$w = \frac{k_i}{p^{wt} - p^w}(p^{wt} f_i - f_i p)$$

Thus, we can inductively reduce the size of T until $T \subset \widehat{W}_\lambda$, so we can assume that $f_i \in F_\lambda$. □

This shows that $\widehat{F}_\lambda$ is the completion of F_λ with respect to the topology induced by the basis of neighborhoods of the identity $F_\lambda\mathfrak{m}_\lambda^N + \mathfrak{m}_\lambda^N F_\lambda$. Alternatively, we can think about this topology by noting that F_λ is finitely generated over $\Lambda_\lambda = \Lambda^{\widehat{W}_\lambda}$. Furthermore, Λ_λ is central in F_λ, since it commutes with $K \cdot \widehat{W}_\lambda$; in fact, by Lemma 2.12 above and [FO10, Th. 4.1(4)], it is the full center of this algebra. Let $\mathfrak{n}_\lambda = \mathfrak{m}_\lambda \cap \Lambda_\lambda$. Since λ is fixed by $\widehat{W}_\lambda$ (by definition), the ideal $\mathfrak{n}_\lambda\Lambda$ still only vanishes at λ, that is, $\mathfrak{n}_\lambda\Lambda \supset \mathfrak{m}_\lambda^k$ for some k.

Thus, if we let $\widehat{\Lambda}$ and $\widehat{\Lambda}_\lambda$ be the completion with of the respective rings in the $\mathfrak{m}_\lambda$-adic and $\mathfrak{n}_\lambda$-adic topologies, then:

Lemma 2.13. *We have an isomorphism of topological rings*

$$\widehat{F}_\lambda = F_\lambda \otimes_{\Lambda_\lambda} \widehat{\Lambda}_\lambda$$

and the ring $\widehat{F}_\lambda$ is a Galois order for $\mathcal{M} = \widehat{W}_\lambda$ and the ring $\widehat{\Lambda}$.

Proof. The tensor product $F_\lambda \otimes_{\Lambda_\lambda} \widehat{\Lambda}_\lambda$ is the completion of F_λ with respect to the topology with basis of 0 given by the 2-sided ideals $F_\lambda \mathfrak{n}_\lambda^N$. Since $\Lambda \mathfrak{n}_\lambda \supset \mathfrak{m}_\lambda^k$ for some k, we have that

$$F_\lambda \mathfrak{m}_\lambda^{kN} + \mathfrak{m}_\lambda^{kN} F_\lambda \subset F_\lambda \mathfrak{n}_\lambda^N \subset F_\lambda \mathfrak{m}_\lambda^N + \mathfrak{m}_\lambda^N F_\lambda$$

which shows the equivalence of the topologies, and thus the isomorphism of completions. Faithful base changes by a central subalgebra obviously preserves the properties of being a Galois order, so this follows from Lemma 2.12. □

We can use these result to also understand the fiber for U as well for any principal Galois order. By Lemma 2.5, we can choose a principal flag order with $U = eFe$. The algebra F_λ contains the stabilizer W_λ and its symmetrizing idempotent e_λ. As before, let γ be the image of λ in MaxSpec(Γ). Let $U_\gamma = e_\lambda F_\lambda e_\lambda$, and $\widehat{U}_\gamma$ the completion $\varprojlim U/\big(U\mathsf{m}_\gamma^N + \mathsf{m}_\gamma^N U\big)$.

Lemma 2.14. *The algebra U_γ surjects onto $U/\big(U\mathsf{m}_\gamma^N + \mathsf{m}_\gamma^N U\big)$ for any N, and thus has dense image in $\widehat{U}_\gamma \cong e_\lambda \widehat{F}_\lambda e_\lambda$.*

This is sufficiently similar to Lemma 2.11 and [FO14, Lemma 5.3] that we leave it as an exercise to the reader.

2.3. **Universal modules.** While this is largely redundant with [DFO94], it will be helpful to explain how we construct simple Gelfand-Tsetlin modules

Definition 2.15. *Fix an integer N. The* **central universal Gelfand-Tsetlin module** *of weight λ and length N is the quotient $P_\lambda^{(N)} = F/F\mathfrak{n}_\lambda^N$.*

Consider the quotient algebra $F_\lambda^{(N)} := F_\lambda / F_\lambda \mathfrak{n}_\lambda^N$.

Theorem 2.16. *The module $P_\lambda^{(N)}$ is a Gelfand-Tsetlin module such that*

$$\mathcal{W}_\lambda(P_\lambda^{(N)}) \cong \operatorname{End}(P_\lambda^{(N)}) \cong F_\lambda^{(N)}.$$

More generally, we have that

$$\operatorname{Hom}_F(P_\lambda^{(N)}, M) = \{m \in M \mid \mathfrak{n}_\lambda^N m = 0\}. \tag{2.1}$$

Note that "length N" refers to the maximal length of a Jordan block of an element of $\mathfrak{n}_\lambda$, not of $\mathfrak{m}_\lambda$. Since $\mathfrak{n}_\lambda$ is central in F_λ, the ideal $\mathfrak{n}_\lambda^N$ acts trivially on $P_\lambda^{(N)}$. On the other hand the nilpotent length of the action of $\mathfrak{m}_\lambda$ on $F_\lambda/F_\lambda {\mathfrak{m}_\lambda}^N$ is typically more than N.

Proof. Equation (2.1) is a basic property of left ideals. This is a Gelfand-Tsetlin module by [Har, Lem. 3.2].

Note that the map $F_\lambda \to \mathcal{W}_\lambda(P_\lambda^{(N)})$ is surjective by 2.11. Of course, the kernel of this map is $F_\lambda \cap F\mathfrak{n}_\lambda^N = F_\lambda \mathfrak{n}_\lambda^N$. This shows that $\mathcal{W}_\lambda(P_\lambda^{(N)}) \cong F_\lambda / F_\lambda \mathfrak{n}_\lambda^N$. Since $\mathfrak{n}_\lambda^N$ is central in F_λ, it acts trivially on this weight space, and the identification with $\operatorname{End}(P_\lambda^{(N)})$ follows from (2.1). □

It follows immediately from [DFO94, Th. 18] that:

Theorem 2.17. *The map sending $S \mapsto \mathcal{W}_\lambda(S)$ is a bijection between the isoclasses of simple Gelfand-Tsetlin F-modules in the fiber over λ and simple $F_\lambda^{(1)}$-modules.*

Since $F_\lambda \otimes_\Gamma K$ is $(\#\widehat{W}_\lambda)^2$ dimensional over K, we have that $F_\lambda^{(1)}$ is at least length $(\#\widehat{W}_\lambda)^2$ over Λ. and $U_\gamma^{(N)} = e_\lambda F_\lambda{}^{(N)} e_\lambda$

Similarly, we can define a U-module $Q_\gamma^{(N)} = eP_\lambda^{(N)} e_\lambda = e(P_\lambda^{(N)})^{W_\lambda}$ such that

$$\mathcal{W}_\gamma(Q_\gamma^{(N)}) \cong \operatorname{End}(Q_\gamma^{(N)}) \cong U_\lambda^{(N)} = e_\lambda F_\lambda^{(N)} e_\lambda.$$

More generally, we have that

$$\operatorname{Hom}_F(Q_\lambda^{(N)}, M) = \{n \in N \mid (\Lambda \mathfrak{n}_\lambda^N \cap \Gamma)m = 0\}. \tag{2.2}$$

Applying [DFO94, Th. 18] again shows that the map sending $S \mapsto \mathcal{W}_\gamma(S)$ is a bijection between the isoclasses of simple Gelfand-Tsetlin U-modules in the fiber over γ and simple $U_\gamma^{(1)}$-modules.

2.4. Weightification and canonical modules.

There is another natural way to try to construct Gelfand-Tsetlin modules. Consider any F-module M, and fix an $\widehat{W}$-invariant subset $\mathscr{S} \subset \operatorname{MaxSpec}(\Lambda)$.

Definition 2.18. *Consider the sums*

$$M^{\mathscr{S}} = \bigoplus_{\lambda \in \mathscr{S}} \{m \in M \mid \mathfrak{n}_\lambda m = 0\} \qquad M_{\mathscr{S}} = \bigoplus_{\lambda \in \mathscr{S}} M/\mathfrak{n}_\lambda M$$

Theorem 2.19. *The action of F on M induces a Gelfand-Tsetlin F-module structure on $M^{\mathscr{S}}$ and $M_{\mathscr{S}}$.*

Note that even if M is a finitely generated module, the modules $M^{\mathscr{S}}$ and $M_{\mathscr{S}}$ may not be finitely generated, though the individual weight spaces

$$\mathcal{W}_\lambda(M^{\mathscr{S}}) = \{m \in M \mid \mathfrak{n}_\lambda m = 0\} \qquad \mathcal{W}_\lambda(M_{\mathscr{S}}) = M/\mathfrak{n}_\lambda M$$

will be finitely generated over $\Lambda_\lambda^{(1)} = \Lambda/\Lambda\mathfrak{n}_\lambda$.

Proof. Consider any element $f \in F$. By the Harish-Chandra property, $\Lambda f \Lambda$ is finitely generated as a right Λ-module, so $\Lambda f \Lambda \otimes_\Lambda \Lambda_\lambda^{(1)}$ is a finite length left Λ-module. Thus, we can assume without loss of generality that the image of f in the quotient is a generalized weight vector of weight μ.

Let ${}_\mu W_\lambda$ be the set of elements of $\widehat{W}$ such that $w \cdot \lambda = \mu$. Let ${}_\mu F_\lambda = F \cap K \cdot {}_\mu W_\lambda$ be the elements of F which are in the K-span of ${}_\mu W_\lambda$. Thus, we can reduce to the case where $f \in {}_\mu F_\lambda$. Every element of ${}_\mu W_\lambda$ induces the same isomorphism $\sigma \colon \Lambda_\lambda \to \Lambda_\mu$ such that $\sigma(\mathfrak{n}_\lambda) = \mathfrak{n}_\mu$, so we have that for any $a \in \mathfrak{n}_\lambda$, then $af = f\sigma^{-1}(a)$.

Thus, if $\mathfrak{n}_\lambda m = 0$, we have that $\mathfrak{n}_\mu f m = 0$, so $f \mathfrak{n}_\lambda m = 0$ and $fm \in \mathcal{W}_\lambda(M^{\mathscr{S}})$. This shows that we have an induced action. Similarly, given $m \in M/\mathfrak{n}_\lambda M$, the image fm is thus a well-defined element of $M/\mathfrak{n}_\mu M$. This completes the proof. $\square$

We could similarly consider "thicker" versions of these modules where we replace $\mathfrak{n}_\lambda$ with powers of this ideal, and direct/inverse limits of the resulting modules. Since we have no application in mind for these modules, we will leave discussion of them to another time.

One particularly interesting module to apply this result to is Λ itself. In this case, $\Lambda_{\mathscr{S}}$ is a Gelfand-Tsetlin module such that $\mathcal{W}_\lambda(\Lambda_{\mathscr{S}}) = \Lambda_\lambda^{(1)}$ for all $\lambda \in \mathscr{S}$. This same module has been constructed by Mazorchuk and Vishnyakova [MV, Th. 4]. The dual version of this construction given by taking the vector space dual $\Lambda^* = \mathrm{Hom}_{\Bbbk}(\Lambda, \Bbbk)$ for some subfield $\Bbbk$ and considering $(\Lambda^*)^{\mathscr{S}}$ has been studied by several authors, including Early-Mazorchuk-Vishnyakova [EMV], Hartwig [Har] and Futorny-Grantcharov-Ramirez-Zadunaisky [FGRZ18]; in particular, it appears to the author that $e(\Lambda^*)^{\mathscr{S}}$ is precisely the $U = eFe$ module $V(\Omega, T(v))$ defined in [FGRZ18, Def. 7.3] when $\mathscr{S} = \widehat{W} \cdot v$ and Ω is a base of the group $\widehat{W}_\lambda$ for any $\lambda \in \mathscr{S}$.

Based on the structure of this module, we can construct a "canonical" module as in [EMV, Har]; the author is not especially fond of this name as the embedding of F in $\mathcal{F}$ is not itself canonical, if the algebra F is the object of interest. For every $\lambda \in \mathscr{S}$, we can consider the submodule C'_λ of $\Lambda_{\mathscr{S}}$ generated by $\mathcal{W}_\lambda(\Lambda_{\mathscr{S}})$ which is clearly finitely (in fact, cyclically) generated.

Lemma 2.20. *The submodule C'_λ has a unique simple quotient C_λ, and corresponds to the unique simple quotient of $\Lambda_\lambda^{(1)}$ as a $F_\lambda^{(1)}$-module under Theorem 2.17.*

Proof. Given any proper submodule $M \subset C'_\lambda$, consider $M \cap \mathcal{W}_\lambda(\Lambda_{\mathscr{S}}) \subset \Lambda_\lambda^{(1)}$. This must be a proper submodule, because $\mathcal{W}_\lambda(\Lambda_{\mathscr{S}})$ generates. As a $\Lambda_\lambda^{(1)}$-module, $\Lambda_\lambda^{(1)}$ has a unique maximal submodule, the ideal $\mathfrak{m}_\lambda/\mathfrak{n}_\lambda$, which thus contains $M \cap \mathcal{W}_\lambda(\Lambda_{\mathscr{S}})$. Thus, the sum of two proper submodules has the same property and thus is again proper. This shows there is a unique maximal proper submodule, and thus a unique simple quotient. □

In the terminology of [Har], the canonical module is actually the right module C^*_λ obtained by dualizing this construction with respect to a subfield $\Bbbk$. Note that since we avoid dualizing, our result here is both a bit stronger and a bit weaker than [Har, Thm. 3.3]. That result does not depend on the finiteness of $\widehat{W}_\lambda$, though as a result, one pays the price of not knowing whether $\mathcal{W}_\lambda$ is finite dimensional. However, our construction applies when Λ is arbitrary, making no assumption on characteristic or linearity over a field.

2.5. **Interaction between weight spaces.** In this section, we continue to assume that every weight considered has finite stabilizer in $\widehat{W}$. Of course, we are also interested in the overall classification of modules. Consider two different weights λ and μ.

Let ${}_\lambda W_\mu$ be the set of elements of $\widehat{W}$ such that $w \cdot \mu = \lambda$. Let ${}_\lambda F_\mu = F \cap K \cdot {}_\lambda W_\mu$ be the elements of F which are in the K-span of ${}_\lambda W_\mu$. This is clearly a F_λ - F_μ-bimodule, and we have a multiplication ${}_\lambda F_\mu \otimes_{F_\mu} {}_\mu F_\nu \to {}_\lambda F_\nu$. Thus, we can define a matrix algebra:

$$F(\lambda_1, \dots, \lambda_k) = \begin{bmatrix} F_{\lambda_1} & {}_{\lambda_1}F_{\lambda_2} & \cdots & {}_{\lambda_1}F_{\lambda_k} \\ {}_{\lambda_2}F_{\lambda_1} & F_{\lambda_2} & \cdots & {}_{\lambda_2}F_{\lambda_k} \\ \vdots & \vdots & \ddots & \vdots \\ {}_{\lambda_k}F_{\lambda_1} & {}_{\lambda_k}F_{\lambda_2} & \cdots & F_{\lambda_k} \end{bmatrix} \tag{2.3}$$

More generally, for any subset $\mathsf{S} \subset \operatorname{MaxSpec}(\Lambda)$, we let $F(\mathscr{S})$ be the direct limit of this matrix algebra over all finite subsets. Note that if $\mathscr{S}$ is not finite, this is not a unital algebra, but is locally unital. This acts by natural transformations on the functor $\bigoplus_{\lambda\in\mathscr{S}} \mathcal{W}_\lambda$.

Note that if λ and μ are not in the same orbit of $\widehat{W}$, then ${}_\lambda F_\mu = 0$, so $F(\mathsf{S})$ naturally breaks up as a direct sum over the different $\widehat{W}$ orbits these weights lie in.

If λ and μ are in the same orbit, then we have a canonical isomorphism $\Lambda_\lambda \cong \Lambda_\mu$ induced by any element of ${}_\lambda W_\mu$, which identifies the ideals $\mathfrak{n}_\lambda$ and $\mathfrak{n}_\mu$. Thus for $\mathscr{S}$ a $\widehat{W}$-orbit, we can identify these with a single algebra $\Lambda_{\mathscr{S}} \supset \mathfrak{n}$.

Proposition 2.21. *If $\mathsf{S} \subset \mathscr{S}$, then $\Lambda_{\mathscr{S}}$ is the center of $F(\mathsf{S})$.*

Proof. As discussed before $F_\lambda \otimes_\Gamma K \cong \widehat{W}_\lambda \ltimes L$, and ${}_{\lambda_1}F_{\lambda_2} \otimes_\Gamma K$ is just the bimodule induced by an isomorphism between these algebras. Thus $F(\mathsf{S}) \otimes K$ is Morita equivalent to $\widehat{W}_\lambda \ltimes L$, and its center is the subfield $L^{\widehat{W}_\lambda} \subset L$. We have that $Z(F(\lambda_1, \dots, \lambda_k)) = F(\mathsf{S}) \cap Z(\widehat{W}_\lambda \ltimes L) = \Lambda_{\mathscr{S}}$. □

Let

$$F^{(N)}(\mathsf{S}) = F(\mathsf{S})/\mathfrak{n}^N F(\mathsf{S})$$
$$\widehat{F}(\mathsf{S}) = F(\mathsf{S}) \otimes_{\Lambda_{\mathscr{S}}} \widehat{\Lambda}_{\mathscr{S}}.$$

As a consequence of [DFO94, Th. 17], we can easily extend Theorem 2.17 to incorporate any number of weight spaces.

Theorem 2.22. *The simple Gelfand-Tsetlin F-modules S such that $\mathcal{W}_\lambda(S) \neq 0$ for some $\lambda \in \mathsf{S}$ are in bijection with simple modules over $F^{(1)}(\mathsf{S})$, sending $S \mapsto \bigoplus_{\lambda\in\mathscr{S}} \mathcal{W}_\lambda(S)$.*

We can also extend this to an equivalence of categories: let $\Gamma Ц(\mathsf{S})$ be the category of all Gelfand-Tsetlin modules modulo the subcategory of modules such that $\mathcal{W}_{\lambda_i}(M) = 0$ for all i, and $\Gamma Ц(\mathscr{S})$ the category of Gelfand-Tsetlin modules where if $\lambda \notin \mathscr{S}$, we have $\mathcal{W}_\lambda(M) = 0$.

For any finite set S, we have that:

Theorem 2.23. *The functor $S \mapsto \oplus_{i=1}^k \mathcal{W}_{\lambda_i}(S)$ gives an equivalence between $\Gamma Ц(\mathsf{S})$ and finite dimensional modules over the completion $\widehat{F}(\mathsf{S})$.*

As before, let $\mathscr{S}$ be a $\widehat{W}$-orbit in $\operatorname{MaxSpec}(\Lambda)$.

Definition 2.24. *We call a set of weights $\mathsf{S} \subset \mathscr{S}$* **complete** *for the orbit $\mathscr{S}$ if $\Gamma Ц(\mathsf{S}) = \Gamma Ц(\mathscr{S})$, that is, if any module M with $\mathcal{W}_{\lambda_i}(M) = 0$ for all i satisfies $\mathcal{W}_\lambda(M) = 0$ for all $\lambda \in \mathscr{S}$.*

Note that if S is a finite complete set for the orbit $\mathscr{S}$, then $\Gamma Ц(\mathscr{S}) \cong \widehat{F}(\mathsf{S})$-fdmod.

Of course, many readers will be more interested in understanding modules of the original principal Galois order. For simplicity, assume that S only contains at most one element of each W-orbit. We can derive weight spaces of U from those of F by taking invariants under the stabilizer $\widehat{W}_\lambda$. Let e_λ be the idempotent in $\widehat{F}_\lambda$ which projects to the invariants of W_λ, and $e_{\boldsymbol{\lambda}} \in \widehat{F}(\mathsf{S})$ the matrix with these as diagonal entries for the different $\lambda \in \mathsf{S}$. Let $U^{(1)}(\mathsf{S}) = e_{\boldsymbol{\lambda}} F^{(1)}(\mathsf{S}) e_{\boldsymbol{\lambda}}$.

Theorem 2.25. *The simple Gelfand-Tsetlin U-modules S such that $\mathcal{W}_\gamma(S) \neq 0$ for γ in the image of S are in bijection with simple modules over $U^{(1)}(\mathsf{S})$, sending $S \mapsto \oplus_{\lambda \in \mathsf{S}} e_\lambda \mathcal{W}_\lambda(S)$.*

3. The reflection case

While we worked in Section 2 in the same generality as [Har] so the results we can prove in this generality are available there, we wish to specialize to a much simpler case. Let V be a $\mathbb{C}$-vector space with an action of a complex reflection group W, and $\mathcal{M}$ a finitely generated (over $\mathbb{Z}$) subgroup of V^*. We assume from now on that $\Lambda = \mathrm{Sym}^\bullet(V)$ is the symmetric algebra on this vector space, with the obvious induced $\mathcal{M}$-action. Note that the stabilizer $\widehat{W}_\lambda$ for any $\lambda \in V^*$ is finite, and in fact a subgroup of W via the usual quotient map $\widehat{W} \to W$. It is generated by the $\mathcal{M}$-translates of root hyperplanes containing λ, and thus is again a complex reflection group, acting by the translation of a linear action.

This simplifies matters in one key way: the module Λ is a free Frobenius extension over Λ_λ and over Γ. Recall that we call a ring extension $A \subset B$ **free Frobenius** if B is a free A-module, and $\mathrm{Hom}_A(B, A)$ is a free B module of rank 1; a **Frobenius trace** is a generator of $\mathrm{Hom}_A(B, A)$.

The fact that Λ is free Frobenius over Γ is well-known, and easily derived from results in [Bro10]: following the notation of *loc. cit.*, we have a map $\Lambda \to \Gamma$ defined by $D(J^*)$, which is the desired trace. In slightly more down to earth terms, we have a unique element $J \in \Lambda$ of minimal degree that transforms under the determinant character for the action on V^*; this obtained by taking a suitable power of the linear form defining each root hyperplane. The Frobenius trace is uniquely characterized by sending this element to $1 \in \Gamma$ and killing all other isotypic components for the action of W.

In particular, this means that $D = \mathrm{End}_\Gamma(\Lambda)$, the **nilHecke algebra** of W, is Morita equivalent to Γ; see for example [Gin18, Lemma 7.1.5].

Definition 3.1. *We call a flag order F* **Morita** *if the symmetrization idempotent gives a Morita equivalence between $U = eFe$ and F; that is if $F = FeF$.*

Recall that for a fixed principal Galois order U, we have an associated flag Galois order F_D. Since $D = DeD$ when $D = \mathrm{End}_\Gamma(\Lambda)$ in the complex reflection case, we have that the flag order F_D is Morita for any principal Galois order in this case.

Thus, for any principal Galois order, we can study the representation theory of the corresponding flag order instead. This approach is implicit in much recent work in the subject, which uses the nilHecke algebra, such as [FGRZ, FGR16, RZ18], but many issues are considerably simplified if we think of the flag order as the basic object.

It's easy to see how Gelfand-Tsetlin modules behave under this equivalence. We can strengthen Lemma 2.8 to:

Lemma 3.2. *If F is Morita, then we have isomorphisms*

$$\mathcal{W}_\gamma(eM) \cong \mathcal{W}_\lambda(M)^{W_\lambda} \qquad \mathcal{W}_\lambda(M) \cong (\mathcal{W}_\gamma(eM))^{\oplus \# W_\lambda}.$$

The additional information we learn from the fact that F is Morita is that $\mathcal{W}_\lambda(M)$ is free as a $\mathbb{C}W_\lambda$-module.

Note that $\Lambda^{(1)}_\lambda = \Lambda/\Lambda\mathfrak{n}_\lambda$ is a local commutative subalgebra of $F^{(1)}_\lambda$. Thus, in any simple $F^{(1)}_\lambda$-module, there is a vector where $\mathfrak{m}_\lambda$ acts trivially. As discussed before, this means that:

Proposition 3.3. *Any simple $F^{(1)}_\lambda$-module appears as a quotient of $F_\lambda/F_\lambda\mathfrak{m}_\lambda$. If $\widehat{F}_\lambda$ is a free module over $\widehat{\Lambda}$ (necessarily of rank $\#\widehat{W}_\lambda$) then* $\dim F_\lambda/F_\lambda\mathfrak{m}_\lambda = \#\widehat{W}_\lambda$.

Combining this with Theorem 2.17 above, we have that:

Corollary 3.4. *The dimensions of the λ-weight spaces in the simples over F in the fiber over λ have sum $\leq \dim F_\lambda/F_\lambda\mathfrak{m}_\lambda$, and thus $\leq \#\widehat{W}_\lambda$ if F_λ is a free module over Λ.*

The dimensions of the γ-weight spaces in the simple U-modules in the fiber over γ have sum $\leq \frac{1}{\#W_\lambda}\dim F_\lambda/F_\lambda\mathfrak{m}_\lambda$, and thus $\leq \frac{\#\widehat{W}_\lambda}{\#W_\lambda}$ if F_λ is a free module over Λ.

As mentioned in the introduction, this is essentially a repackaging of the techniques in [FO14].

The reflection hypothesis also allows us to define a dual version of the canonical module $C;\lambda$. We can consider the quotient $\tilde{C}'_\lambda$ of the module $\Lambda_{\mathscr{S}}$ by all submodules having trivial intersection with $\mathcal{W}_\lambda(\Lambda_{\mathscr{S}})$.

The algebra $\Lambda^{(1)}_\lambda$ is a Frobenius algebra, so its socle as a $\Lambda^{(1)}_\lambda$-module is 1-dimensional, and every non-zero submodule of $\tilde{C}'_\lambda$ has non-trivial intersection with $\mathcal{W}_\lambda(\Lambda_{\mathscr{S}})$, and thus contains this socle. This shows that the intersection of all non-zero submodules is non-trivial, giving a simple socle $\tilde{C}_\lambda \subset \tilde{C}'_\lambda$ This will sometimes be isomorphic to C_λ, and sometimes not.

3.1. Special cases of interest.

Definition 3.5. *We call a weight* **non-singular** *if $\widehat{W}_\lambda = \{1\}$ and more generally p-***singular** *if $\widehat{W}_\lambda$ has a minimal generating set of p reflections.*

Corollary 3.6. *If λ is non-singular, there is a unique simple Gelfand-Tsetlin module S with $\mathcal{W}_\lambda(S) \cong \mathbb{C}$ and for all other simples S' we have $\mathcal{W}_\lambda(S) = 0$.*

Of course, a natural question to consider is when two non-singular weights λ, μ have the same simple, and when they do not. Of course, they can only give the same simple if $\mu = w \cdot \lambda$ for some $w \in \widehat{W}$. In this case, ${}_\mu F_\lambda$ is the elements of the form $w\ell$, and similarly ${}_\lambda F_\mu$ the elements of the form $w^{-1}\ell'$.

Corollary 3.7. *Given λ and μ as above, we have a simple Gelfand-Tsetlin module S with $\mathcal{W}_\lambda(S) \cong \mathcal{W}_\mu(S) \cong \mathbb{C}$ if and only if ${}_\lambda F_\mu \cdot {}_\mu F_\lambda \not\subset \mathfrak{m}_\lambda$.*

Now assume λ is 1-singular and F_λ is a free module over Λ. In this case, $\widehat{W}_\lambda \cong S_2$, so $F^{(1)}_\lambda$ is 4-dimensional. Thus, there are 3 possibilities for the behavior of such a weight:

Corollary 3.8. *Exactly 1 of the following holds:*

(1) $F^{(1)}_\lambda \cong M_2(\mathbb{C})$ *and there is a unique simple Gelfand-Tsetlin module S with $\mathcal{W}_\lambda(S) \cong \mathbb{C}^2$ and for all other simples it is 0.*

(2) *the Jacobson radical of $F^{(1)}_\lambda$ is 2-dimensional and there are two simple Gelfand-Tsetlin modules S_1, S_2 with $\mathcal{W}_\lambda(S_i) \cong \mathbb{C}$ and for all other simples it is 0.*

(3) *the Jacobson radical of $F^{(1)}_\lambda$ is 3-dimensional and there is a unique simple Gelfand-Tsetlin module S with $\mathcal{W}_\lambda(S) \cong \mathbb{C}$ and for all other simples it is 0.*

4. Coulomb branches

4.1. **Coulomb branches and principal orders.** One extremely interesting example of principal Galois orders are the Coulomb branches defined by Braverman, Finkelberg and Nakajima [BFNb]. These algebras have attracted considerable interest in recent years, and subsume most examples of interesting principal Galois orders known to the author.

There is a Coulomb branch attached to each connected reductive complex group G and representation N. Let $G[t]$ be the Taylor series points of the group G, and $G((t))$ its Laurent series points. Let

$$\mathcal{Y} = (G((t)) \times N[t])/G[t],$$

equipped with its obvious map $\pi\colon \mathcal{Y} \to N((t))$; we can think of this as a vector bundle over the affine Grassmannian $G((t))/G[t]$. Readers who prefer moduli theoretic interpretations can think of this as the moduli space of principal bundles on a formal disk with choice of section and trivialization away from the origin.

Let $H = N_{GL(N)}(G)^\circ$ be the connected component of the identity in the normalizer of G. Let T_G, T_H be compatible maximal tori in the two groups, and B_G, B_H compatible choices of Borels, and $G \subset Q \subset H$ the subgroup generated by G and T_H. Note that $\mathcal{Y}$ has an H-action via $h \cdot (g(t), n(t)) = (hg(t)h^{-1}, hn(t))$. It also carries a canonical principal Q-bundle $\mathcal{Y}_Q$ given by the quotient $G((t)) \times Q \times N[t]$ via the action $g(t) \cdot (g'(t), q, n(t)) = (g'(t)g^{-1}(t), qg^{-1}(0), g(t)n(t)$. We can extend this to an action of $Q \times \mathbb{C}^*$ where the factor of $\mathbb{C}^*$ acts by the loop scaling.

Definition 4.1. *The (quantum) Coulomb branch is the convolution algebra*

$$\mathcal{A} = H_*^{Q\times\mathbb{C}^*}(\pi^{-1}(N[t])),$$

It might not be readily apparent what the algebra structure on this space is. However, it is unique determined by the fact that it acts on $H_*^{Q\times\mathbb{C}^*}(N[t]) = H^*_{Q\times\mathbb{C}^*}(*)$ by

$$a \star b = \pi_*(a \cap \iota(b)) \tag{4.1}$$

where ι is the inclusion of this algebra into $\mathcal{A}$ as the Chern classes of the principal bundle $\mathcal{Y}_Q$ and the obvious inclusion of $\mathbb{C}[h] \cong H_*^{\mathbb{C}^*}(N[t])$. We further obtain a module structure on the $Q \times \mathbb{C}^*$-equivariant homology of any $G[t]$-invariant subvariety in $N[t]$; applying this to $\{0\}$, we obtain an action on Γ, which sends the subalgebra discussed above to multiplication operators. Obviously there are a lot of technical issues that are being swept under the rug here; a reader concerned on this point should refer to [BFNb] for more details.

Let $T_F = Q/G = T_H/T_G$, and $\mathfrak{t}_F$ the Lie algebra of this group. The subalgebra $H^*_{Q/G\times\mathbb{C}^*}(*) = \mathrm{Sym}(\mathfrak{t}_F^*)[h] \subset \mathcal{A}$ induced by the $Q\times\mathbb{C}^*$-action is central; borrowing terminology from physics, we call these **flavor parameters**. We can thus consider the quotient of $\mathcal{A}$ by a maximal ideal in this ring. This quotient is what is called the "Coulomb branch" in [BFNb, Def. 3.13] and our Definition 4.1 matches the deformation constructed in [BFNb, §3(viii)]. We'll distinguish this situations by referring to them fixed/generic flavor parameters.

We let W be the Weyl group of G (which is also the Weyl group of Q), let $V = \mathfrak{t}_H^*\oplus\mathbb{C}\cdot h$ where $\mathfrak{t}_H$ is the (abstract) Cartan Lie algebra of H and let $\mathcal{M}$ the cocharacter lattice of T_G, acting by the h-scaled translations

$$\chi\cdot(\nu+kh) = \nu + k\langle\chi,\nu\rangle + kh.$$

Note that the action has finite stabilizers on any point where $h\neq 0$, but any point with $h=0$ will have infinite stabilizer. We'll ultimately only be interested in modules over the specialization $h=1$, so this will not cause an issue for the moment. Note that

$$\Lambda\cong H^*_{T_H\times\mathbb{C}^*}(*) = \mathrm{Sym}^\bullet(\mathfrak{t}_H)[h] \qquad \Gamma\cong H^*_{Q\times\mathbb{C}^*}(*) = \mathrm{Sym}^\bullet(\mathfrak{t}_H)^W[h],$$

and $\mathcal{M}\ltimes W$ is the extended affine Weyl group of Q. Localization in equivariant cohomology shows that the action of (4.1) induces an inclusion $\mathcal{A}\hookrightarrow\mathcal{K}_\Gamma$ for the data above; see [BFNb, (5.18) & Prop. 5.19]. Thus, it immediately follows that:

Proposition 4.2. *The Coulomb branch is a principal Galois order for this data.*

If we fix the flavor parameters, the result will also be a principal Galois order for appropriate quotient of Λ.

The flag order attached to this data also has an interpretation as the flag BFN algebra from [Weba, Def. 3.2]. Let $\mathfrak{X} = (G((t))\times N[t])/I$, where I is the standard Iwahori, $\pi_{\mathfrak{X}}\colon \mathfrak{X}\to N((t))$ the obvious map and ${}_0\mathfrak{X}_0 = \pi_{\mathfrak{X}}^{-1}(N[t])$.

Definition 4.3. *The* **Iwahori Coulomb branch** *is the convolution algebra*

$$F = H_*^{T_H\times\mathbb{C}^*}({}_0\mathfrak{X}_0).$$

This is the Morita flag order F_D associated to $\mathcal{A}$ with $D = \mathrm{End}_\Gamma(\Lambda)$ the nilHecke algebra of W, as is shown in [Weba, Thm. 3.3].

As mentioned before, we wish to consider the specializations of these algebras where $h=1$. These are again principal/flag Galois orders in their own right, but are harder to interpret geometrically. Note that by homogeneity, the specializations of this algebra at all different non-zero values of h are isomorphic. The specialization $h=0$ is quite different in nature, since in this case, the action of $\mathcal{M}$ is trivial.

4.2. **Representations of Coulomb branches.** In this case, the algebra $F_\lambda^{(1)}$ has a geometric interpretation. Since we assume that $h=1$, when we interpret λ as an element of the Lie algebra $\mathfrak{t}_H\oplus\mathbb{C}$, the second component is 1. Let G_λ be the Levi subgroup of G which only contains the roots which are integral at λ, and N_λ the span of the weight spaces for weights integral on λ. Let B_λ be the Borel in G_λ such that $\mathrm{Lie}(B_\lambda)$ is generated by the roots α such that $\langle \lambda, \alpha\rangle$ is negative and those in the fixed Borel $\mathfrak{b}_G$ such that $\langle \lambda, \alpha\rangle = 0$; this is the unique Borel in $glsPla_\lambda$ such that $G_\lambda \cap B_\lambda = G_\lambda \cap B_G$.

The element λ integrates to a character acting on N_λ. Let N_λ^- be the subspace of N_λ which is non-positive for the cocharacter corresponding to λ; this subspace is preserved by the action of B_λ. Consider the associated vector bundle $X_\lambda = (G_\lambda \times N_\lambda^-)/B_\lambda$ and p_λ the associated map $p\colon X_\lambda \to N_\lambda$. If $W_\lambda \neq \{1\}$, then there is also a parabolic version of these spaces. Let $P_\lambda \subset G_\lambda$ be the parabolic corresponding to W_λ, and let $Y_\lambda = (G_\lambda \times N_\lambda^-)/P_\lambda$.

As usual, we have associated Steinberg varieties:

$$\mathbb{X}_\lambda = X_\lambda \times_{N_\lambda} X_\lambda = \{(g_1B_\lambda, g_2B_\lambda, n) \mid n \in g_1N_\lambda^- \cap g_2N_\lambda^-\}$$
$$_\lambda\mathbb{X}_\mu = X_\lambda \times_{N_\lambda} X_\mu = \{(g_1B_\lambda, g_2B_\mu, n) \mid n \in g_1N_\lambda^- \cap g_2N_\mu^-\}$$
$$\mathbb{Y}_\lambda = Y_\lambda \times_{N_\lambda} Y_\lambda = \{(g_1P_\lambda, g_2P_\lambda, n) \mid n \in g_1N_\lambda^- \cap g_2N_\lambda^-\}$$
$$_\lambda\mathbb{Y}_\mu = Y_\lambda \times_{N_\lambda} Y_\mu = \{(g_1P_\lambda, g_2P_\mu, n) \mid n \in g_1N_\lambda^- \cap g_2N_\mu^-\}$$

Recall that the **Borel-Moore homology** of an algebraic variety X over $\mathbb{C}$ is the hypercohomology of the dualizing sheaf $\mathbb{D}\mathbb{C}_X$ indexed backwards. We use the same convention for equivariant Borel-Moore homology:

$$H_i^{BM}(X) = \mathbb{H}^{-i}(X_{\mathrm{an}}; \mathbb{D}\mathbb{C}_X) \qquad H_i^{BM,G}(X) = \mathbb{H}_G^{-i}(X_{\mathrm{an}}; \mathbb{D}\mathbb{C}_X).$$

Note that this convention makes $H_*^{BM,G}(X)$ into a module over $H_G^*(X)$ which is homogenous when this ring is given the negative of its usual homological grading; similarly, the group $H_i^{BM,G}(X)$ must be 0 if $i > \dim_{\mathbb{R}} X$, but this can be non-zero in infinitely many negative degrees. We let $\widehat{H}_*^{BM,G_\lambda}(X)$ denote the completion of G_λ-equivariant Borel-Moore homology this respect to its grading, with all elements of degree $\leq k$ being a neighborhood of the identity for all k.

The Borel-Moore homology $H_*^{BM}(\mathbb{X}_\lambda)$ has a convolution algebra structure, and $H_*^{BM}(_\lambda\mathbb{X}_\mu)$ a bimodule structure defined by [CG97, (2.7.9)].

Theorem 4.4. *We have isomorphisms of algebras and bimodules*

(4.2) $$F_\lambda^{(1)} \cong H_*^{BM}(\mathbb{X}_\lambda) \qquad {_\lambda F_\mu^{(1)}} \cong H_*^{BM}(_\lambda\mathbb{X}_\mu)$$

(4.3) $$\widehat{F}_\lambda \cong \widehat{H}_*^{BM,G_\lambda}(\mathbb{X}_\lambda) \qquad {_\lambda \widehat{F}_\mu} \cong \widehat{H}_*^{BM,G_\lambda}(_\lambda\mathbb{X}_\mu)$$

(4.4) $$U_\lambda^{(1)} \cong H_*^{BM}(\mathbb{Y}_\lambda) \qquad {_\lambda U_\mu^{(1)}} \cong H_*^{BM}(_\lambda\mathbb{Y}_\mu)$$

(4.5) $$\widehat{U}_\lambda \cong \widehat{H}_*^{BM,G_\lambda}(\mathbb{Y}_\lambda) \qquad {_\lambda \widehat{U}_\mu} \cong \widehat{H}_*^{BM,G_\lambda}(_\lambda\mathbb{Y}_\mu)$$

This theorem is a consequence of [Weba, Thm. 4.2], which is proven purely algebraically. H. Nakajima has also communicated a more direct geometric proof to the author. We will include a sketch of that argument here, but there are some slightly subtle points about infinite dimensional topology which we will skip over.

Proof (sketch). Note first how the left and right actions of Λ on F operate. The left action is simply induced by the equivariant cohomology of a point, whereas the right action is by the Chern classes of tautological bundles on $G((t))/I$.

Consider the 1-parameter subgroup $\mathbb{T}$ of $G\times\mathbb{C}^*$ obtained by exponentiating λ. By the localization theorem in equivariant cohomology, the completion $\varinjlim F/\mathfrak{n}_\lambda^N F$ is isomorphic to the completion of the T_H-equivariant Borel-Moore homology of ${}_0\mathfrak{X}_0^{\mathbb{T}}$, completed with respect to the usual grading. This is easily seen from [GKM98, (6.2)(1)]: the T_H-equivariant Borel-Moore homology of the complement of the fixed points is a torsion module whose support avoids λ, since the action of $\mathbb{T}$ is locally free. Thus, after completion, the long exact sequence in Borel-Moore homology gives the desired result. Note that here we also use that since the action of $\mathbb{T}$ on the fixed points is trivial, the completion at any point in $\mathfrak{t}$ gives the same result.

First note that the fixed points $N[t]^{\mathbb{T}}$ are isomorphic to N_λ^- via the map $\tau_\lambda\colon N_\lambda\to N((t))$ sending an element n of weight $-a$ in N_λ to $t^a n$.

We can also apply this to the adjoint representation, and find that the fixed points of the 1-parameter subgroup on $\mathfrak{g}((t))$; this is a copy of $\mathfrak{g}_\lambda$, embedded according the description above. Accordingly, the centralizer of this 1-parameter subgroup in $G((t))$ is a copy of G_λ generated by the roots SL_2's of the roots $t^{-\langle\lambda,\alpha\rangle}\alpha$. The Borel B_λ is the intersection of this copy of G_λ with the Iwahori I.

Now consider the fixed points of $\mathbb{T}$ in $G((t))/I$. Each component of this space is a G_λ-orbit, and these components are in bijection with elements of the orbit $\hat{W}\cdot\lambda$; that is, wI and $w'I$ are in the same orbit if and only if $w\cdot\lambda=w'\cdot\lambda$. If w is of minimal length with $\mu=w\cdot\lambda$, the stabilizer of wI under the action of G_λ is the Borel B_μ. Considering the vector bundles induced by the tautological bundles shows that elements of $\mathfrak{n}_\mu$ act by elements with trivial degree 0 term, i.e. that the homology of this component is ${}_\lambda\widehat{F}_\mu$

Thus, the fixed points $\mathfrak{X}^{\mathbb{T}}$ break into components corresponding to these orbits as well, with the fiber over gwI for $g\in G_\lambda$ and w as defined above is given by gN_μ^-, via the map $g\cdot\tau_\mu$. The map $\pi_{\mathfrak{X}}$ maps this to $N((t))$ via the map $\tau_\lambda\circ\tau_\mu^{-1}\circ g^{-1}$, so its intersection with the preimage of $N[t]$ is $N_\lambda^-\cap gN_\mu^-$.

The relevant T_H-equivariant homology group is thus

$$H_*^{T_H}(\{(gB_\mu,x)\mid g\in G_\mu, x\in N_\lambda^-\cap gN_\mu^-\})\cong H_*^{G_\lambda}({}_\lambda\mathbb{X}_\mu).$$

Taking quotient by $\mathfrak{n}_\lambda$, we obtain the non-equivariant Borel-Moore homology of this variety as desired. This shows that we have a vector space isomorphism in (4.2).

The row of isomorphisms (4.4) follow from the same argument applied to $\pi^{-1}(N[t])$ and the affine Grassmannian.

Note that we have not checked that the resulting isomorphism is compatible with multiplication, and doing so is somewhat subtle. For a finite dimensional manifold X, we have two isomorphisms between $H_*^{\mathbb{T}}(X)$ and $H_*^{\mathbb{T}}(X^{\mathbb{T}})$ after completion at any nonzero point in $\mathfrak{t}$: pullback (defined using Poincaré duality) and pushforward, which differ by the (invertible) Euler class of the normal bundle by the adjunction formula. To obtain an isomorphism $H_*^{\mathbb{T}}(X\times X)$ and $H_*^{\mathbb{T}}(X^{\mathbb{T}}\times X^{\mathbb{T}})$ that commutes with convolution, one must take the middle road between these, using pullback times the inverse of the Euler class of the normal bundle along the first factor, which is the same as the inverse of pushforward times the Euler class of the normal bundle along the second factor (effectively, we use the pushforward isomorphism in the first factor, and the pullback in the second factor). Due to the infinite dimensionality of the factors $\mathfrak{X}$ and $\mathfrak{Y}$, and the nature of the cycles we use, neither the pushforward nor the pullback isomorphisms make sense, but this intermediate isomorphism does.

As we said above, we will not give a detailed account of this isomorphism, since we have already constructed a ring isomorphism using the algebraic arguments of [Weba]. Savvy readers will notice the Euler class we need to invert in [Weba, (4.3a)] □

The stabilizer $\widehat{W}_\lambda$ is always isomorphic to a parabolic subgroup of the original Weyl group W.

Definition 4.5. *We call an orbit* **integral** *if* $\widehat{W}_\lambda \cong W$ *and* $N = N_\lambda$.

One especially satisfying consequence of Theorem 4.4 is that the category of modules with weights in the non-integral orbit is equivalent to that same category for integral orbit but of the Coulomb branch for the corresponding Levi subgroup G_λ and subrepresentation N_λ.

More precisely, fix an orbit $\mathscr{S}$ of $\widehat{W}$, and let $G' = G_\lambda$ and $N' = N_\lambda$ for arbitrary $\lambda \in \mathscr{S}$. Let $\mathscr{S}' \subset \mathscr{S}$ be an orbit of the subgroup $\widehat{W}' \subset \widehat{W}$ generated by the Weyl group of G' and the subgroup $\mathcal{M}$. Let $\Gamma Ц'(\mathscr{S}')$ be the category of weight modules with all weights concentrated $\mathscr{S}'$for the Coulomb branch of (G', N'). Note that since all the orbits of $\mathscr{S}' \subset \mathscr{S}$ are conjugate under the action of W, this category only depends on $\mathscr{S}$. Of course, for this smaller group, $\mathscr{S}'$ is an integral orbit. By Theorem 4.4, we have that:

Corollary 4.6. *We have an equivalence of categories* $\Gamma Ц(\mathscr{S}) \cong \Gamma Ц'(\mathscr{S}')$.

This equivalence does not change the underlying vector space and its weight space decomposition; it simply multiplies the action of elements of F by elements of the appropriate completion of Γ to adjust the relations. This can be proven in the spirit of Theorem 4.4 by presenting the Coulomb branch of (G_λ, N_λ) as the homology of the fixed points of the torus action, and noting that the Euler class of the normal bundle acts invertibly on all the modules in the relevant subcategory.

4.3. **Gradings.** This is a particularly nice description since the convolution algebras in question are graded, and a simple geometric argument shows that they are graded free over the subalgebra $\Lambda^{(1)}_\lambda$, with the degrees of the generators read off from the dimensions of the preimages of the orbits in $\mathbb{X}_\lambda$. For reasons of Poincaré duality, we grade $H^{BM}_*(\mathbb{X}_\lambda)$ so that a cycle of dimension d has degree $\dim X_\lambda - d$, and $H^{BM}_*({}_\lambda\mathbb{X}_\mu)$ so that a cycle of dimension d has degree $\frac{\dim X_\lambda + \dim X_\mu}{2} - d$. This is homogeneous by [CG97, (2.7.9)].

Proposition 4.7. $F^{(1)}_\lambda$ *has a set of free generators with degrees given by* $\dim(N^-_\lambda) - \dim(wN^-_\lambda \cap N^-_\lambda) - \ell(w)$ *ranging over* $w \in \widehat{W}_\lambda$, *identified with the Weyl group of* G_λ.

Proof. The product $(G_\lambda/B_\lambda)^2$ breaks up into finitely many G_λ-orbits, each one of which contains (B_la, wB_λ) for a unique $w \in \widehat{W}_\lambda$. This orbit is isomorphic to an affine bundle over G_λ/B_λ with fiber $B_\lambda/(B_\lambda \cap wB_\lambda w^{-1})$, which is an affine space of dimension $\ell(w)$. Furthermore, the preimage of this orbit in $\mathbb{X}_\lambda$ is a vector bundle of dimension $\dim(wN^-_\lambda \cap N^-_\lambda)$. his means that under the usual grading on the convolution algebra, the fundamental class has degree equal to $\dim X_\lambda$ minus the dimension of this orbit.

These fundamental classes give free generators over $\Lambda_\lambda^{(1)}$, since the homology of each of these vector bundles is free of rank 1. □

In particular, if these degrees are always non-negative, then all elements of positive degree are in the Jacobson radical.

Corollary 4.8. *If* $\dim(N_\lambda^-) - \dim(wN_\lambda^- \cap N_\lambda^-) - \ell(w) \geq 0$ *for all* $w \in \widehat{W}_\lambda$*, then the sum of* $(\dim \mathcal{W}_\lambda(S))^2$ *over all simple Gelfand-Tsetlin modules is*

$$\leq \#\{w \in \widehat{W}_\lambda \mid \dim(N_\lambda^-) - \dim(wN_\lambda^- \cap N_\lambda^-) = \ell(w)\}.$$

Note that the fact that the algebra $F^{(1)}(\mathsf{S})$ is graded allows us to define a graded lift $\widetilde{\Gamma Ц}$ of the category of Gelfand-Tsetlin modules by considering graded modules over $F^{(1)}(\lambda_1, \dots, \lambda_k)$.

Following Ginzburg and Chriss [CG97, 8.6.7], we can restate Theorem 4.4 as

$$F_\lambda^{(1)} \cong \operatorname{Ext}^\bullet\left((p_\lambda)_*\mathbb{C}_{X_\lambda}, (p_\lambda)_*\mathbb{C}_{X_\lambda}\right)$$

$$F^{(1)}(\mathsf{S}) \cong \operatorname{Ext}^\bullet\left(\bigoplus_{i=1}^k (p_{\lambda_i})_*\mathbb{C}_{X_{\lambda_i}}, \bigoplus_{i=1}^k (p_{\lambda_i})_*\mathbb{C}_{X_{\lambda_i}}\right) \tag{4.6}$$

The geometric description of (4.6) has an important combinatorial consequence when combined with the Decomposition Theorem of Beilinson-Bernstein-Deligne-Gabber [CG97, Thm. 8.4.8]:

Theorem 4.9. *The simple Gelfand-Tsetlin modules* S *such that* $\mathcal{W}_{\lambda_i}(S) \neq 0$ *for some* i *are in bijection with simple perverse sheaves* $\mathrm{IC}(Y, \chi)$ *appearing as summands up to shift of* $\oplus_i (p_{\lambda_i})_*\mathbb{C}_{X_{\lambda_i}}$*, with the dimension of* $\mathcal{W}_{\lambda_i}(S)$ *being the multiplicity of all shifts of* $\mathrm{IC}(Y, \chi)$*.*

Note that this result is implicit in [CG97, §8.7] and [Sau, pg. 9] but unfortunately is not stated clearly in either source.

Proof. By the Decomposition Theorem, $(p_\lambda)_*\mathbb{C}_{X_\lambda}$ is a direct sum of shifts of simple perverse sheaves. In the notation of [CG97, Thm. 8.4.8], we have

$$(p_\lambda)_*\mathbb{C}_{X_\lambda} \cong \bigoplus_{(i,Y,\chi)} L_{Y,\chi}(i, \lambda) \otimes \mathrm{IC}(Y, \chi)[i].$$

Let $L_{Y,\chi} \cong \oplus_{i,\lambda_j} L_{Y,\chi}(i, \lambda_j)$ be the $\mathbb{Z}$-graded vector space obtained by summing the multiplicity spaces. Thus, the algebra $F^{(1)}(\mathsf{S})$ is Morita equivalent to

$$A = \operatorname{Ext}^\bullet\Big(\bigoplus_{L_{Y,\chi}\neq 0} \mathrm{IC}(Y, \chi)\Big) \qquad B = \operatorname{Ext}^\bullet\Big(\bigoplus_j (p_{\lambda_j})_*\mathbb{C}_{X_{\lambda_j}}, \bigoplus_{L_{Y,\chi}\neq 0} \mathrm{IC}(Y, \chi)\Big)$$

via the bimodule B. By [CG97, Cor. 8.4.4], this algebra is a positively graded basic algebra with irreps indexed by pairs (Y, χ) such that $L_{Y,\chi} \neq 0$. Thus, the simple representations of $F^{(1)}(\mathsf{S})$ are the images of these 1-dimensional irreps under the Morita equivalence, that is, the multiplicity spaces $L_{Y,\chi}$, with the dimension of the different weight spaces is given by $\dim L_{Y,\chi}(*, \lambda)$, the multiplicity of all shifts of $\mathrm{IC}(Y, \chi)$ in $(p_\lambda)_*\mathbb{C}_{X_\lambda}$. □

The additive category of perverse sheaves given by sums of shifts of summands of $(p_{\lambda_i})_*\mathbb{C}_{X_\lambda}$ satisfies the hypotheses of [Web15, Lem. 1.18], and so by [Web15, Lem. 1.13 & Cor. 2.4], we have that (as proven in [Weba, Cor. 2.20]):

Theorem 4.10. *The classes of the simple Gelfand-Tsetlin modules form a dual canonical basis (in the sense of* [Web15, §2]*) in the Grothendieck group of* $\widetilde{\Gamma Ц}$.

For those who dislike geometry, we only truly need the Decomposition theorem to prove a single purely algebraic, but extremely non-trivial fact:

Corollary 4.11. *The graded algebra $F^{(1)}(\mathsf{S})$ is graded Morita equivalent to an algebra which is non-negatively graded and semi-simple in degree 0.*

This property is called "mixedness" in [BGS96, Web15]; the celebrated recent work of Elias and Williamson [EW] gives an algebraic proof of this fact in some related contexts and could possibly be applied here as well.

4.4. **Applications.** As before, this description is particularly useful in the 1-singular case. In this case, we must have $G_\lambda/B_\lambda \cong \mathbb{P}^1$.

Corollary 4.12. *For a 1-singular weight, we are in situation (1) of Corollary 3.8 if $N_\lambda^- = sN_\lambda^-$, situation (2) if $N_\lambda^- \cap sN_\lambda^-$ is codimension 1 in N_λ^-, and situation (3) otherwise.*

Geometrically, these correspond to the situations where the map $X_\lambda \to G_\lambda \cdot N_\lambda^-$ is (1) the projection $X_\lambda = \mathbb{P}^1 \times N_\lambda^- \to N_\lambda^-$, (2) strictly semi-small or (3) small.

Of course, in the non-singular case, there is no difficulty in classifying the simple modules where a given weight appears: there is always a unique one. However, it is still an interesting question when these simples are the same for 2 different weights. Note that if λ, μ are in the same orbit of $\widehat{W}$, then $N_\lambda = N_\mu$, but the positive subspaces are not necessarily equal.

Corollary 4.13. *Assume that λ, μ are non-singular and in the same $\widehat{W}$-orbit. Then there is a simple Gelfand-Tsetlin module with $\mathcal{W}_\lambda(S)$ and $\mathcal{W}_\mu(S)$ both non-zero if and only if $N_\lambda^- = N_\mu^-$.*

Outside the non-singular case, we have that:

Lemma 4.14. *If λ, μ are non-singular and in the same $\widehat{W}$-orbit, $B_\lambda = B_\mu$ and $N_\lambda^- = N_\mu^-$, then the weight spaces $\mathcal{W}_\lambda(M)$ and $\mathcal{W}_\mu(M)$ are canonically isomorphic for all modules M.*

Proof. The diagonal class $(G_\lambda \times N_\lambda^-)/B_\lambda$ gives the desired isomorphism. □

Since only finitely many subspaces may appear as N_λ^- as λ ranges over an orbit of $\widehat{W}$:

Corollary 4.15. *Every $\widehat{W}$-orbit has a finite complete set in the sense of Definition 2.24.*

Note that this result is not true for a general principal Galois order.

A **seed** is a weight $\gamma \in \mathrm{MaxSpec}(\Gamma)$ which is the image of $\lambda \in \mathrm{MaxSpec}(\Lambda)$ such that $P_\lambda = G_\lambda$.

Theorem 4.16. *If λ is a seed, there is a unique simple Gelfand-Tsetlin U-module S with $\mathcal{W}_\gamma(S) \cong \mathbb{C}$, and for all other simples S' we have $\mathcal{W}_\gamma(S') = 0$. The weight spaces of S satisfy $\dim \mathcal{W}_{\gamma'}(S) \leq \# W_\lambda / W_{\lambda'}$, and this bound is sharp if $N_\lambda^- = N_{\lambda'}^-$.*

Proof. First, we note that $U_\lambda^{(1)} \cong \mathbb{C}$, so this shows the desired uniqueness. The module $eP_\lambda^{(1)}$ is a weight module with S as cosocle satisfying $\dim \mathcal{W}_{\gamma'}(eP_\lambda^{(1)}) \leq \# W_\lambda / W_{\lambda'}$ whenever $\lambda' \in \widehat{W} \cdot \lambda$. This shows that desired upper bound.

We has that $\dim \mathcal{W}_{\gamma'}(S) = \# W_\lambda / W_{\lambda'}$ if and only if S is also the only Gelfand-Tsetlin module with this weight space non-zero, i.e. if and only if ${}_\lambda U_{\lambda'}^{(1)}$ is a Morita equivalence. This is clear if $N_\lambda^- = N_{\lambda'}^-$, since in this case $F_\lambda^{(1)} = F_{\lambda'}^{(1)}$ with ${}_\lambda F_{\lambda'}^{(1)}$ giving the obvious Morita equivalence. □

Note that this shows that the module S discussed above has all the properties proven for the socle of the tableau module in [FGRZ, Th. 1.1]. Using the numbering of that paper,

(ii) The weight γ itself lies in the essential support.
(iii) This follows from Corollary 3.4.
(iv) This follows from Theorem 4.16.
(v) For any parabolic subgroup $W' \subset W$, we can find a λ' such that $N_{\lambda'} = N_\lambda$, and $W' = W_\lambda$. The result then follows from Corollary 3.4.

5. The case of orthogonal Gelfand-Tsetlin algebras

Let us now briefly describe how one can interpret the results of this paper for orthogonal Gelfand-Tsetlin algebras [Maz99] over $\mathbb{C}$ in terms of [KTW+]. As in the introduction, choose a dimension vector $\mathbf{v} = (v_1, \dots, v_n)$ and fix complex numbers $(\lambda_{n,1}, \dots \lambda_{n,v_n}) \in \mathbb{C}^{v_n}$. Let

$$\Omega = \{(i, r) \mid 1 \leq i \leq n, 1 \leq r \leq v_i\}.$$

Let U be the associated orthogonal Gelfand-Zetlin algebra modulo the ideal generated by specializing $x_{n,r} = \lambda_{n,r}$. This is a principal Galois order with the data:

- The ring Λ given by the polynomial ring generated by $x_{i,j}$ with $(i, j) \in \Omega$ and $i < n$. Note that we have not included the variables $x_{n,1}, \dots, x_{n,v_n}$, since these are already specialized to scalars.
- The monoid $\mathcal{M}$ given by the subgroup of $\mathrm{Aut}(\Lambda)$ generated by $\varphi_{i,j}$, the translation satisfying

$$\varphi_{i,j}(x_{k,\ell}) = (x_{k,\ell} + \delta_{ik}\delta_{j\ell})\varphi_{i,j}$$

- The group $W = S_{v_1} \times \cdots \times S_{v_{n-1}}$, acting by permuting each alphabet of variables.

By definition, U is the subalgebra of $\mathcal{K}$ generated by $\Gamma = \Lambda^W$ and the elements

$$X_i^{\pm} = \mp \sum_{j=1}^{v_i} \frac{\prod_{k=1}^{v_{i\pm 1}} (x_{i,j} - x_{i\pm 1,k})}{\prod_{k\neq j} (x_{i,j} - x_{i,k})} \varphi_{i,j}^{\pm}$$

We let $F = F_D$ be the corresponding Morita flag order. This is the subalgebra of $\mathcal{F}$ generated by U embedded in $e\mathcal{F}e \cong \mathcal{K}$ and the nilHecke algebra $D = \mathrm{End}_\Gamma(\Lambda)$.

As mentioned in the introduction, it is proven in [Wee] that:

Theorem 5.1. *We have an isomorphism between the OGZ algebra attached to the dimension vector* $\mathbf{v}$ *and the Coulomb branch for the* (G, N) *introduced in (1.1a–1.1b) at* $h = 1$*, with the variables* $x_{n,1}, \dots, x_{n,v_n}$ *corresponding to the flavor parameters. Thus,* U *is isomorphic to the Coulomb branch with the flavor parameters fixed by* $z_r = \lambda_{n,r} - \frac{n}{2}$*.*

Thus, we can apply the results of Section 4 to OGZ algebras. An element $\lambda \in \mathrm{MaxSpec}(\Lambda)$ is exactly choosing a numerical value $x_{i,r} = \lambda_{i,r}$ for all $(i, r) \in \Omega$, and the corresponding $\gamma \in \mathrm{MaxSpec}(\Gamma)$ only remembers these values up to permutation of the second index. A choice of λ partitions the set Ω according to which coset of $\mathbb{Z}$ the value $\lambda_{i,r}$ lies in. Given a coset $[a] \in \mathbb{C}/\mathbb{Z}$, let

$$\Omega_{[a]} = \{(i, r) \in \Omega \mid \lambda_{i,r} \equiv a \pmod{\mathbb{Z}}\}.$$

The maximal ideal λ has an integral orbit if there is one coset such that $\Omega = \Omega_{[a]}$. In general, let $X = \{[a] \in \mathbb{C}/\mathbb{Z} \mid \Omega_{[a]} = \emptyset\}$.

Note that the representation N is spanned by the dual basis to the matrix coefficients of the maps $\mathbb{C}^{v_k} \to \mathbb{C}^{v_{k+1}}$, which we denote $h^{(k)}_{r,s}$ for $1 \leq r \leq v_k$ and $1 \leq s \leq v_{k+1}$.

Proposition 5.2. *Given* $\lambda \in \mathrm{MaxSpec}(\Lambda)$*, we have that* N_λ *is the span the elements* $h^{(k)}_{r,s}$ *such that* $\lambda_{k,r} - \lambda_{k+1,s} \in \mathbb{Z}$*, and* N_λ^- *is the span of these elements with* $\lambda_{k,r} - \lambda_{k+1,s} \in \mathbb{Z}_{\geq 0}$*.*

Remark 5.3. *Note that equivalence classes of weights in a* $\widehat{W}$*-orbit with* N_λ^- *fixed also appears in the discussion of generic regular modules in* [EMV, §3.3]*. That is, the subspace* N_λ^- *changes precisely when the numerator of one of the Gelfand-Tsetlin formulae vanishes.*

We can encapsulate this with an order on the set Ω which is the coarsest such that $(i, r) \prec (i+1, s)$ if $\lambda_{i,r} - \lambda_{i+1,s} \in \mathbb{Z}_{<0}$ and $(i, r) \succ (i+1, s)$ if $\lambda_{i,r} - \lambda_{i+1,s} \in \mathbb{Z}_{\geq 0}$. Lemma 4.14 then shows that:

Proposition 5.4. *We have a natural isomorphism* $\mathcal{W}_\lambda(M) \cong \mathcal{W}_{\lambda'}(M)$ *for any Gelfand-Tsetlin module* M *over* U *if for all pairs* (i, r) *and* $r \in [1, v_i]$*, we have* $\lambda_{i,r} - \lambda'_{i,r} \in \mathbb{Z}$*, and the induced order on the set* Ω *is the same.*

More generally, the Gelfand-Tsetlin modules over this module are controlled by KLRW algebras, as shown in [KTW$^+$]. If $\mathscr{S}$ is not integral, then by Corollary 4.6, the category $\Gamma\text{Ц}(\mathscr{S})$ is equivalent to the category of Gelfand-Tsetlin modules supported on the same orbit for a tensor product $\otimes_{[a]\in\mathbb{C}/\mathbb{Z}} U_{[a]}$ where $U_{[a]}$ is the OGZ algebra attached to the set $\Omega_{[a]}$, that is, to the dimension vector $\mathbf{v}^{(a)}$ given by the number of indices k such that $\lambda_{i,k} \equiv a \pmod{\mathbb{Z}}$. Since the simple Gelfand-Tsetlin modules over this tensor product are just an outer tensor product of the simple Gelfand-Tsetlin modules over the individual factors (and in fact, the category $\Gamma Ц(\mathscr{S})$ is a Deligne tensor product of the corresponding category for the factors), let us focus attention on the integral case.

5.1. **The integral case.** Let $\mathscr{S}_\mathbb{Z}$ be the $\widehat{W}$-orbit where $\lambda_{i,r} \in \mathbb{Z}$ for all $(i,r) \in \Omega$, and we fix integral values $\lambda_{n,1} \leq \cdots \leq \lambda_{n,v_n}$.

In this case, we are effectively rephrasing [KTW$^+$, Th. 5.2] in slightly different language, and the notation of this paper. Identify $I = \{1, \dots, n-1\}$ with the Dynkin diagram of $\mathfrak{sl}_n$ as usual. Let $\tilde{T}_\mathbf{v}$ be the block of the KLRW algebra as discussed in [KTW$^+$, §3.1], attached to the sequence $(\omega_{n-1}, \cdots, \omega_{n-1})$ with this fundamental weight appearing v_n times and where v_i black strands have label i for all $i \in I$. Note that this algebra contains a central copy of the algebra

$$\Lambda_{\mathscr{S}_\mathbb{Z}} = \bigotimes_{i=1}^{n-1} \mathbb{C}[x_{i,1}, \dots, x_{i,v_i}]^{S_{v_i}},$$

given by the polynomials in the dots which are symmetric under permutation of all strands.

Fix a very small real number $0 < \epsilon \ll 1$. Given a weight λ, we define a map

$$x \colon \Omega \to \mathbb{R} \qquad x(i,s) = \lambda_{i,s} - i\epsilon - s\epsilon^2.$$

Note that under this map, the partial order $\prec$ is compatible with the usual order on $\mathbb{R}$; this map thus gives a canonical way to refine $\prec$ and the order on Ω induced by the usual partial order on $\lambda_{i,s}$ to a total order on Ω. The ϵ term is very important for assuring the compatibility with $\prec$, whereas the ϵ^2 term is essentially arbitrary, and is only there to avoid issues when two strands go to the same place.

Definition 5.5. *Let $w(\lambda)$ be the word in $[1,n]$ given by ordering the elements of Ω according to the function x, and then projecting to the first index.*

Now, consider the idempotent $e(\lambda)$ in $\tilde{T}_\mathbf{v}$ where we place a red strand with label ω_{n-1} at $x(n,r)$ for all $r = 1, \dots, v_n$, and a black strand with label i at $x(i,s)$ for all $i \in I$ and $s = 1, \dots, v_i$. The labels of strands read left to right are just the word $w(\lambda)$.

Note that the isomorphism type of this idempotent only depends on the partial order $\prec$, and it would be the same for any map x that preserves this order. For example, we would match [KTW$^+$] more closely if we used $x(i,s) = 2\lambda_{i,s} - i$ (again with a perturbation to assure all elements have distinct images) which works equally well. This choice matches better with the parameterization of Γ by the variables $w_{i,k}$ used in [BFNa].

Let $\mathsf{S} \subset \mathscr{S}_\mathbb{Z}$ be a finite set. For simplicity, we assume that this set has no pairs of weights that correspond as in Proposition 5.4, up to the action of W. Of course, this set will be complete if every possible partial order $\prec$ that appears in the orbit $\mathscr{S}_\mathbb{Z}$ is realized. Let e_S be the sum of these idempotents in $\tilde{T}_\mathbf{v}$

Theorem 5.6. *The algebra $\widehat{F}_\mathsf{S}$ is isomorphic to the completion with respect to its grading of $e_\mathsf{S}\tilde{T}_\mathbf{v}e_\mathsf{S}$, and $F_\mathsf{S}^{(1)}$ is isomorphic to $e_\mathsf{S}\tilde{T}_\mathbf{v}e_\mathsf{S}$ modulo all positive degree elements of $\Lambda_{\mathscr{S}_\mathbb{Z}}$.*

This is truly a restatement of [KTW$^+$, Th. 5.2], but can also be derived from Theorem 4.4, using the convolution description of $\tilde{T}_{\mathbf{v}}$ as a convolution algebra based on [Webc, Th. 4.5 & 3.5]. If you prefer to keep $x_{n,r}$ as variables rather than specializing them, then the resulting algebra is the deformation of $\tilde{T}_{\mathbf{v}}$ defined in [Webb, Def. 4.1].

This reduces the question of understanding Gelfand-Tsetlin modules to studying the simple representations of these algebras. The usual theory of translation functors shows that the structure of this category only depends on the stabilizer under the action of S_{v_n} on the element $(\lambda_{n,1}, \dots, \lambda_{n,v_n})$. This is a Young subgroup of the form $S_{\mathbf{g}} = S_{g_1} \times \cdots \times S_{g_\ell}$; of course, a regular block will have all $g_k = 1$. Consider the sequence of dominant weights $\mathbf{g} = (g_1\omega_{n-1}, \dots, g_\ell\omega_{n-1})$. This corresponds to the tensor product $\mathrm{Sym}^{g_1}(\mathbb{C}^{n-1}) \otimes \mathrm{Sym}^{g_2}(\mathbb{C}^{n-1}) \otimes \cdots \otimes \mathrm{Sym}^{g_\ell}(\mathbb{C}^{n-1})$, and so by [KTW$^+$, Prop. 3.1], we have that: $K^0(\tilde{T}^{\mathbf{g}}_{\mathbf{v}}) \cong U(\mathbf{g})$ where $\mathfrak{n}_-$ is the algebra of $n \times n$ strictly lower triangular matrices and

$$U(\mathbf{g}) := U(\mathfrak{n}_-) \otimes \mathrm{Sym}^{g_1}(\mathbb{C}^{n-1}) \otimes \mathrm{Sym}^{g_2}(\mathbb{C}^{n-1}) \otimes \cdots \otimes \mathrm{Sym}^{g_\ell}(\mathbb{C}^{n-1}).$$

While we have a general theorem connecting simples over $\tilde{T}^{\mathbf{g}}_{\mathbf{v}}$ to the dual canonical basis of $U(\mathbf{g})$, because we are looking at a particularly simple special case, this combinatorics simplifies.

Following the work of Leclerc [Lec04] and the relation of this work to KLR algebras discussed in [KR11], we can give a simple indexing set of this dual canonical basis. Consider a simple Gelfand-Tsetlin module S, and the set $\mathbb{L}(S)$ of words $w(\lambda)$ for $\mathcal{W}_\lambda(S) \neq 0$. We order words in the set $[1, n]$ lexicographically, with the rule that $(i_1, \dots, i_{k-1}) > (i_1, \dots, i_k)$.

Definition 5.7. *We call a word* **good** *if it is minimal in lexicographic order amongst $\mathbb{L}(S)$ for some simple S. Since $\mathbb{L}(S)$ is finite, obviously every simple has a unique good word.*

Let $\mathcal{GL}$ be the set of words of the form $(k, k-1, \cdots, k-p)$ for $k \leq n-1$, and $0 \leq p < k$, and $\mathcal{GL}'$ be the set of words of the form $(n, n-1, \cdots, n-p)$ for $0 \leq p < n$; as noted in [Lec04, §6.6], these together form the good Lyndon words of the A_n root system in the obvious order on nodes in the Dynkin diagram (which we identify with $[1, n]$).

Definition 5.8. *We say a word $\mathbf{i}$ is* **goodly** *if it is the concatenation $\mathbf{i} = a_1 \cdots a_p b_1 \cdots b_{v_n}$ of words for $a_k \in \mathcal{GL}$, and $b_k \in \mathcal{GL}'$ satisfying $a_1 \leq a_2 \leq \cdots \leq a_p$ in lexicographic order.*

Assume for simplicity that the central character $(\lambda_{n,1}, \dots, \lambda_{n,v_n})$ is regular, that is, $S_{\mathbf{g}} = \{1\}$. In this case, a goodly word can always be realized as $w(\lambda^{(i)})$ for a weight $\lambda^{(\mathbf{i})}$ chosen as follows: pick integers $\mu_1, \dots, \mu_p$ so that $\mu_1 < \cdots < \mu_p < \lambda_{n,1} < \cdots < \lambda_{n,v_n}$. Now, choose the set $\lambda^{(\mathbf{i})}_{i,*}$ so that μ_k appears (always with multiplicity 1) if and only if i appears as a letter in a_k, and $\lambda_{n,q}$ if and only if i appears as a letter in b_q. This weight depends on the choice of μ_*, but all these choices are equivalent via Lemma 4.14.

Theorem 5.9. *The map sending a simple Gelfand-Tsetlin module to its good word is a bijection, and a word is good if and only if it is goodly.*

Note that implicit in the theorem above is that we consider the set of all good words for all different $\mathbf{v}$'s, but $\mathbf{v}$ is easily reconstructed from the word, by just letting v_i be the number of times i appears.

Proof. Note that the words in $\mathcal{GL}$ index cuspidal representations of the KLR algebra of type A in the sense of Kleshchev-Ram [KR11]; thus concatenations of these words in increasing lexicographic order give the good words for type A, and the lex maximal word in the different simple representations of the KLR algebra of type A by [KR11, Th. 7.2].

On the other hand, the words $\mathcal{GL}'$ give the idempotents corresponding to the different simples over the cyclotomic quotient $T^{\omega_{n-1}}$, which are all 1-dimensional. By [Web17, Cor. 5.23], every simple over $\tilde{T}_{\mathbf{v}}$ is the unique simple quotient of a standardization of a simple module over the KLR algebra $\tilde{T}^{\emptyset}$ and v_n simple modules over $T^{\omega_{n-1}}$. The former module gives the desired words $a_1 \cdots a_p$ as described above, and the latter v_n simples give the words in $\mathcal{GL}'$. By construction, the resulting concatenation is lex minimal amongst those with $e(\mathbf{i})$ not killing the standard module, and survives in the simple quotient since the image of $e(\mathbf{i})$ generates. Let L be the corresponding simple $\tilde{T}_{\mathbf{v}}\tilde{T}$-module.

The image $e_{\mathsf{S}}L$ gives a simple module over $F_{\mathsf{S}}^{(1)}$ for any set S containing the weight $\lambda^{(\mathbf{i})}$ and thus a simple Gelfand-Tsetlin-module S by Theorem 2.23. We claim that $\mathbf{i}$ is the good word for this simple, since for any other word that appears as $w(\lambda) < \mathbf{i}$, we can add λ to S, and see that by the properties of L, we have that $\mathcal{W}_\lambda(S) = e(\lambda)L = 0$. Similarly, this shows that S is the unique Gelfand-Tsetlin-module with this property since L is uniquely characterized by this property; any other simple S' comes from a simple $\tilde{T}_{\mathbf{v}}$ representation L', which is the quotient of the standardization of a different word $\mathbf{i}'$ of the form in the statement of the theorem. As we've already argued, this means that $\mathbf{i}' \neq \mathbf{i}$ is its good word. This shows uniqueness and completes the proof. □

Example 5.10. *For example, the case of integral Gelfand-Tsetlin modules of $\mathfrak{sl}_3$ corresponds to $\mathbf{v} = (1, 2, 3)$. Thus, the good words are of the form:*

$$(1|2|2|3|3|3) \qquad (2,1|2|3|3|3)$$

$$(1|2|3,2|3|3) \qquad (1|2|3|3,2|3) \qquad (1|2|3|3|3,2)$$

$$(2,1|3,2|3|3) \qquad (2,1|3|3,2|3) \qquad (2,1|3|3|3,2)$$

$$(1|3,2|3,2|3) \qquad (1|3|3,2|3,2) \qquad (1|3,2|3|3,2)$$

$$(3,2,1|3,2|3) \qquad (3|3,2,1|3,2) \qquad (3,2,1|3|3,2)$$

$$(3,2|3,2,1|3) \qquad (3|3,2|3,2,1) \qquad (3,2|3|3,2,1)$$

We've included vertical bars | between the Lyndon factors of each word.

In order to construct the actual weights appearing, we choose

$$\mu_1 = -2 < \mu_2 = -1 < \mu_3 = 0 < \lambda_{3,1} = 1 < \lambda_{3,2} = 2 < \lambda_{3,3} = 3.$$

In the usual notation for Gelfand-Tsetlin weights, we have the corresponding weight spaces $\lambda^{(\mathbf{i})}$ for the words above are:

$$\begin{matrix} 1 & & 2 & & 3 \\ & -1 & & 0 & \\ & & -2 & & \end{matrix} \qquad \begin{matrix} 1 & & 2 & & 3 \\ & -1 & & 0 & \\ & & -1 & & \end{matrix}$$

$$\begin{array}{ccccc} 1 & & 2 & & 3 \\ & -1 & & 1 & \\ & & -2 & & \end{array} \qquad \begin{array}{ccccc} 1 & & 2 & & 3 \\ & -1 & & 2 & \\ & & -2 & & \end{array} \qquad \begin{array}{ccccc} 1 & & 2 & & 3 \\ & -1 & & 3 & \\ & & -2 & & \end{array}$$

$$\begin{array}{ccccc} 1 & & 2 & & 3 \\ & -2 & & 1 & \\ & & -2 & & \end{array} \qquad \begin{array}{ccccc} 1 & & 2 & & 3 \\ & -2 & & 2 & \\ & & -2 & & \end{array} \qquad \begin{array}{ccccc} 1 & & 2 & & 3 \\ & -2 & & 3 & \\ & & -2 & & \end{array}$$

$$\begin{array}{ccccc} 1 & & 2 & & 3 \\ & 1 & & 2 & \\ & & -2 & & \end{array} \qquad \begin{array}{ccccc} 1 & & 2 & & 3 \\ & 2 & & 3 & \\ & & -2 & & \end{array} \qquad \begin{array}{ccccc} 1 & & 2 & & 3 \\ & 1 & & 3 & \\ & & -2 & & \end{array}$$

$$\begin{array}{ccccc} 1 & & 2 & & 3 \\ & 1 & & 2 & \\ & & 1 & & \end{array} \qquad \begin{array}{ccccc} 1 & & 2 & & 3 \\ & 2 & & 3 & \\ & & 2 & & \end{array} \qquad \begin{array}{ccccc} 1 & & 2 & & 3 \\ & 1 & & 3 & \\ & & 1 & & \end{array}$$

$$\begin{array}{ccccc} 1 & & 2 & & 3 \\ & 1 & & 2 & \\ & & 2 & & \end{array} \qquad \begin{array}{ccccc} 1 & & 2 & & 3 \\ & 2 & & 3 & \\ & & 3 & & \end{array} \qquad \begin{array}{ccccc} 1 & & 2 & & 3 \\ & 1 & & 3 & \\ & & 3 & & \end{array}$$

Thus, each generic integral block for $\mathfrak{gl}_3$ has 17 simple Gelfand-Tsetlin modules.

This Theorem is a little more awkward to state for the singular case where $S_{\mathbf{g}} \neq \{1\}$. For slightly silly reasons, the good words as we have defined them depend on the choice of $\lambda_{n,*}$, but we can still consider goodly words $\mathbf{i} = a_1 \cdots a_p b_1 \cdots b_{v_n}$ and the associated weight $\lambda^{(\mathbf{i})}$. Note that this now only depends on the choice of $b_1, \dots, b_{v_n}$ up to permutations under $S_{\mathbf{g}}$. Using the fact that weight spaces of $\mathrm{Sym}^{g_i}(\mathbb{C}^n)$ are all 1 dimensional, we can similarly argue that:

Proposition 5.11. *For each word $\mathbf{i} = a_1 \cdots a_p b_1 \cdots b_{v_n}$ which is lex maximal in its $S_{\mathbf{g}}$-orbit, there is a unique simple Gelfand-Tsetlin module S such that $\mathcal{W}_{\lambda^{(\mathbf{i})}}(S) \neq 0$, and $\mathcal{W}_{\lambda^{(\mathbf{i}')}}(S) = 0$ for all $\mathbf{i}'$ of the same form with $\mathbf{i}' < \mathbf{i}$.*

We will not prove this fact since it involves a considerable investment in combinatorics we do not want to take the space for here, but one can show that translation from a regular central character to the singular one fixed above kills the simples whose good word is not lex-maximal in their $S_{\mathbf{g}}$-orbit, and induces a bijection between the remaining simples.

Note that in the course of these proofs, we have also shown that:

Proposition 5.12. *If S is a complete set, then $\widehat{F}_{\mathsf{S}}$ is Morita equivalent to the completion with respect to its grading of $\tilde{T}^{\mathbf{g}}_{\mathbf{v}}$ for $\mathbf{g} = (g_1\omega_{n-1}, \dots, g_\ell\omega_{n-1})$, and $F^{(1)}_{\mathsf{S}}$ to the quotient of this algebra by positive degree elements of $\Lambda_{\mathscr{S}_{\mathbb{Z}}}$.*

Proof. Since we will never have a black strand between red strands that correspond to $\lambda_{n,k} = \lambda_{n,k+1}$, we have that $e(\lambda) \in \tilde{T}^{\mathbf{g}}_{\mathbf{v}}$ embedded as in [Web17, Prop. 4.21] by "zipping" the red strands. Thus, $\widehat{F}_{\mathsf{S}}$ maps into the completion of this algebra, and to show Morita equivalence, we need to show that the idempotents $e(\lambda)$ for $\lambda \in \mathsf{S}$ generate $\tilde{T}^{\mathbf{g}}_{\mathbf{v}}$ as a 2-sided ideal. This follows because Theorem 5.9 and Proposition 5.11 show that the number of distinct simple Gelfand-Tsetlin-modules is equal to the number of graded simple $\tilde{T}^{\mathbf{g}}_{\mathbf{v}}$-modules. □

Glossary

Λ	A Noetherian algebra with a W-action. After Section 3, assumed to be the symmetric algebra $\mathrm{Sym}^\bullet(V) = \mathbb{C}[V^*]$	1–3, 5–8, 10–17, 22, 23, 28, 29
Γ	The W-invariants Λ^W	1, 2, 4–7, 9, 10, 12, 13, 15, 16, 19, 22, 28, 29
U	A principal Galois order, usually satisfying $U = eFe$.	1, 2, 6, 7, 9–11, 13, 14, 22, 28, 29
$\widehat{U}_\gamma$	The completion $\varprojlim U/(U\mathsf{m}_\gamma^N + \mathsf{m}_\gamma^N U)$	1, 2
F	A flag Galois order, usually satisfying $U = eFe$.	2, 5–7, 9–14, 16–19, 23, 28
W	A finite group acting on Λ. After Section 3, assumed to be a complex reflection group acting on V by a reflection representation	2, 5–7, 13, 16, 19, 22, 24, 28, 29
$\widehat{F}_\lambda$	The completion of F_λ in the $\mathfrak{n}_\lambda$-adic topology	2, 7–9, 13, 14, 17
$\widehat{W}_\lambda$	The stabilizer of $\lambda \in \mathrm{MaxSpec}(\Lambda)$ under the action of $\widehat{W}$	2, 7–15, 19, 20, 28, 29
$\widehat{W}$	The semi-direct product $\widehat{W} = \mathcal{M} \ltimes W$	2–5, 7, 8, 10–14, 19, 21–24, 28, 29
$\widehat{\Lambda}$	The completion of Λ at a fixed maximal ideal $\mathfrak{mm}_\lambda$	2, 8, 14
W_λ	The stabilizer of $\lambda \in \mathrm{MaxSpec}(\Lambda)$ under the action of W	2, 7, 9, 10, 13, 14, 17, 22
$\widehat{\Lambda}_\lambda$	The completion of Λ_λ in the $\mathfrak{n}_\lambda$-adic topology	2, 8, 9, 12
$F_\lambda^{(N)}$	The quotient algebra $F_\lambda/F_\lambda\mathfrak{n}_\lambda^N = \mathrm{End}(P_\lambda^{(N)})$	2, 9–11, 14–17, 19, 20, 22
$\mathfrak{m}_\lambda$	The maximal ideal in Λ corresponding to $\lambda \in \mathrm{MaxSpec}(\Lambda)$	2, 6–9, 11, 14, 28, 29
G	The gauge group of the Coulomb branch.	3, 4, 15–18, 23, 28
N	The matter representation of the Coulomb branch.	3, 4, 15, 16, 18, 19, 23, 29
G_λ	The Levi subgroup in G corresponding to $\widehat{W}_\lambda \subset W$.	4, 16–19, 21, 22, 29
P_λ	The parabolic subgroup in G corresponding to negative weights of λ.	4, 17, 22, 28, 29

N_λ^-	The P_λ-submodule of N_λ where λ acts by non-positive weights.	4, 17–23
$\tilde{T}_{\mathsf{v}}$	The block of the KLRW algebra as discussed in [KTW+, §3.1], attached to the sequence $(\omega_{n-1}, \cdots, \omega_{n-1})$ with this fundamental weight appearing v_n times and where v_i black strands have label i for all $i \in I$	4, 24–27
L	The fraction field of Λ	5, 12, 28
K	The fraction field of Γ, which is also the fixed field L^W	5, 7, 8, 10–12, 29
$\mathcal{M}$	A fixed submonoid of $\mathrm{Aut}(\Lambda)$ which is normalized by W	5, 7, 16, 22, 28
$\mathcal{L}$	The smash product $L\#\mathcal{M}$	5, 6, 28, 29
$\mathcal{F}$	The smash product $\mathcal{L}\#W$	5, 6, 8, 11, 23, 29
$\mathcal{K}$	The invariants $\mathcal{L}^W$	5–7, 23, 29
$\mathcal{K}_\Gamma$	The standard order $\{X \in \mathcal{K} \mid X(\Gamma) = \Gamma\}$	5, 6, 16
$\mathcal{F}_\Lambda$	The standard flag order $\{X \in \mathcal{F} \mid X(\Lambda) = \Lambda\}$	5, 6
D	A subalgebra satisfying $\Lambda\#W \subset D \subset \mathrm{End}_\Gamma(\Lambda)$.	6, 13, 16, 29
F_D	A flag Galois order canonically constructed from U and D by considering $De \otimes_\Gamma U \otimes_\Gamma eD$.	6, 13, 16, 23
$\mathcal{W}_\lambda$	The functor of taking generalized weight space for a maximal ideal in Λ or Γ	6, 7, 9, 10, 12–15, 20–23, 25, 27, 29
Λ_λ	The fixed points $\Lambda_\lambda = \Lambda^{\widehat{W}_\lambda}$	7–10, 12, 13, 28, 29
F_λ	The intersection $F \cap K \cdot \widehat{W}_\lambda$	7–12, 14, 15, 28
$\mathfrak{n}_\lambda$	The maximal ideal in Λ_λ given by $\mathfrak{m}_\lambda \cap \Lambda_\lambda$	8–12, 14, 18, 28, 29
$P_\lambda^{(N)}$	The quotient module $F/F\mathfrak{n}_\lambda^N$	9, 10, 22, 28
$\Lambda_\lambda^{(1)}$	The quotient $\Lambda/\Lambda\mathfrak{n}_\lambda$	10, 11, 14, 19, 20
${}_\lambda W_\mu$	The elements of $\widehat{W}$ such that $w \cdot \mu = \lambda$	10–12, 29
${}_\lambda F_\mu$	The intersection $F \cap K \cdot {}_\lambda W_\mu$	10–12, 14, 17
$F(\mathsf{S})$	The algebra defined in (2.3) that naturally acts on $\bigoplus_{\lambda \in \mathscr{S}} \mathcal{W}_\lambda$.	12
$F^{(N)}(\mathsf{S})$	The quotient $F(\mathsf{S})/\mathfrak{n}^N F(\mathsf{S})$	12, 20, 21
V	A vector space equipped with a W action that we hold fixed	13, 16, 28
Q	The subgroup in $GL(V)$ generated by G and a torus of the normalizer $N(G)$.	16
N_λ	The subspace of N where the cocharacter λ acts by integral weights.	16–19, 21–23, 28
B_λ	The unique Borel subgroup in P_λ such that $B_\lambda \cap G = B \cap G$.	17–19, 21
$\mathbf{v}$	The dimension vector $\mathbf{v} = (v_1, \ldots, v_n)$ corresponding to the quiver gauge theory that gives an OGZ algebra	22–26, 29
Ω	The index set $\Omega = \{(i, r) \mid 1 \leq i \leq n, 1 \leq r \leq v_i\}$	22–24, 29
x	The function $\Omega \to \mathbb{R}$ defined by $x(i, s) = \lambda_{i,s} - i\epsilon - s\epsilon^2$ for some $0 < \epsilon \ll 1$	24, 29

$w(\lambda)$ The word given by the first indices of the elements of Ω, orderd according to the function x. 24, 25

References

[BDGH] Mathew Bullimore, Tudor Dimofte, Davide Gaiotto, and Justin Hilburn, *Boundaries, mirror symmetry, and symplectic duality in 3d $\mathcal{N} = 4$ gauge theory.*

[BFNa] A. Braverman, M. Finkelberg, and H. Nakajima, *Coulomb branches of 3d $\mathcal{N} = 4$ quiver gauge theories and slices in the affine Grassmannian (with appendices by Alexander Braverman, Michael Finkelberg, Joel Kamnitzer, Ryosuke Kodera, Hiraku Nakajima, Ben Webster, and Alex Weekes)*, arXiv:1604.03625.

[BFNb] Alexander Braverman, Michael Finkelberg, and Hiraku Nakajima, *Towards a mathematical definition of Coulomb branches of 3-dimensional $\mathcal{N} = 4$ gauge theories, II*, arXiv:1601.03586.

[BGS96] Alexander Beilinson, Victor Ginzburg, and Wolfgang Soergel, *Koszul duality patterns in representation theory*, J. Amer. Math. Soc. **9** (1996), no. 2, 473–527.

[BLPW12] Tom Braden, Anthony Licata, Nicholas Proudfoot, and Ben Webster, *Hypertoric category $\mathcal{O}$*, Adv. Math. **231** (2012), no. 3-4, 1487–1545.

[Bro10] Michel Broué, *Introduction to complex reflection groups and their braid groups*, Lecture Notes in Mathematics, vol. 1988, Springer-Verlag, Berlin, 2010. MR 2590895

[CG97] Neil Chriss and Victor Ginzburg, *Representation theory and complex geometry*, Birkhäuser Boston Inc., Boston, MA, 1997. MR 98i:22021

[DFO94] Yu. A. Drozd, V. M. Futorny, and S. A. Ovsienko, *Harish-Chandra subalgebras and Gel′fand-Zetlin modules*, Finite-dimensional algebras and related topics (Ottawa, ON, 1992), NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci., vol. 424, Kluwer Acad. Publ., Dordrecht, 1994, pp. 79–93. MR 1308982

[EMV] Nick Early, Volodymyr Mazorchuk, and Elizaveta Vishnyakova, *Canonical Gelfand-Zeitlin modules over orthogonal Gelfand-Zeitlin algebras*, to appear in IMRN; arXiv:1709.01553.

[EW] Ben Elias and Geordie Williamson, *The Hodge theory of Soergel bimodules*, arXiv:1212.0791.

[FGR16] Vyacheslav Futorny, Dimitar Grantcharov, and Luis Enrique Ramirez, *Singular Gelfand-Tsetlin modules of $\mathfrak{gl}(n)$*, Adv. Math. **290** (2016), 453–482. MR 3451929

[FGRZ] Vyacheslav Futorny, Dimitar Grantcharov, Luis Enrique Ramirez, and Pablo Zadunaisky, *Bounds of Gelfand-Tsetlin multiplicities and tableaux realizations of Verma modules*, arXiv:1710.04186.

[FGRZ18] Vyacheslav Futorny, Dimitar Grantcharov, Luis Enrique Ramírez, and Pablo Zadunaisky, *Gelfand-Tsetlin Theory for Rational Galois Algebras*, arXiv:1801.09316 (2018).

[FO10] Vyacheslav Futorny and Serge Ovsienko, *Galois orders in skew monoid rings*, Journal of Algebra **324** (2010), no. 4, 598–630.

[FO14] ———, *Fibers of characters in Gelfand-Tsetlin categories*, Trans. Amer. Math. Soc. **366** (2014), no. 8, 4173–4208. MR 3206456

[Gin18] Victor Ginzburg, *Nil-Hecke Algebras and Whittaker $\mathscr{D}$-Modules*, Lie Groups, Geometry, and Representation Theory, Springer, 2018, pp. 137–184.

[GKM98] Mark Goresky, Robert Kottwitz, and Robert MacPherson, *Equivariant cohomology, Koszul duality, and the localization theorem*, Invent. Math. **131** (1998), no. 1, 25–83.

[Har] Jonas T. Hartwig, *Principal Galois orders and Gelfand-Zeitlin modules*, arXiv:1710.04186.

[KL09] Mikhail Khovanov and Aaron D. Lauda, *A diagrammatic approach to categorification of quantum groups. I*, Represent. Theory **13** (2009), 309–347.

[KN] Ryosuke Kodera and Hiraku Nakajima, *Quantized Coulomb branches of Jordan quiver gauge theories and cyclotomic rational Cherednik algebras*, arXiv:1608.00875.

[KR11] Alexander Kleshchev and Arun Ram, *Representations of Khovanov-Lauda-Rouquier algebras and combinatorics of Lyndon words*, Math. Ann. **349** (2011), no. 4, 943–975. MR 2777040 (2012b:16078)

[KTW⁺] Joel Kamnitzer, Peter Tingley, Ben Webster, Alex Weekes, and Oded Yacobi, *On category $\mathcal{O}$ for affine Grassmannian slices and categorified tensor products*, `arXiv:1806.07519`.

[Lec04] Bernard Leclerc, *Dual canonical bases, quantum shuffles and q-characters*, Math. Z. **246** (2004), no. 4, 691–732. MR 2045836 (2005c:17019)

[Maz99] Volodymyr Mazorchuk, *Orthogonal Gelfand-Zetlin Algebras I*, Contributions to Algebra and Geometry **40** (1999), no. 2, 399–415.

[MV] Volodymyr Mazorchuk and Elizaveta Vishnyakova, *Harish-chandra modules over invariant subalgebras in a skew-group ring.*

[Nak] Hiraku Nakajima, *Towards a mathematical definition of Coulomb branches of* 3*-dimensional* $\mathcal{N} = 4$ *gauge theories, I*, `arXiv:1503.03676`.

[RZ18] L. E. Ramírez and P. Zadunaisky, *Gelfand-Tsetlin modules over* $\mathfrak{gl}(n)$ *with arbitrary characters*, J. Algebra **502** (2018), 328–346. MR 3774896

[Sau] Julia Sauter, *A survey on Springer theory*, `arXiv:1307.0973`.

[Weba] Ben Webster, *Koszul duality between Higgs and Coulomb categories* $\mathcal{O}$, `arXiv:1611.06541`.

[Webb] ______, *Unfurling Khovanov-Lauda-Rouquier algebras*, `arXiv:1603.06311`.

[Webc] ______, *Weighted Khovanov-Lauda-Rouquier algebras*, `arXiv:1209.2463`.

[Web15] ______, *Canonical bases and higher representation theory*, Compos. Math. **151** (2015), no. 1, 121–166.

[Web17] ______, *Knot invariants and higher representation theory*, Mem. Amer. Math. Soc. **250** (2017), no. 1191, 141.

[Wee] Alex Weekes, *Generators for Coulomb branches of quiver gauge theories*, `arXiv:1903.07734`.

B. Webster: Department of Pure Mathematics, University of Waterloo & Perimeter Institute for Theoretical Physics, Canada

Email address: `ben.webster@uwaterloo.ca`