HOMOMORPHISMS BETWEEN WEYL MODULES FOR $\mathrm{SL}_3(k)$

ANTON COX AND ALISON PARKER

ABSTRACT. We classify all homomorphisms between Weyl modules for $\mathrm{SL}_3(k)$ when k is an algebraically closed field of characteristic at least three, and show that the Hom-spaces are all at most one-dimensional. As a corollary we obtain all homomorphisms between Specht modules for the symmetric group when the labelling partitions have at most three parts and the prime is at least three. We conclude by showing how a result of Fayers and Lyle on Hom-spaces for Specht modules is related to earlier work of Donkin for algebraic groups.

1. INTRODUCTION

Let G be a reductive algebraic group over an algebraically closed field of characteristic $p > 0$. An important class of modules for such a group are the Weyl modules $\Delta(\lambda)$, labelled by dominant weights; these can be constructed (relatively) explicitly, and their heads provide a full set of simple modules for G. (Equivalently one can study the duals of these modules, denoted $\nabla(\lambda)$ which have the advantage of being induced from one-dimensional modules for a Borel subgroup). In determining the structure of such modules, or indeed their cohomology, the calculation of Hom-spaces between them is a useful tool.

Relatively little is known in general about such Hom-spaces. In type A, when λ and μ are related by a (suitable) single reflection, explicit non-zero homomorphisms from $\Delta(\lambda)$ to $\Delta(\mu)$ were constructed (with some restrictions) by Carter and Lusztig [3], and (more generally) by Carter and Payne [4]. The corresponding cases in other types were considered by Franklin [14]. While it is clear that there should be a hierarchy of families of homomorphisms corresponding to different powers of p, the only case where the above results provide a complete classification is when G is SL_2, where it is relatively easy to determine all Hom-spaces exactly [6].

The only other general results in this area, by Andersen [1] and Koppinen [18], concern homomorphisms between modules labelled by weights which are 'close together'. Typically such results show that certain Hom-spaces are non-zero, or in some cases one-dimensional. For weights which are far apart and not related by a single reflection almost nothing is known.

In this paper we will determine all homomorphisms between Weyl modules for SL_3 when $p \geq 3$, and provide a recursive procedure for determining the composition factors arising in the image (or kernel) of such maps in most cases. From these results we will also classify all homomorphisms between Specht modules for the symmetric groups corresponding to three part partitions, when $p \geq 3$.

After a section of preliminaries, we review the SL_3 data concerning p-filtrations that we will need from [21]. This describes certain filtrations of induced modules which will allow us to proceed by induction, together with the set of p-good homomorphisms which will be fundamental in our later constructions. We also recall a theorem of Carter and Payne [4] on the existence of certain homomorphisms. These will be the two key sets of data which we need to determine all possible homomorphisms.

With the notation developed up to that point in place, in Section 4 we can give the strategy to be followed in the remainder of the paper, and in particular the translation functor arguments

The first author was partially supported by Nuffield grant scheme NUF-NAL 02.

2000 Mathematics Subject classification: 20G05 (primary) and 20C30 (secondary).

that form the basis of our argument. This section also contains a more precise description of the main results that will be obtained.

The rest of the paper takes the form of a single inductive argument on the weight labelling our induced module. In Section 5 we show (Theorem 5.1) that $\mathrm{Hom}(\nabla(\lambda), \nabla(\mu))$ is at most one-dimensional. Unfortunately the proof relies on the fact that for one very specific configuration of weights τ and ν there are no homomorphisms from $\nabla(\tau)$ to $\nabla(\nu)$ (Assumption 5.2). Thus we cannot complete this proof until we have classified all possible homomorphisms from induced modules labelled by smaller weights, and hence it forms the first step in our inductive argument.

We then come to the heart of the paper, Sections 6–9. In Section 6 we analyse the homomorphisms constructed by Carter and Payne in the SL_3 case. We begin by giving (up to a controllable ambiguity in certain non-generic cases) a recursive procedure for describing such homomorphisms (Proposition 6.3), and use this to determine which composites of such maps are non-zero (Corollary 6.4).

We begin in Section 7 to determine precisely which Hom-spaces are non-zero, by starting with those obtained inductively from composites of Carter-Payne maps. This uses translation functor arguments, together with the explicit construction of homomorphisms from the previous section. For weights close to the boundary of the dominant region it is more convenient to use a different argument (Proposition 8.3), based on Doty's description of the structure of symmetric powers [10] as reinterpreted in [5].

As well as the Carter-Payne maps, there is another class of maps obtained by 'twisting', which are easy to describe. Unfortunately not every map comes from Carter-Payne composites or twisting, and hence in Section 9 we must construct the remaining maps by hand (by gluing together maps defined on appropriate parts of the p-filtrations of the two modules), to complete the classification. With this we see that Assumption 5.2 does indeed hold, and the main results in the paper now follow by induction.

In Section 10 we illustrate our main results with some examples for the case $p = 3$.

Using a theorem of Carter and Lusztig there is a close relationship between homomorphisms between induced modules for $\mathrm{SL}_3(k)$ and between Specht modules for symmetric groups $k\Sigma_d$. In Section 11 we show how our results also give a classification of homomorphisms between Specht modules labelled by partitions with at most three non-zero parts when $p \geq 3$. Finally, we recall a result of Donkin relating Ext-groups between induced modules for a reductive group and for its Levi factors, and show how this can be used to derive a tensor product theorem of Fayers and Lyle [13] if p is not two.

2. Preliminaries

In this section we shall review the basic results required in this paper. Except where otherwise indicated, this material can be found in found in [17, II, Chapters 1–6]. Although we shall state these results for an arbitrary reductive algebraic group G, defined over an algebraically closed field k of characteristic $p > 0$, for most of this paper we will only be interested in the case when $G = \mathrm{SL}_3(k)$.

We fix a maximal torus $T \subset G$ and hence a weight lattice $X(T)$. There is an associated root system R, in which we choose a set of positive roots R^+. The corresponding set of simple roots will be denoted S. If G is semisimple and simply–connected there is a basis $\{\varpi_\alpha \mid \alpha \in S\}$ of fundamental weights for $X(T)$. Let s_α be the reflection on $E = X(T) \otimes_{\mathbb{Z}} \mathbb{R}$ given by $s_\alpha \lambda = \lambda - \langle \lambda, \alpha^\vee \rangle \alpha$ where $\alpha^\vee$ is the coroot associated to α in $X(T)^*$ and $\langle -, - \rangle$ is the usual bilinear form. The Weyl group W is the group generated by such reflections, while the affine Weyl group W_p is the semidirect product of W with the group $p\mathbb{Z}R$ (acting by translation on E).

Setting $\rho = \frac{1}{2}\sum_{\alpha\in R^+}\alpha \in E$, we define the dot action $w.\lambda = w(\lambda+\rho)-\rho$ of W_p on $X(T)$. Associated to this action is a system of facets; these are all sets of the form

$$F = \{\lambda \in E \mid \langle \lambda+\rho, \alpha^\vee\rangle = n_\alpha p \quad \forall\, \alpha \in R_0^+(F), \\ (n_\alpha - 1)p < \langle \lambda+\rho, \alpha^\vee\rangle < n_\alpha p \quad \forall\, \alpha\in R_1^+(F)\}$$

for some integers n_α and a disjoint decomposition $R^+ = R_0^+(F)\cup R_1^+(F)$. A facet F is called an alcove if $R_0^+(F) = \emptyset$ and a wall if $|R_0^+(F)| = 1$. We shall always assume that $p \geq h$ (the maximum of the Coxeter numbers of the connected components of G), which ensures that every alcove contains a weight. The closure (respectively upper/lower closure) of a facet F is the set obtained from F by replacing both (respectively the right/lefthand) strict inequalities occurring in the defining equations for F by non-strict inequalities. The closure of any alcove is a fundamental domain for the (dot) action of the affine Weyl group, and we call the alcove containing the origin the fundamental alcove. Occasionally we will need to consider p^e-facets; these are defined as above but with all occurrences of p replaced by p^e (and hence are associated to the action of W_{p^e}).

We will also need to consider the set of dominant weights

$$X^+ = \{\lambda\in X(T) \mid 0 < \langle\lambda+\rho,\alpha^\vee\rangle \text{ for all } \alpha\in S\}$$

and its subset of p-restricted weights $X_1(T)$ defined by imposing the additional constraint that $\langle\lambda+\rho,\alpha^\vee\rangle \leq p$ for all $\alpha\in S$. Any weight λ can be uniquely written in the form $\lambda = \lambda' + p\lambda''$ with $\lambda'\in X_1(T)$, and *any decomposition of a weight in this way will be assumed to be of this form.*

Given a Borel $T\subset B\subset G$ we can define the modules $\nabla(\lambda) = \mathrm{ind}_B^G k_\lambda$ where k_λ is the one-dimensional B-module of weight λ. By choosing B appropriately, we may arrange that $\nabla(\lambda)$ is non–zero precisely when λ is dominant. There is a contravariant duality on G, and the Weyl module $\Delta(\lambda)$ is the contravariant dual of $\nabla(\lambda)$. Both $\nabla(\lambda)$ and $\Delta(\lambda)$ has the same character, given by Weyl's character formula. Clearly we have $\mathrm{Hom}(\Delta(\mu),\Delta(\lambda)) \cong \mathrm{Hom}(\nabla(\lambda),\nabla(\mu))$, and hence for the remainder of this paper we will consider only induced rather than Weyl modules. (Note that throughout we will abuse notation and write Hom for Hom_G.)

For each $\lambda\in X^+$, the module $\nabla(\lambda)$ has simple socle $L(\lambda)$, and all simple modules arise in this manner. All other composition factors $L(\mu)$ of $\nabla(\lambda)$ satisfy $\mu<\lambda$ in the dominance order determined by R^+, and the strong linkage principle implies that $\mu\in W_p.\lambda$. Thus $\mathrm{Hom}(\nabla(\lambda),\nabla(\mu))$ is non–zero only if $\mu\leq\lambda$ and $\mu\in W_p.\lambda$. Note that for a given weight λ, the highest weight of any composition factor of $\nabla(\lambda)$ can be specified by determining the facet in which this highest weight lies. By Steinberg's tensor product Theorem, for any dominant weight λ we have an isomorphism $L(\lambda)\cong L(\lambda')\otimes L(\lambda'')^{\mathrm{F}}$, where F is the Frobenius morphism. In the special case $\lambda = (p-1)\rho$ we have that $\nabla(\lambda)\cong L(\lambda)$, and we call this the Steinberg module St. All simple modules are contravariantly self-dual.

Given weights λ,μ in the closure of some alcove F, there is a unique dominant weight ν in $W(\mu-\lambda)$. We define the translation functor T_λ^μ from λ to μ on a module V by $T_\lambda^\mu V = \mathrm{pr}_\mu(L(\nu)\otimes \mathrm{pr}_\lambda V)$, where $\mathrm{pr}_\tau V$ is the largest submodule of V all of whose composition factors have highest weights in $W_p.\tau$. We summarise the properties of translation functors that we shall require in the following proposition (proofs can be found in [17, II, Chapter 7]).

Proposition 2.1. *Let λ and μ be in the closure of some facet F. Then*

(i) *T_λ^μ is an exact functor, and is adjoint to T_μ^λ.*
(ii) *If λ and μ lie in the same facet then T_λ^μ is an equivalence of categories from the category of modules V with $\mathrm{pr}_\lambda V = V$ to the category of modules V with $\mathrm{pr}_\mu V = V$.*
(iii) *If μ is in the closure of the facet containing λ then $T_\lambda^\mu\nabla(\lambda)\cong\nabla(\mu)$.*
(iv) *More generally, $T_\lambda^\mu\nabla(\lambda)$ has a filtration $0 = T_0\subset\cdots\subset T_{t-1}\subset T_t = T_\lambda^\mu\nabla(\lambda)$ such that the factors T_i/T_{i-1} are all of the form $\nabla(w_i.\mu)$ where $w_i\in\mathrm{Stab}_{W_p}(\lambda)$ with $w_i.\mu\in X^+$, each such factor occurring exactly once. Further, this filtration may be taken such that $w_i.\mu < w_j.\mu$ implies that $i<j$.*

We say that a module V has a good filtration if it has a filtration with all factors of the form $\nabla(\lambda_i)$ with $\lambda_i \in X^+$, and denote by $(V : \nabla(\mu))$ the number of factors of the form $\nabla(\mu)$ in such a filtration. (This is independent of our choice of good filtration.) We call any module V such that both V and its dual have a good filtration a tilting module. By [8] there is for each dominant weight λ a unique indecomposable tilting module $T(\lambda)$ with that highest weight, and every indecomposable tilting module arises in this way.

Similarly we say that a module V has a good p-filtration if it has a filtration all of whose factors of the form $\nabla_p(\lambda) = \nabla(\lambda_i'')^{\mathrm{F}} \otimes L(\lambda_i')$ (recall our standing convention for such a decomposition of λ). In [17, II, Proposition 9.11] Jantzen gives a criterion for $\nabla(\lambda)$ to have a good p-filtration, while for $p \geq 2h-2$ Andersen [2, 3.7 Corollary] has shown that such good p-filtrations always exist. For our purposes it is enough to note [16, 3.13] that $\nabla(\lambda)$ with $\lambda \in X^+$, always has a good p-filtration in the case of $\mathrm{SL}_3(k)$. As these are the only type of p-filtrations which we will consider, we will usually omit the qualifier 'good'.

For any dominant weight λ we have that $\nabla(p\lambda + (p-1)\rho) \cong \nabla(\lambda)^{\mathrm{F}} \otimes \mathrm{St}$. Indeed the functor $V \longmapsto V^{\mathrm{F}} \otimes \mathrm{St}$ induces [17, II, 10.5(1)] an equivalence of categories from the category of all modules to the category of those modules all of whose composition factors are of the form $L(\lambda)$ with $\lambda' = (p-1)\rho$. More generally, the functor $V \longmapsto V^{\mathrm{F}} \otimes L(\lambda)$ where $\lambda \in X_1(T)$ induces an isomorphism of submodule lattices [12, Lemma 1.2].

We denote the (scheme theoretic) kernel of the Frobenius morphism by G_1. The simple modules for this subgroup are precisely the restrictions of the simple G-modules with p-restricted highest weights. As $\mathrm{Hom}_G(U,V) \cong (V \otimes U^*)^G \cong ((V \otimes U^*)^{G_1})^{G/G_1}$, and $G \cong G/G_1$ via F, we deduce that

$$\mathrm{Hom}(\nabla_p(\lambda), \nabla_p(\mu)) \cong \begin{cases} \mathrm{Hom}(\nabla(\lambda''), \nabla(\mu'')) & \text{if } \lambda' = \mu' \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

for all dominant weights λ and μ.

Any G-module W which is trivial as a G_1-module is of the form V^{F} for some G-module V, and we define $W^{(-1)} = V$. The (usual) dual of $\nabla(\lambda)$ is a Weyl module, which we denote by $\Delta(\lambda^*)$. Finally, if U, V, X, Y are finite dimensional G-modules then we have

$$\mathrm{Ext}^i_{G_1}(U \otimes X^{\mathrm{F}}, V \otimes Y^{\mathrm{F}}) \cong \mathrm{Ext}^i_{G_1}(U,V) \otimes (X^*)^{\mathrm{F}} \otimes Y^{\mathrm{F}} \tag{2}$$

(as a G-module) for all $i \geq 0$.

3. SL$_3$ data

In this section we will describe explicitly the good p-filtration for $\nabla(\lambda)$ in the case of $\mathrm{SL}_3(k)$. The detailed structure of these filtrations was given in [21, Theorem 4.12], and we summarise those results here. Further information about these filtrations also appears in [19, Chapter 2]. We will use this to give an explicit description of *p-good homomorphisms* (defined below) between induced modules, and recall a general result about homomorphisms due to Carter and Payne which will be needed in what follows.

Definition 3.1. A *p-good homomorphism* is a homomorphism ϕ between two modules which both have p-good filtrations, such that the image and kernel of ϕ both have good p-filtrations.

For $\mathrm{SL}_3(k)$ we shall denote the two simple roots by α_1 and α_2. A weight $\lambda \in X(T)$ can be identified with a pair of integers (a,b) via $\lambda = a\varpi_{\alpha_1} + b\varpi_{\alpha_2}$. With this convention we can identify X^+ with the subset $\mathbb{N}^2$ of $\mathbb{Z}^2$. The set of p-restricted weights $X_1(T)$ contains two alcoves; the fundamental alcove C_0 is the alcove containing the origin. Any alcove which is a translate of the fundamental alcove will be called a *down* alcove, while translates of the remaining p-restricted alcove will be called *up* alcoves. We extend our earlier notation and write $\lambda = p(a'', b'') + (a', b')$ where $(a', b') \in X_1(T)$. In order to describe the good p-filtration of $\nabla(\lambda)$, it is enough (by the strong linkage principle) to indicate the facets in which the highest weights of the factors in the p-filtration occur. By the translation principle (Proposition 2.1(ii)) the configuration of facets obtained depends only on the facet containing λ and not on the weight itself. Thus we will represent the factors $\nabla_p(\lambda_i)$ by numbers corresponding to an appropriate labelling of the facets in each case.

In the diagrams representing good p-filtrations, the factors of a filtration will be represented by labels with lines connecting them. Two factors will be joined if and only if there is a nontrivial extension (in $\nabla(\lambda)$) between them. The diagrams are oriented so that sections that embed in $\nabla(\lambda)$ occur at the bottom of the diagram. If λ_i'' is not dominant for some factor $\nabla_p(\lambda_i)$ occurring in one of our diagrams then we interpret this factor as the zero module (i.e. we ignore that part of the diagram).

We now list the various cases that can arise, according to [21, Theorem 4.12]. If a'' and $b'' \geq 1$ and (a', b') lies in the fundamental alcove then there are nine factors in the p-filtration as indicated in Figure 1a(i), and the structure of this filtration (when a'' and $b'' \not\equiv 0 \pmod p$) is given in Figure 1a(ii). If $a'' \equiv 0 \pmod p$ (respectively $b'' \equiv 0 \pmod p$) then there is an additional extension of $\nabla_p(\lambda_5)$ by $\nabla_p(\lambda_6)$ (respectively of $\nabla_p(\lambda_3)$ by $\nabla_p(\lambda_8)$). If both a'' and $b'' \equiv 0 \pmod p$ then the structure is given in Figure 1a(iii); the other cases are similar.

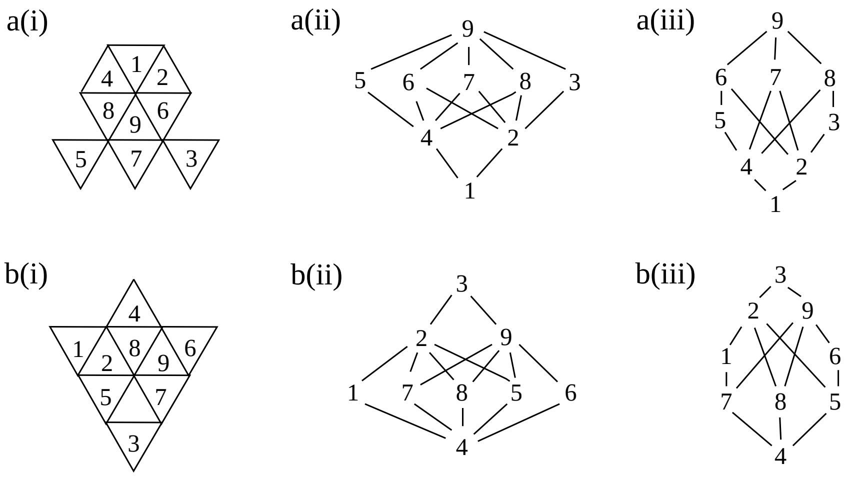

FIGURE 1.

If (a', b') lies in the interior of the other p-restricted alcove then there are also nine factors in the p-filtration, as indicated in Figure 1b(i), and the structure of the filtration again depends on a'' and b''. If both a'' and $b'' \not\equiv -1 \pmod p$ then it is given in Figure 1b(ii) while if $a'' \equiv -1 \pmod p$ (respectively if $b'' \equiv -1 \pmod p$) then there is an additional extension of $\nabla_p(\lambda_5)$ by $\nabla_p(\lambda_6)$ (respectively of $\nabla_p(\lambda_7)$ by $\nabla_p(\lambda_1)$). If both a'' and $b'' \equiv -1 \pmod p$ then the structure is illustrated in Figure 1b(iii); the other cases are similar.

There are two remaining alcove cases: where λ lies in a down alcove and either $a'' = 0$ or $b'' = 0$. As these cases are symmetric we only consider the latter. So suppose that $\lambda = p(a'', 0) + (a', b')$ with $(a', b') \in C_0$. Then $\nabla(\lambda)$ has three factors in the p-filtration, as indicated in Figure 2a(i) (where the shaded region represents the boundary of the dominant region), and the structure of the filtration is given in Figure 2a(ii). When we need to distinguish between the two down alcove cases we shall refer to those as in Figure 1a(i) as *internal* down alcoves, and those as in Figure 2a(i) (and the symmetric version thereof) as *just dominant* down alcoves. We shall also refer to the up alcoves immediately above just dominant down alcoves as just dominant, and similarly the walls between two just dominant alcoves. All other walls will be referred to as interior walls. To distinguish further the just dominant down alcove cases we shall refer to that illustrated in Figure 2a(i) as the *right-hand* case, and to the symmetric version (which we number in the symmetric fashion) as the *left-hand* case.

Next we consider the various wall cases. If $a' + b' = p - 2$ then λ lies on one of the horizontal walls in our diagrams. There are now four factors in our filtration, as indicated in Figure 2b(i), and the structure of the filtration is given in Figure 2b(ii). The remaining two wall cases are symmetric,

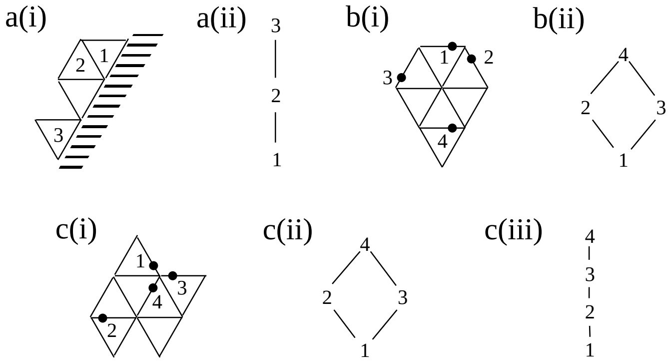

FIGURE 2.

so we only consider the case where $\lambda = p(a'', b'') + (p-1, b')$. We refer to this case as the *left-hand diagonal wall*. Once again there are four factors in the filtration as indicated in Figure 2c(i). If $a'' \not\equiv -1 \pmod p$ then the structure is as in Figure 2c(ii), otherwise it is as in Figure 2c(iii). For right-hand walls we will use the numbering of the four factors given by reflecting the diagram (so that the top and bottom factors in the filtration are also labelled by 4 and 1 respectively). The only remaining case is when λ lies on a vertex, i.e. $\lambda' = (p-1)\rho$. But in this case there is only a single factor in the p-filtration, as noted in the previous section.

In what follows we will frequently make use of

Lemma 3.2. *Suppose that $G = \mathrm{SL}_3(k)$. For all $\lambda \in X^+$ the module $\nabla_p(\lambda)$ has a simple head.*

Proof: Jantzen has shown [16, 6.9] that $\nabla(\mu)$ has simple head for all dominant weights μ, provided that $p > 3$, and in [21, Proposition 4.11] the restriction on p has been removed. Now the result follows from [12, Lemma 1.2]. □

Our goal is to determine $\mathrm{Hom}(\nabla(\lambda), \nabla(\mu))$ for all pairs of dominant weights λ and μ. In order to construct homomorphisms we will need the following result of Carter and Payne.

Theorem 3.3. *Suppose that $G = \mathrm{SL}_n(k)$ and that $\lambda, \mu \in X^+$ satisfy the following conditions:*

(i) $\mu < \lambda$.
(ii) *There exists some integer $e > 0$ such that:*
  (a) *λ and μ are mirror images in some p^e-wall L and*
  (b) *L is the unique p^e-wall between λ and μ (possibly containing λ or μ) parallel to L.*

Then $\mathrm{Hom}(\nabla(\lambda), \nabla(\mu)) \neq 0$.

Proof: See [4]. □

When we refer to such a homomorphism from $\nabla(\lambda)$ to $\nabla(\mu)$ as reflection about a p^e-wall, we shall assume that e is chosen to be minimal with that property. For SL_3 we shall refer to the three cases as reflection to the left (L), right (R), and below (B), depending on the relative positions of λ and μ.

Theorem 3.3 allows us to construct homomorphisms between widely separated weights, but provides no information as to their structure. For $\mathrm{SL}_3(k)$ we will be able to provide this extra information using the explicit description of p-good morphisms given in [21]. Thus we conclude this section by reviewing this data.

In the following lemma we describe all non-zero p-good homomorphisms between induced modules. We will consider each of the cases illustrated in Figures 1 and 2 in turn (the remaining cases are symmetric). In each case we will describe homomorphisms from $\nabla(\lambda)$ *in terms of the labelling of facets given in the figure corresponding to λ*. The image of such a homomorphism $\nabla(\lambda) \to \nabla(\lambda_i)$ will then be isomorphic to the smallest quotient of the p-filtration of $\nabla(\lambda)$ containing $\nabla_p(\lambda_i)$, which can be read off the corresponding diagrams in Figures 1 and 2.

Lemma 3.4. *For a dominant weight λ there exists a non-zero p-good homomorphism $\nabla(\lambda) \to \nabla(\lambda_i)$ if and only if λ_i is dominant and the pair (λ, i) occurs in the following table (where we identify λ with the type of facet in which it occurs):*

λ	i
Fig 1a: internal down alcove	*1, 2, 4, 6, 8 and 9*
Fig 1b: up alcove	*1, 2, 3, 4, 6, 8 and 9*
Fig 2a: just dominant down alcove	*1, 2 and 3*
Fig 2b: horizontal wall	*1, 2, 3 and 4*
Fig 2c: interior diagonal wall	*1, 3 and 4*
Fig 2c: just dominant diagonal wall	*1 and 2*

Proof: Argue as in the proof of [21, Lemma 5.1], or [17, II, 9.14 Remark 4]. □

Note that when all terms in the p-filtration of $\nabla(\lambda)$ are simple, the above Lemma gives all weights μ such that $\mathrm{Hom}(\nabla(\lambda), \nabla(\mu)) \neq 0$. All but four of the p-good homomorphisms above can be constructed as composites of Carter-Payne maps. The four exceptions are the map from $\nabla(\lambda)$ to $\nabla(\lambda_3)$ in the interior up alcove case, the map from $\nabla(\lambda)$ to $\nabla(\lambda_4)$ in the interior horizontal wall case, the map from $\nabla(\lambda)$ to $\nabla(\lambda_2)$ in the just dominant diagonal wall case, and the map from $\nabla(\lambda)$ to $\nabla(\lambda_3)$ in the just dominant down alcove case. We will call these latter two maps the *exceptional* p-good maps. (It can happen that these exceptional maps can be constructed as the composite of Carter-Payne maps, as in second example in section 10. This is non-generic behaviour. Generically, the exceptional p-good maps will not be the composite of Carter-Payne maps.) Note that all of these p-good maps which are generically composites of Carter-Payne maps are the maps with image the G_1-head of $\nabla(\lambda)$.

4. Overall strategy

As the remainder of this paper involves some rather intricate calculations, we devote this section to an overview of the strategy behind the proof. From this point on we will assume that $G = \mathrm{SL}_3(k)$ and that $p > 2$ (so that all our facets are non-empty).

Given a dominant weight λ, we assume that all homomorphisms from $\nabla(\tau)$ with $\tau < \lambda$ have been classified, and that all the corresponding Hom-spaces are one-dimensional. First we claim that $\dim \mathrm{Hom}(\nabla(\lambda), \nabla(\mu))$ is at most one dimensional for all weights $\mu \in X^+$. Second we claim that we can find all weights $\mu < \lambda$ such that $\mathrm{Hom}(\nabla(\lambda), \nabla(\mu))$ is non-zero. By induction this will calculate all possible Hom-spaces for SL_3 with $p > 2$.

The basic idea is to use translation functors. For weights on a vertex, the result is known by induction using (1), or the equivalence of categories discussed just before that equation. Next suppose that λ lies in the closure of a down alcove whose lowest vertex ν is dominant. By induction the Hom-spaces for $\nabla(\nu)$ are known. Let

$$\Theta_\nu = \{\theta \in X^+ \mid \mathrm{Hom}(\nabla(\nu), \nabla(\theta)) \neq 0\}.$$

We 'translate' all the weights in Θ_ν to a new set Γ_λ with

$$\Gamma_\lambda = \{\gamma \in X^+ \mid (T_\nu^\lambda \nabla(\theta) : \nabla(\gamma)) \neq 0 \text{ for some } \theta \in \Theta_\nu\}. \tag{3}$$

Now suppose λ lies in an up alcove, and choose τ on the horizontal wall below it. If τ lies in the closure of a down alcove with lower vertex ν, and ν is dominant, then we further translate the weights in Γ_τ to form

$$\Gamma_\lambda = \{\gamma \in X^+ \mid (T_\tau^\lambda \nabla(\theta) : \nabla(\gamma)) \neq 0 \text{ for some } \theta \in \Gamma_\tau\}. \tag{4}$$

In both cases (3) and (4) we now remove all weights in Γ_λ which are not composition factors of $\nabla(\lambda)$. We would like to claim that the resulting set

$$\Upsilon_\lambda = \{\gamma \in \Gamma_\lambda \mid [\nabla(\lambda) : L(\gamma)] \neq 0\}$$

is precisely the set of weights η for which $\mathrm{Hom}(\nabla(\lambda), \nabla(\eta)) \cong k$, and that if $\eta \notin \Upsilon_\lambda$ then $\mathrm{Hom}(\nabla(\lambda), \nabla(\eta)) \cong 0$. However, this is not quite the case and will need to be modified.

In the course of showing that all Hom-spaces are at most one-dimensional (Theorem 5.1) we easily show that Γ_λ is an upper bound for the set of weights for which non-zero homomorphisms can exist, and it is obvious that this can then be refined to give Υ_λ as an upper bound.

We already have one family of homomorphisms: those constructed by Carter and Payne. There is another obvious class of maps, which we will now describe. Suppose that $\nabla_p(\tau)$ is the top term in a p-filtration of some $\nabla(\lambda)$. Given any map from $\nabla(\tau'')$ to $\nabla(\mu)$ (which we will know by induction) we obtain a map from $\nabla_p(\tau)$ to $\nabla_p(\tau' + p\mu)$, and hence from $\nabla(\lambda)$ to $\nabla(\tau' + p\mu)$ (by killing all other terms in the p-filtration of $\nabla(\lambda)$). We will refer to maps arising in this way as *twisted maps*. (Of course maps can be both twisted maps and a composite of Carter-Payne maps.)

We would like to claim that all homomorphisms are either composites of Carter-Payne maps, or twisted maps. However, this too is not quite correct. To see why, we must consider those dominant λ not covered by the cases above. First suppose that λ lies in the closure of a just dominant down alcove. Although we cannot apply the translation functor approach outlined above, it is in fact possible to determine all homomorphisms directly (Theorem 8.3), using the explicit description of the structure of symmetric powers given by Doty [10]. (For λ in an up alcove above a just dominant down alcove, the determination of homomorphisms is now straightforward by translation arguments.)

We now return to the sets Υ_λ. If ν and θ (as used in the definition of these sets) are related by a composite of Carter-Payne maps then for each weight in Υ_λ we can either (i) explicitly construct all possible homomorphisms using composites of Carter-Payne maps or twisted maps, or (ii) show no map exists. The latter case is rare: for 'generic' weights case (ii) never occurs for λ in a down alcove, and only in very special configurations for λ in an up alcove. For the precise statement for down alcoves see Theorem 7.2, and for up alcoves Theorem 7.3.

If ν and θ are instead related by a twisted map, then in most cases the analogues of Theorems 7.2 and 7.3 are straightforward, as every weight in Υ_λ can be obtained by a twisted map. However, when (for example) $\nu = (p-1)\rho + p\lambda$ with λ close to the boundary of the dominant region, there are additional weights in Υ_λ which can be reached neither by composites of Carter-Payne maps nor twisted maps. Unfortunately in these cases maps do exist, and Section 9 of the paper is devoted to the construction of these *exceptional maps*.

We will show that there are a pair of terms at the top of the p-filtration for $\nabla(\lambda)$ which individually map to a pair of terms at the bottom of the p-filtration for $\nabla(\mu)$. By considering pullbacks and pushouts we will show that this pair of maps can be 'glued together' to give the required exceptional map. This will complete our classification of homomorphisms.

5. Dimensions of Hom-spaces

In this section we begin an inductive procedure that will continue for the rest of the paper. We show that all Hom-spaces between induced modules are at most one-dimensional, provided that there are no homomorphisms between certain special pairs of induced modules labelled by smaller weights. The verification of this hypothesis will follow from the remainder of the paper.

We say that two weights λ, μ are in the *same* ∇_p*-class* if $\lambda' = \mu'$ and λ'' and μ'' are in the same G-block. Any weight lies in a unique translate of the set of p-restricted weights; note that λ and μ can only be in the same ∇_p-class if they both lie in the same position in the respective translates (and also have the Steinberg weights immediately above each translate lying in the same W_{p^2}-orbit). Clearly if λ and μ are in distinct ∇_p-classes, then $\nabla_p(\lambda)$ and $\nabla_p(\mu)$ have no common composition factors. We will also say that two modules are in the same ∇_p-class if the highest weights of their composition factors are all in the same ∇_p-class.

Theorem 5.1. *For all $\lambda, \mu \in X^+$ we have*

$$\dim \operatorname{Hom}(\nabla(\lambda), \nabla(\mu)) \leq 1.$$

Further, for those λ for which the set Υ_λ from Section 4 has been defined we have

$$\operatorname{Hom}(\nabla(\lambda), \nabla(\mu)) = 0$$

for all $\mu \notin \Upsilon_\lambda$.

Proof: We proceed by induction on λ, and may assume that $\mu \in W.\lambda$ with $\mu \leq \lambda$. If λ is in the fundamental alcove then $\nabla(\lambda)$ is simple and so we are done by [17, II 2.8 Proposition]. Now suppose that λ is in the closure of a just dominant down alcove (i.e there *does not* exist a dominant weight θ on the vertex immediately below this). By symmetry, it is enough to consider the case where λ is near to the right hand boundary of the dominant region.

We first consider the two wall cases. By Figure 2b or c we have a short exact sequence

$$0 \to \nabla_p(\lambda_1) \to \nabla(\lambda) \to \nabla_p(\lambda_a) \to 0 \tag{5}$$

where $a = 3$ in case (b) and $a = 2$ in case (c). (Recall that two of the four possible terms in the p-filtration are zero when λ is just dominant.) Hence there exists a long exact sequence

$$0 \to \mathrm{Hom}(\nabla_p(\lambda_a), \nabla(\mu)) \to \mathrm{Hom}(\nabla(\lambda), \nabla(\mu)) \to \mathrm{Hom}(\nabla_p(\lambda_1), \nabla(\mu)) \to \cdots .$$

It will be enough to show that the first and third Hom-spaces $\mathrm{Hom}(\nabla_p(\lambda_i), \nabla(\mu))$ are each at most one dimensional, as they can never both be non-zero as λ_1 and λ_a are in different ∇_p-classes.

We begin by showing that the image in either case of any non-zero map ϕ from $\nabla_p(\lambda_i)$ to $\nabla(\mu)$ must be a submodule of $\nabla_p(\mu)$. Clearly $\nabla_p(\lambda_i)$ and $\nabla_p(\mu)$ must be in the same ∇_p-class, as the socle of $\nabla(\mu)$ must be a common composition factor. By Figure 2b or c, the only other term in the p-filtration of $\nabla_p(\mu)$ in this ∇_p-class occurs when μ is as in case (b), in which case it is the factor $\nabla_p(\mu_4)$. (Recall that to be in the same ∇_p-class two weights must lie in the same position in their respective translates of the set of p-restricted weights.) As all composition factors of $\mathrm{im}\,\phi$ lie in the same ∇_p-class, it is enough to show that none of the composition factors in $\nabla_p(\mu_4)$ occur in this image. However, the extension of $\nabla_p(\mu_2)$ by $\nabla_p(\mu_4)$ has a simple socle, as this extension is the image under the map $\nabla(\mu_1) \to \nabla(\mu_2)$. Therefore if any composition factor of $\nabla_p(\mu_4)$ does occur in $\mathrm{im}\,\phi$ then so must $\mathrm{soc}\,\nabla_p(\mu_2)$, which is impossible.

Thus we have shown that $\mathrm{im}\,\phi \leq \nabla_p(\mu)$. Now (1), and the induction hypothesis, immediately implies that the desired Hom-spaces are at most one dimensional, and cannot both be non-zero.

Next suppose that λ is in the interior of a just dominant down alcove, (and hence that there exists a weight θ on the diagonal wall below λ). We have a short exact sequence

$$0 \to \nabla(\nu) \to T_\theta^\lambda \nabla(\theta) \to \nabla(\lambda) \to 0 \tag{6}$$

for some ν. This gives rise to an exact sequence

$$0 \to \mathrm{Hom}(\nabla(\lambda), \nabla(\mu)) \to \mathrm{Hom}(T_\theta^\lambda \nabla(\theta), \nabla(\mu)) \tag{7}$$

and it is enough to show that this final Hom-space is at most one-dimensional. However

$$\mathrm{Hom}(T_\theta^\lambda \nabla(\theta), \nabla(\mu)) \cong \mathrm{Hom}(\nabla(\theta), T_\lambda^\theta \nabla(\mu)) \cong \mathrm{Hom}(\nabla(\theta), \nabla(\tau))$$

for some τ (by Proposition 2.1(iii)), and we are done by induction.

We next consider the case where λ is in the lower closure of an internal down alcove (so there exists a dominant weight θ on the vertex immediately below this). If $\lambda = \theta$ then $\nabla(\lambda) \cong \nabla(\lambda'')^{\mathrm{F}} \otimes \mathrm{St}$ and block considerations show that to have a non-zero homomorphism we must have $\nabla(\mu) \cong \nabla(\mu'')^{\mathrm{F}} \otimes \mathrm{St}$. Now by (1) we have that

$$\mathrm{Hom}(\nabla(\lambda'')^{\mathrm{F}} \otimes \mathrm{St}, \nabla(\mu'')^{\mathrm{F}} \otimes \mathrm{St}) \cong \mathrm{Hom}(\nabla(\lambda''), \nabla(\mu''))$$

and we are done by induction.

Now suppose that λ is in the interior of an internal down alcove. By Proposition 2.1 we have that $T_\theta^\lambda \nabla(\theta)$ has six factors of the form $\nabla(\alpha)$, and there is a short exact sequence

$$0 \to X \to T_\theta^\lambda \nabla(\theta) \to \nabla(\lambda) \to 0$$

where X has a good filtration. Hence we have an exact sequence as in (7), and the argument follows exactly as above. For later use we note that (in this case) if the final Hom-space in (7) is non-zero then μ must have been in one of the six alcoves adjacent to τ. If λ lies on one of the two walls then the argument is similar, but there are only three weights μ which could give rise to the weight τ.

When λ is on the wall in the upper closure of an internal down alcove the proof is a little more complicated, so we shall first show how the up alcove case follows from it. Thus we suppose that λ lies in the interior of an up alcove and let θ be a weight on the wall immediately below λ. We again have a short exact sequence as in (6), and the argument proceeds as in that case.

Finally we consider the case when λ is on the wall in the upper closure of an internal down alcove. We have a short exact sequence

$$0 \to M_1 \to \nabla(\lambda) \to M_2 \to 0$$

where M_1 and M_2 are defined by the sequences

$$0 \to \nabla_p(\lambda_1) \to M_1 \to \nabla_p(\lambda_2) \to 0 \quad \text{and} \quad 0 \to \nabla_p(\lambda_3) \to M_2 \to \nabla_p(\lambda_4) \to 0.$$

Arguing as for the sequence (5), we see that $\dim \mathrm{Hom}(M_i, \nabla(\mu)) \leq 1$ for $i = 1, 2$, and thus from the sequence

$$0 \to \mathrm{Hom}(M_2, \nabla(\mu)) \to \mathrm{Hom}(\nabla(\lambda), \nabla(\mu)) \to \mathrm{Hom}(M_1, \nabla(\mu))$$

we see that we are done unless both these Hom-spaces are non-zero.

Thus we may assume that $\mathrm{Hom}(M_i, \nabla(\mu)) \neq 0$ for $i = 1, 2$. When $i = 1$ this implies that $\mu' = \lambda_1'$ or λ_2', and for $i = 2$ that $\mu' = \lambda_3'$ or λ_4'. Therefore we must have $\mu' = \lambda_1' = \lambda_4'$, and so μ lies on a horizontal wall. Let θ be a weight lying in the interior of the alcove immediately below λ. We have that

$$\mathrm{Hom}(\nabla(\lambda), \nabla(\mu)) \cong \mathrm{Hom}(T_\theta^\lambda \nabla(\theta), \nabla(\mu)) \cong \mathrm{Hom}(\nabla(\theta), T_\lambda^\theta \nabla(\mu))$$

and an exact sequence

$$0 \to \mathrm{Hom}(\nabla(\theta), \nabla(\mu_l)) \to \mathrm{Hom}(\nabla(\theta), T_\lambda^\theta \nabla(\mu)) \to \mathrm{Hom}(\nabla(\theta), \nabla(\mu_u))$$

where μ_l (respectively μ_u) in $W.\theta$ lies in the lower (respectively upper) alcove adjacent to μ. By induction the outer Hom-spaces in this sequence are at most one-dimensional, so it is enough to prove that they cannot both be non-zero.

As $\theta' = \mu_l'$ (as μ is in the same block of λ and they are both on horizontal walls), we will be done if we can show

Assumption 5.2. Suppose that $\gamma \in X^+$ is in an internal down alcove and $\tau \in X^+$ lies on the horizontal wall above γ. Let $w \in W_p$ with $w \neq 1$. If $\nu = w.\gamma$ is in a dominant down alcove such that $w.\tau$ lies on the horizontal wall above ν then $\mathrm{Hom}(\nabla(\gamma), \nabla(\nu)) = 0$.

Assumption 5.2 will follow from the results in Sections 7–9, which will thus complete the proof of Theorem 5.1. (Note that the second part of Theorem 5.1 is clear from the translation arguments we have used.) Consequently, in what follows we may only apply Theorem 5.1 to weights that are strictly smaller than λ. □

6. Composites of Carter-Payne maps

In this section we give a recursive description of the homomorphisms constructed by Carter and Payne (for SL_3). Using this we can determine inductively which composites of such maps are non-zero.

Lemma 6.1. *All factors in the p-filtration of $\nabla(\lambda)$ are in distinct ∇_p-classes if and only if we are in one of the following situations:*

- (i) *λ lies in a down alcove and either the alcove is just dominant or both*
  - (a) $\langle \lambda'' + \rho, \check{\alpha}_i \rangle \not\equiv 1 \pmod p$ *for $i = 1, 2$.*
  - (b) $\langle \lambda'' + \rho, \check{\alpha}_1 + \check{\alpha}_2 \rangle \not\equiv 1 \pmod p$.
- (ii) *λ lies in an up alcove and either the alcove is just dominant or both*
  - (a) $\langle \lambda'' + \rho, \check{\alpha}_i \rangle \not\equiv 0 \pmod p$ *for $i = 1, 2$.*
  - (b) $\langle \lambda'' + \rho, \check{\alpha}_1 + \check{\alpha}_2 \rangle \not\equiv 1 \pmod p$.

(iii) *λ lies on a left-hand diagonal wall and either is just dominant or* $\langle \lambda''+\rho, \check{\alpha}_1\rangle \not\equiv 0 \pmod p$.
(iv) *λ lies on a right-hand diagonal wall and either is just dominant or* $\langle \lambda''+\rho, \check{\alpha}_2\rangle \not\equiv 0 \pmod p$.
(v) *λ lies on a horizontal wall and either is just dominant or* $\langle \lambda''+\rho, \check{\alpha}_1+\check{\alpha}_2\rangle \not\equiv 1 \pmod p$.
(vi) *λ lies on a vertex.*

Proof: This is an elementary calculation using the description of p-filtrations given in the previous section if $p \geq 5$. If $p=3$ a little more care is needed for the up and down alcove case as there are only two p-restricted weights which lie inside alcoves. But looking at the λ'' that can occur it is clear that we get the same result as for $p \geq 5$ as the λ'' are mostly too close together to be in the same G-block. □

We will say that λ is *generic* if it satisfies the conditions of this lemma. We will say that λ is *sufficiently generic* unless either

(i) λ lies on a diagonal wall and is not generic, or
(ii) λ lies in a down alcove with $\langle \lambda''+\rho, \check{\alpha}_i\rangle \equiv 1 \pmod p$ for $i=1$ or 2.

We will say that λ is *recursively generic* if λ is sufficiently generic and either all terms in the p-filtration of $\nabla(\lambda)$ are simple, or all such terms are of the form $\nabla_p(\lambda_i)$ with λ_i'' recursively generic.

When λ is non-generic, we will have to consider in detail the structure of extensions between factors in the p-filtration in the same ∇_p-class. The basic properties of these extensions are summarised in the following lemma.

Lemma 6.2. *Suppose that $\nabla_p(\lambda_i)$ and $\nabla_p(\lambda_j)$ are two factors in the p-filtration of $\nabla(\lambda)$ in the same ∇_p-class with $\lambda_i > \lambda_j$ and there is a non-split extension appearing between them in $\nabla(\lambda)$. Then*

$$\mathrm{Ext}^1_G(\nabla_p(\lambda_i), \nabla_p(\lambda_j)) \cong k$$

and the non-split extension is isomorphic to $(T_\mu^{\lambda_i''}\nabla(\mu))^{\mathrm{F}} \otimes L(\lambda_i')$ (where μ lies on the wall between λ_i'' and λ_j'') and has simple socle $L(\lambda_j)$. This extension also has simple head provided λ_i'' is not just dominant.

Proof: By considering the various cases that can arise, it is easy to verify that λ_i'' and λ_j'' are related by a single left or right reflection. By [21, Lemma 4.2] we have

$$\mathrm{Ext}^1_G(\nabla_p(\lambda_i), \nabla_p(\lambda_j)) \cong \mathrm{Ext}^1_G(\nabla(\lambda_i''), \nabla(\lambda_j''))$$

and this latter Ext-group is k using the corresponding $\mathrm{SL}_2(k)$ result [11, Theorem 3.6 and Corollary 4.3].

Now consider the translate $T_\mu^{\lambda_i''}\nabla(\mu)$ where μ lies on the wall between λ_i'' and λ_j''. Without loss of generality we can pick μ so that $T_\mu^{\lambda_i''} = \mathrm{pr}_{\lambda_i''}(-\otimes E)$ or $T_\mu^{\lambda_i''} = \mathrm{pr}_{\lambda_i''}(-\otimes E^*)$ where $E = \nabla(1,0)$, the natural representation and depending on whether λ_i'' and λ_j'' are related by a right hand or left hand reflection, respectively. We want to check that this translate $T_\mu^{\lambda_i''}\nabla(\mu)$ has simple head and socle. Firstly this translate has short exact sequence

$$0 \to \nabla(\lambda_j'') \to T_\mu^{\lambda_i''}\nabla(\mu) \to \nabla(\lambda_i'') \to 0.$$

So its socle is contained in $L(\lambda_i'') \oplus L(\lambda_j'')$ and its head is contained in $\mathrm{hd}(\nabla(\lambda_i'')) \oplus \mathrm{hd}(\nabla(\lambda_j''))$ where hd denotes the head of a module. Thus it is enough to show that one of each of these respective simple modules cannot be in the socle (or head).

If λ_i'' and λ_j'' are in adjacent alcoves then μ is in the closure of the facet containing λ_i'' so then this translate has simple socle $L(\lambda_j'')$ and simple head by [17, II, 7.19 Proposition (b)].

We now suppose that λ_i'' and λ_j'' are related by a right hand reflection and lie on walls as shown in Figure 3. The argument for the other cases are similar.

Now

$$\mathrm{Hom}(L(\lambda_i''), T_\mu^{\lambda_i''}\nabla(\mu)) \cong \mathrm{Hom}(L(\lambda_i''), \nabla(\mu)\otimes E) \cong \mathrm{Hom}(L(\lambda_i'')\otimes E^*, \nabla(\mu))$$

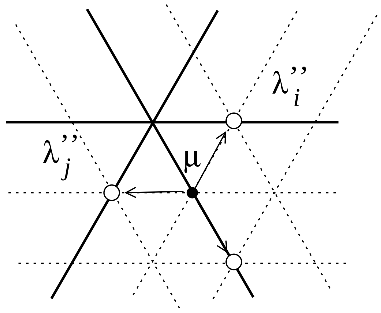

FIGURE 3.

and $\lambda_i'' = p\lambda_i'''' + (0, p-2)$. So $L(\lambda_i'') \otimes E^*$ has character $\mathrm{ch}(L(\lambda_i'''')^{\mathrm{F}} \otimes \nabla(0, p-1)) + \mathrm{ch}(L(\lambda_i'''')^{\mathrm{F}} \otimes \nabla(1, p-3))$ and neither of these factors has G_1-type the same as that of $L(\mu)$, which is $(p-1, p-2)$. Thus the last Hom space must be zero and the socle of $T_\mu^{\lambda_i''}\nabla(\mu)$ is simple. This implies that the translate must be indecomposable and hence by uniqueness and [12, Lemma 1.2], our desired extension is isomorphic to $\left(T_\mu^{\lambda_i''}\nabla(\mu)\right)^{\mathrm{F}} \otimes L(\lambda_j'')$ and has simple socle $L(\lambda_j)$.

Similarly for the head (when μ is an internal weight) we have

$$\mathrm{Hom}(T_\mu^{\lambda_i''}\nabla(\mu), \mathrm{hd}(\nabla(\lambda_j''))) \cong \mathrm{Hom}(\nabla(\mu) \otimes E, \mathrm{hd}(\nabla(\lambda_j''))) \cong \mathrm{Hom}(\nabla(\mu), \mathrm{hd}(\nabla(\lambda_j'')) \otimes E^*)$$

and $\mathrm{hd}(\nabla(\lambda_j'')) = L(\eta) = L(p\eta'' + (p-1, 0)$. So $L(\eta) \otimes E^*$ has character $\mathrm{ch}(L(\eta'')^{\mathrm{F}} \otimes \nabla(p-1, 1)) + \mathrm{ch}(L(\eta'')^{\mathrm{F}} \otimes \nabla(p-2, 0))$ and neither of these factors has G_1-type the same as that of $\mathrm{hd}(\nabla((\mu)))$, which is $(0, p-1)$. Thus the last Hom space must be zero and the head of $T_\mu^{\lambda_i''}\nabla(\mu)$ is simple.

If μ is a just dominant weight then the head of $\nabla(\mu)$ has G_1-type $(p-2, 0)$ so our argument fails. In fact in this case the translate has non-simple head so the condition on λ_i'' in the lemma is necessary. □

If λ and μ satisfy the conditions of Theorem 3.3 for some $e > 0$, with μ to the left (respectively to the right, below) λ we denote the corresponding Carter-Payne map ϕ by ϕ_L^e (respectively ϕ_R^e, ϕ_B^e). To each of these maps there is a corresponding local (i.e. $e = 1$) map starting at λ (which if λ is on a wall may by the identity map). These local maps are p-good homomorphisms, and we refer to the set of terms in the p-filtration of $\nabla(\lambda)$ which survive (necessarily completely) under such a map ϕ as *the p-factors (of $\nabla(\lambda)$) associated to ϕ*, unless both λ lies in an up alcove or on a horizontal wall, and $\phi = \phi_B^{>1}$. In these remaining cases we refer only to the *top term* in the p-filtration of $\nabla(\lambda)$ as such a p-factor.

Proposition 6.3. *Suppose that λ and μ satisfy the conditions of Theorem 3.3 for some $e > 0$. In all cases the Carter-Payne map ϕ is non-zero only on the p-factors associated to ϕ, and on such a factor $\nabla_p(\lambda_i)$ it is induced via twisting from the corresponding Carter-Payne map from $\nabla(\lambda_i'')$ about an $(e-1)$-wall in all but the following cases:*

(i) *λ in a down alcove with λ_3 and λ_8 in the same ∇_p-class, $\phi = \phi_R$, and $i = 3$ or 8.*
(ii) *λ in a down alcove with λ_5 and λ_6 in the same ∇_p-class, $\phi = \phi_L$, and $i = 5$ or 6.*
(iii) *λ on a LH diagonal wall with λ_2 and λ_3 in the same ∇_p-class, $\phi = \phi_L$, and $i = 2$ or 3.*
(iv) *λ on a RH diagonal wall with λ_2 and λ_3 in the same ∇_p-class, $\phi = \phi_R$, and $i = 2$ or 3.*

In each of these four cases the map induced on $\nabla_p(\lambda_i)$ is non-zero, but does not come from the twist of the corresponding Carter-Payne map.

Thus if λ is recursively generic then the composition factors occurring in the image of $\nabla(\lambda)$ under the Carter-Payne homomorphism ϕ can be described inductively using the local data from the preceding section. In the remaining cases the above description will be sufficient for our purposes, so we will not analyse the exceptional cases further.

Proof: We proceed by induction on e. When $e = 1$ we are done by the results in the preceding section. Now suppose that $e > 1$. In what follows we shall assume when presenting lattices in $\nabla(\mu)$ that μ is not too close to the edge of the dominant region. The reader may verify that the modifications necessary in the remaining cases do not affect the form of the answer.

Case (i): Down alcoves. We will begin by consider the case when λ is as in Figure 1a(ii). There are three subcases, corresponding to reflections about a wall to the left, right or below the alcove containing λ. We consider a reflection to a weight μ to the right of λ (i.e. corresponding when $e = 1$ to the map $\nabla(1) \to \nabla(2)$) and denote the corresponding morphism (which will be unique up to scalars once we know the Hom space is one-dimensional) by ϕ. The case of a reflection to the left is similar, while that for a reflection below is even simpler; both are left to the reader.

In the case under consideration μ lies in an up alcove, and corresponds to the situation in Figure 1b. We will renumber the alcoves in this latter diagram to be consistent with the numbering for λ. That is, we will renumber the alcoves so that alcoves with the same number in each diagram are in the same ∇_p-class. Note that as $e > 1$ this numbering is different from that in Figure 1. If λ is generic then this is unambiguous, and we will use the same labelling in the non-generic case. The new labelling is illustrated in Figure 4, together with the corresponding lattices. To avoid confusion we will refer to a weight in $W.\lambda$ in an alcove labelled i in the diagram for λ as λ_i and in the diagram for μ as $\mu[\lambda_i]$.

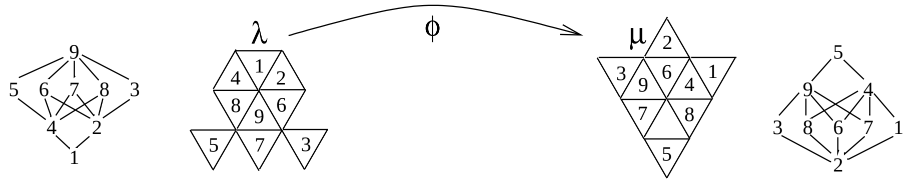

FIGURE 4.

If λ is non-generic then the only terms in the p-filtration which can lie in the same ∇_p-class are λ_1 and λ_7, or λ_3 and λ_8, or λ_5 and λ_6. As we are considering the case in Figure 1a(ii), λ_3 and λ_8 are in distinct ∇_p-classes. Also, for any λ, the terms in the p-filtration labelled by λ_1, λ_4 and λ_5 must lie in the kernel of ϕ as none of their composition factors lie above the socle of $\nabla_p(\lambda_2)$ (by figure 1a) and some composition factor of $\nabla_p(\lambda_2)$ maps to the socle of $\nabla(\mu)$ (as this is the only factor in the same p-class as the lowest term in the filtration of $\nabla(\mu)$).

In a similar fashion we see that the image of ϕ does not involve the terms in the p-filtration of $\nabla(\mu)$ labelled by $\mu[\lambda_1]$, $\mu[\lambda_4]$ and $\mu[\lambda_5]$, as the image of the head of $\nabla(\lambda)$ is a composition factor of $\nabla_p(\lambda_9)$. Thus the six terms in the filtration of $\nabla(\lambda)$ that may survive under ϕ are in distinct ∇_p-classes.

By the above remarks, ϕ induces a (non-zero) map ϕ' on the quotient of $\nabla(\lambda)$ by the submodule with a filtration by $\nabla_p(\lambda_1)$, $\nabla_p(\lambda_4)$ and $\nabla_p(\lambda_5)$. We shall denote this quotient module by $\nabla(\lambda)/[\lambda_1, \lambda_4, \lambda_5]$, and other quotients similarly. Now consider the restriction of ϕ' to $\nabla_p(\lambda_2)$. As this module is the only part of the filtration with any composition factors in common with those of $\nabla_p(\mu[\lambda_2])$ (which contains the socle of $\nabla(\mu)$) and ϕ' is non-zero, this restriction must also be non-zero. Hence we obtain a non-zero map from $\nabla_p(\lambda_2)$ to $\nabla_p(\mu[\lambda_2])$. By (1) we have

$$\mathrm{Hom}(\nabla_p(\lambda_2), \nabla_p(\mu[\lambda_2])) \cong \mathrm{Hom}(\nabla(\lambda_2''), \nabla(\mu[\lambda_2]'')).$$

A simple calculation shows that this latter pair of weights satisfy the conditions of Theorem 3.3 for $e-1$. Using the inductive hypothesis this Hom-space is one-dimensional, and hence any such homomorphism is unique (up to scalars). Therefore, by induction, we can describe the composition factors occurring in the image of such a homomorphism (up to the exceptional ambiguity in Theorem 3.3). As the isomorphism of Hom-spaces is induced by the map $V \longmapsto V^{\mathrm{F}} \otimes L(\lambda_2')$, we can thus describe the composition factors occurring in the image of our restriction of ϕ'.

We next consider the map obtained by following ϕ by the quotient map from $\nabla(\mu)$ which kills $\nabla_p(\mu[\lambda_2])$. As the image under ϕ of $\nabla_p(\lambda_2)$ is killed by this quotient map, this induces a map ϕ

from $\nabla(\lambda)/[\lambda_1, \lambda_2, \lambda_4, \lambda_5]$ to $\nabla(\mu)/[\mu[\lambda_2]]$. We wish to argue as above and show that the restriction of ϕ to each of the submodules $\nabla_p(\lambda_i)$, for $i \in \{3, 6, 7, 8\}$, can be determined by induction from a reflection about an $(e-1)$-wall. For this it is enough to show that the restriction is non-zero, as then the argument given above also holds in this case. Once we have shown this, the only remaining factor in $\nabla(\lambda)$ is $\nabla_p(\lambda_9)$ and a similar argument to the above shows that ϕ induces a (non-zero) map from this factor to $\nabla_p(\mu[\lambda_9])$, which again can be determined by induction.

Thus it only remains to show that the map induced on $\nabla_p(\lambda_i)$ is non-zero for $i \in \{3, 6, 7, 8\}$. Suppose that this map is zero, and hence that this factor is in the kernel of the original map ϕ. As the image under ϕ of $\nabla_p(\lambda_2)$ is non-zero, and this latter module has a simple head, this head must survive. Hence the submodule E_i of $\nabla(\lambda)/[\lambda_1, \lambda_4, \lambda_5]$ with a filtration by $\nabla_p(\lambda_2)$ and $\nabla_p(\lambda_i)$ must have a non simple head. We will show that the head is always simple, contradicting our assumption that the induced map is zero.

When $i = 3$ (respectively $i = 7$) the head of E_i is simple as there is a p-good homomorphism into $\nabla(\lambda)$ from some $\nabla(\tau)$ (which has a simple head) such that the image contains E_i as a quotient. For the remaining cases will imitate an argument in [21, pages 360-1]. As the head of E_i can contain at most two simple modules, it will be enough by [17, II, 3.16 (3)] to show that $\mathrm{Hom}_{G_1}(E_i, L(\lambda_2')) = 0$. From the defining sequence for E_i (and writing $L = L(\lambda_2')$) we have the exact sequence

$$0 \to \mathrm{Hom}_{G_1}(\nabla_p(\lambda_i), L) \to \mathrm{Hom}_{G_1}(E_i, L) \to \mathrm{Hom}_{G_1}(\nabla_p(\lambda_2), L) \to \mathrm{Ext}^1_{G_1}(\nabla_p(\lambda_i), L)$$

which by (2) and the fact that λ_2 and λ_i are in distinct ∇_p-classes gives

$$0 \to \mathrm{Hom}_{G_1}(E_i, L) \to \Delta(\lambda_2''^*)^{\mathrm{F}} \xrightarrow{\theta} \Delta(\lambda_i''^*)^{\mathrm{F}} \otimes \mathrm{Ext}^1_{G_1}(L(\lambda_i'), L).$$

Thus we will be done if we can show that θ is an embedding. First suppose that $\theta = 0$. Then $\mathrm{Hom}_{G_1}(E_i, L) = \Delta(\lambda_2''^*)^{\mathrm{F}}$ and this implies that E_i is semisimple as a G_1-module. By [17, II, 3.16 (3)] we see that $\mathrm{soc}_G(E_i)$ is not simple, which contradicts the fact that E_i embeds in $\nabla(\lambda_2)$. Therefore we must have that $\theta \neq 0$. To show that θ is an embedding, we will show that

$$\theta^{*(-1)} : \nabla(\lambda_i'') \otimes (\mathrm{Ext}^1_{G_1}(L(\lambda_i'), L)^{(-1)})^* \to \nabla(\lambda_2'') \tag{8}$$

is onto.

When $i = 6$ the results of Yehia [22] summarised in [21, Proposition 4.1] give that the Ext^1-group in (8) is isomorphic to k, and $\lambda_6'' = \lambda_2''$. As $\mathrm{Hom}(\nabla(\lambda_6''), \nabla(\lambda_2'')) \cong k$ we deduce that θ is an embedding.

When $i = 8$, [21, Proposition 4.1] gives that the Ext^1-group in (8) is isomorphic to $\nabla(1,0)^{\mathrm{F}}$ and $\lambda_2'' = (a, b-1)$, $\lambda_8'' = (a-1, b)$ for some $a, b > 0$. Thus $\theta^{*(-1)}$ is in

$$\begin{aligned} \mathrm{Hom}(\nabla(a-1,b) \otimes \ & \nabla(1,0)^*, \nabla(a,b-1)) \\ & \cong \mathrm{Hom}(\nabla(a-1,b), \nabla(1,0) \otimes \nabla(a,b-1)) \\ & \cong \mathrm{Hom}(\nabla(a-1,b), \mathrm{pr}_{(a-1,b)}\, \nabla(1,0) \otimes \nabla(a,b-1)) \\ & \cong \mathrm{Hom}(\nabla(a-1,b), T_{(a,b-1)}^{(a-1,b)} \nabla(a,b-1)) \cong k \end{aligned} \tag{9}$$

where the final isomorphism follows because $T_{(a,b-1)}^{(a-1,b)} \nabla(a,b-1) \cong \nabla(a-1,b)$ (as $b \not\equiv 0 \pmod p$). But there is an obvious surjection in the Hom-space containing $\theta^{*(-1)}$, as $\nabla(a-1,b) \otimes \nabla(1,0)^*$ has a good filtration with quotient $\nabla(a,b-1)$, and hence $\theta^{*(-1)}$ is surjective as required.

Next consider the cases when either $a'' \equiv 0 \pmod p$ or $b'' \equiv 0 \pmod p$. It is easy to see that the argument above is unaffected by the former, so we only need consider the case $b'' \equiv 0 \pmod p$. We are in the situation where λ is as in Figure 1a(iii) and μ is as in Figure 1b(iii).

Just as above, we consider the map induced by ϕ on various subquotients of $\nabla(\lambda)$. Everything goes through unchanged except for the cases involving $\nabla_p(\lambda_3)$ and $\nabla_p(\lambda_8)$.

Let $\nabla_p(\tau, \nu)$ denote a non-split extension of $\nabla_p(\nu)$ by $\nabla_p(\tau)$, and similarly for $\nabla(\tau, \nu)$. (We will only apply this notation in situations where the extension is unique.) By assumption, we have to consider the map induced by ϕ from $\nabla_p(\lambda_8, \lambda_3)$ to $\nabla_p(\mu[\lambda_3], \mu[\lambda_8])$. By Lemma 6.2 this notation is well-defined, and each of the modules has simple socle. Arguing as in the generic case we see that the restriction of ϕ to $\nabla_p(\lambda_3)$ must be non-zero. (It is clearly non-zero on $\nabla_p(\lambda_8)$.) There is an obvious homomorphism from $\nabla_p(\lambda_8, \lambda_3)$ to $\nabla_p(\mu[\lambda_3], \mu[\lambda_8])$ induced by the Carter-Payne homomorphism from $\nabla(\lambda_8'')$ to $\nabla(\mu[\lambda_8]'')$. But the restriction of this map to $\nabla_p(\lambda_3)$ is zero, and hence it cannot be the map we require. Thus the map induced by ϕ on $\nabla(\lambda_8'')$ cannot be the twist of the corresponding Carter-Payne map. As the socle of $\nabla_p(\mu[\lambda_3], \mu[\lambda_8])$ is simple it is clear that the same is true of the map induced by ϕ on $\nabla(\lambda_3'')$. This completes the proof in the case of reflection to the right from a down alcove. Clearly, reflection to the left is entirely analogous.

For the reflection ϕ below λ the argument is much more straightforward. The only term in the p-filtration of $\nabla(\lambda)$ which can survive under ϕ is $\nabla_p(\lambda_9)$, and this must map into $\nabla_p(\mu[\lambda_p])$. This map is induced by a Carter-Payne map from $\nabla(\lambda_9'')$ to $\nabla(\mu[\lambda_9]'')$, and hence is known by induction.

Case (ii): Up alcoves. We begin by considering reflection to the right of λ. As in case (i), we renumber the terms in the p-filtration of $\nabla(\mu)$ by their ∇_p-classes with respect to $\nabla(\lambda)$. (Again we use the generic labelling even in the non-generic case.) This is illustrated in Figure 5 (in the generic case).

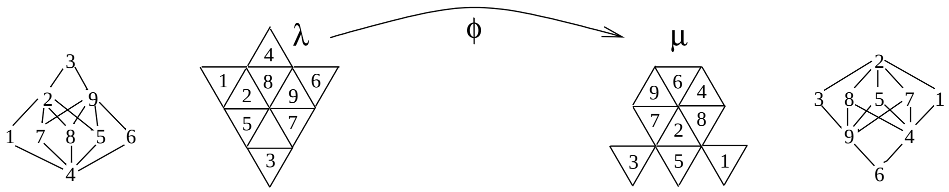

FIGURE 5.

If λ is generic then arguing as in case (i) we see that all but the terms labelled by λ_3, λ_6 and λ_9 must lie in the kernel of ϕ, and the image is contained in the submodule filtered by the terms labelled by $\mu[\lambda_3]$, $\mu[\lambda_6]$, and $\mu[\lambda_9]$. Arguing as in case (i) we see that ϕ is non-zero on each of the three subquotients, and corresponds to a Carter-Payne map from $\nabla(\lambda_i'')$ to $\nabla(\mu[\lambda_i]'')$, which are known by induction.

We wish to argue that the same is true in the non-generic case. Here it is possible that the term labelled by λ_5 may survive under ϕ, and/or the term labelled by $\mu[\lambda_8]$ may contain part of the image. In the former case we must have that $\nabla_p(\lambda_5)$ maps into $\nabla_p(\mu[\lambda_6])$ (not $\nabla_p(\mu[\lambda_5])$ as this would contradict the fact that the image of $\nabla(\lambda)$ has simple head). This implies the existence of a non-zero homomorphism from $\nabla(\lambda_5'')$ to $\nabla(\mu[\lambda_6]'')$. But $\mu[\lambda_6]'' \not\leq \lambda_5''$, so this is impossible. Similarly, if $\mu[\lambda_8]$ intersects $\operatorname{im}\phi$ then this corresponds to a non-zero homomorphism from $\nabla(\lambda_3'')$ to $\nabla(\mu[\lambda_8]'')$, but $\mu[\lambda_8]'' \not\leq \lambda_3''$. Thus the same argument as for the generic case holds here.

So far each of the cases we have considered has corresponded to the corresponding result when $e = 1$. However we shall see in the next case, the reflection below an up alcove, that this is not always the case. Let λ be in an up alcove, and ϕ be a reflection below λ. In Figure 6 we have (as usual) labelled the terms in the p-filtrations using the generic labels for λ.

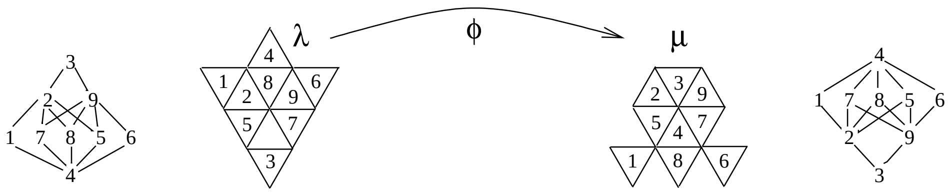

FIGURE 6.

In the generic case we see that ϕ must kill all but the top quotient $\nabla_p(\lambda_3)$, and corresponds to a Carter-Payne map from $\nabla(\lambda_3'')$ to $\nabla(\mu[\lambda_3]'')$. *Note that this differs from the case $e = 1$, when four terms in the filtration survive.* Clearly this map also exists in the non-generic case, but we also have the possibility that (as for $e = 1$) the terms labelled by λ_2, λ_3, λ_8 and λ_9 are not in the kernel of ϕ.

This would contradict the one-dimensionality of Hom-spaces, but we cannot use this to eliminate the possibility as we have not established that result for $\nabla(\lambda)$ at this stage. However, such a map would induce a non-zero homomorphism from $\nabla_p(\lambda_3)$ to $\nabla_p(\mu[\lambda_8])$, and hence a non-zero homomorphism from $\nabla(\lambda_3'')$ to $\nabla(\mu[\lambda_8]'')$. As we are in the non-generic case, it is easy to verify that $\gamma = \lambda_3''$ and $\nu = \mu[\lambda_8]''$ satisfy the conditions in Assumption 5.2. Hence this homomorphism cannot exist, as Assumption 5.2 follows from the results in Sections 7–9 which we know holds for weights smaller that λ by induction. Therefore there are no extra homomorphisms in the non-generic case.

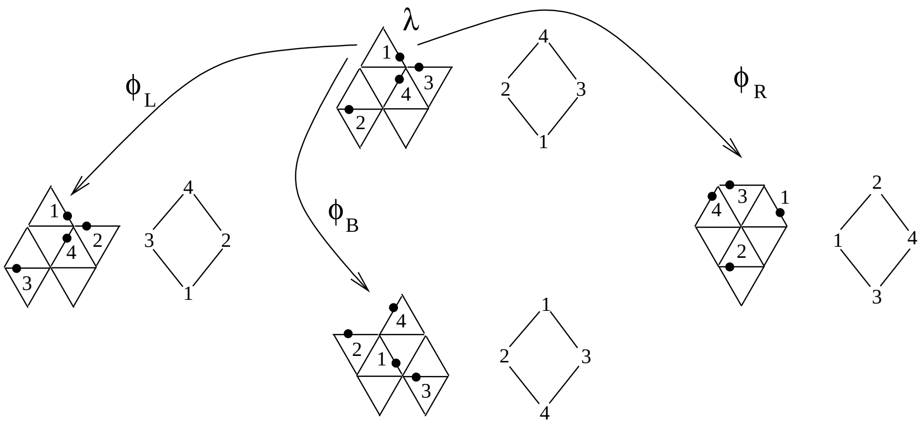

FIGURE 7.

Case (iii): Diagonal walls. We shall consider the case where λ lies on a wall as in Figure 7; the other case is symmetric. As usual there are three possible maps from $\nabla(\lambda)$: to the left (ϕ_L), right (ϕ_R) and below (ϕ_B). As usual we label the terms in the p-filtration of the target module $\nabla(\mu)$ by their (generic) ∇_p-classes with respect to the labelling of $\nabla(\lambda)$. The three cases are illustrated in Figure 7, in the generic case.

When λ is generic it is easy to verify as above that reflections to the right (respectively below) correspond to Carter-Payne maps (which are known by induction) from the two (respectively one) term(s) in the p-filtration that survive under the corresponding $e = 1$ map. The same is true for reflection to the left *provided one realises that the corresponding $e = 1$ map here is the identity morphism, i.e that all four terms survive.*

For non-generic λ the same results hold, but the proofs require more care. First consider reflection to the right. The only possible difference here is that $\nabla_p(\lambda_2)$ might survive under ϕ_R. However, $\mu[\lambda_3]'' \not\leq \lambda_2''$, and so this is impossible (as in the case of reflection to the right from an up alcove).

Next consider reflection to the left. The argument here is just as for the corresponding reflection from a down alcove.

Finally, the argument for reflections below λ is unchanged in the non-generic case, as the only pair of terms in the p-filtration which might be affected are already in the kernel of the map.

Case (iv): Horizontal walls. We conclude by considering the case where λ lies on a horizontal wall as in Figure 8. As usual there are three possible maps: to the left (ϕ_L), right (ϕ_R) and below (ϕ_B). As in the previous cases we number the terms in the p-filtration of $\nabla(\mu)$ by their (generic) ∇_p-classes with respect to the labelling for $\nabla(\lambda)$. By symmetry we only need consider reflections to the right and below.

As in the preceding case, when λ is generic it is routine to verify that reflections to the right correspond to Carter-Payne maps from the two terms in the p-filtration that survive under the

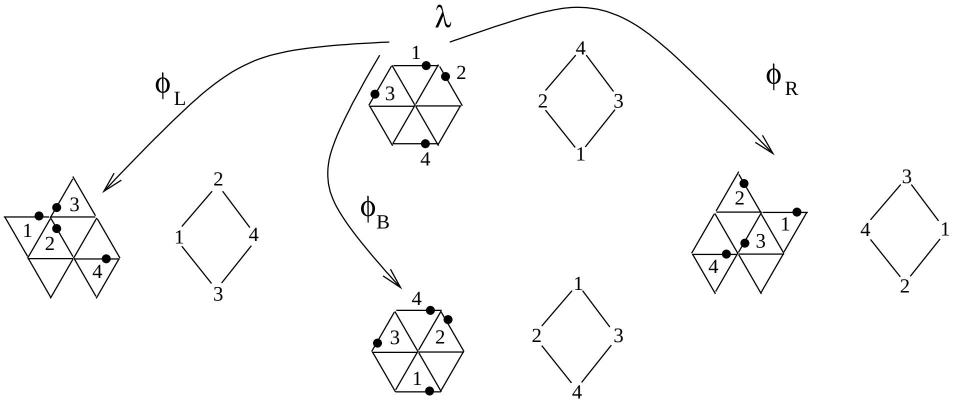

FIGURE 8.

corresponding $e = 1$ map. For reflections below the map corresponds to the Carter-Payne map from the top term of the p-filtration of $\nabla(\lambda)$ into $\nabla_p(\mu)$.

When λ is non-generic and ϕ is a reflection to the right or down then mimicking the argument for the corresponding reflection from an up alcove in case (ii) we see that the result is the same as in the generic case. □

We wish to determine when the composition of two Carter-Payne homomorphisms is non-zero. To describe this it will be convenient to introduce the following notation. Given a weight λ, we denote by $\nabla_p(\tilde{\lambda})$ the term in the p-filtration of $\nabla(\lambda)$ containing the head of $\nabla(\lambda)$. (Note that if $\lambda = p\lambda'' + \lambda'$ where $\lambda'' = (a, b)$ with $a > 0$ and $b > 0$, then $\tilde{\lambda} = p(\lambda'' - \rho) + w_a\lambda'$ where $w_a\lambda'$ is the reflection of λ' in the horizontal wall immediately above the origin.)

Recall that for a dominant weight λ, we let ϕ_L^i (respectively ϕ_R^i, ϕ_B^i) denote the Carter-Payne homomorphism from $\nabla(\lambda)$ about a p^i-wall to the left of (respectively to the right of, below) λ. *Recall further that in using this notation we always chose i minimal for the given wall.* If we do not wish to specify explicitly the direction of the reflection, we will replace L, R, B in the above by $\dagger$, and similarly the degree i of a wall by $*$. If the degree has to be strictly greater than 1, we replace the $*$ by > 1. Given a composition $\theta^t \cdots \theta^1$ of homomorphisms, we shall refer to any composition of the form $\theta^u \cdots \theta^1$ with $u < t$ as a *subcomposition* of the given composition.

We next define for each type of facet a set of *maximal compositions*. (The reason for this terminology will become clear in the next corollary.) Given a weight λ, we define the product $\theta^t \cdots \theta^1$, where each θ^i is of the form $\phi_\dagger^{*,i}$, to mean the composite of the Carter-Payne maps $\theta^i : \nabla(\lambda^i) \to \nabla(\lambda^{i+1})$ where $\lambda^1 = \lambda$ and λ^{i+1} is the weight μ such that the Carter-Payne map $\phi_\dagger^{*,i}$ from $\nabla(\lambda^i)$ has image in $\nabla(\mu)$. (Here we interpret any Carter-Payne map which starts in or leaves the dominant region as the zero map.) Now the set $\max(\lambda)$ of maximal compositions for λ depends on the type of facet in which λ lies:

λ	$\max(\lambda)$
Internal down alcove	$\{\phi_B^*,\ \phi_R^*\phi_L^*\phi_R^*,\ \phi_L^*\phi_B^1\phi_R^*,\ \phi_L^*\phi_R^*\phi_L^*,\ \phi_R^*\phi_B^1\phi_L^*\}$
Up alcove	$\{\phi_B^*,\ \phi_L^*\phi_R^*,\ \phi_R^*\phi_L^{>1}\phi_R^*,\ \phi_R^*\phi_L^*,\ \phi_L^*\phi_R^{>1}\phi_L^*,\ \phi_R^*\phi_B^1,\ \phi_L^*\phi_B^1\}$
RH just dominant down alcove	$\{\phi_R^*\phi_L^{>1},\ \phi_L^*\}$
LH just dominant down alcove	$\{\phi_L^*\phi_R^{>1},\ \phi_R^*\}$
Horizontal wall	$\{\phi_B^*,\ \phi_R^*\phi_L^*\phi_R^*,\ \phi_L^*\phi_R^*\phi_L^*\}$
RH diagonal wall	$\{\phi_L^*\phi_B^*,\ \phi_L^*\phi_R^*\phi_L^*,\ \phi_R^*\phi_L^*\phi_R^*,\ \phi_L^*\phi_B^*\phi_R^*\}$
LH diagonal wall	$\{\phi_R^*\phi_B^*,\ \phi_R^*\phi_L^*\phi_R^*,\ \phi_L^*\phi_R^*\phi_L^*,\ \phi_R^*\phi_B^*\phi_L^*\}$

The following corollary allows us to determine inductively when a composite of Carter-Payne maps can be non-zero. As all $\nabla(\lambda)$ have simple head, condition (i) is clearly both necessary and sufficient. However, the content of this lemma is that we can impose the additional condition (ii).

Corollary 6.4. *Let $\lambda^1, \ldots, \lambda^u$ be a sequence of distinct dominant weights such that for each i we have a Carter-Payne homomorphism $\phi_{\dagger}^{*,i} : \nabla(\lambda^i) \to \nabla(\lambda^{i+1})$. Then the composite $\phi_{\dagger}^{*,u} \cdots \phi_{\dagger}^{*,1}$ is non-zero if and only if both*

(i) *the induced homomorphism*

$$\phi_{\dagger}^{*,u} \cdots \phi_{\dagger}^{*,1} : \nabla_p(\widetilde{\lambda^1}) \to \nabla_p(\lambda^u[\widetilde{\lambda^1}])$$

is non-zero, and

(ii) *the composite $\phi_{\dagger}^{*,u} \cdots \phi_{\dagger}^{*,1}$ is a subcomposition of a composition in* $\max(\lambda)$.

Proof: As noted above, the first condition is itself both necessary and sufficient. That the second condition must also be satisfied for a composite to be non-zero follows in most cases from the description of maps via p-factors in Proposition 6.3. By considering each case in turn, we see that $\nabla_p(\lambda^u[\widetilde{\lambda^1}])$ is never one of the terms in the p-filtration of $\nabla(\lambda^i)$ for which any ambiguity remains in our description of the Carter-Payne homomorphisms from $\nabla(\lambda^i)$.

There are a few additional cases which are not eliminated by the p-factor description in Proposition 6.3. In the cases where λ^1 lies on a diagonal wall we can eliminate the additional possibility of two consecutive reflections in the same direction by reducing to the SL_2 case using (1) and [11, (4.3) Corollary], where the required Hom-spaces have been calculated in [6, Theorem 5.1]. Finally, for left-hand (respectively right-hand) diagonal walls we must eliminate the possibility of a composite of the form $\phi_R^*\phi_L^*\phi_R^*\phi_L^*$ (respectively $\phi_L^*\phi_R^*\phi_L^*\phi_R^*$). But for such a map to exist there must be a corresponding composite on $\nabla_p(\tilde{\lambda})$, which is impossible by the above description of $\max(\mu)$ for each possible type of μ and induction on the sum of the degrees of the maps in the quartet. $\square$

Remark 6.5. The induced morphism in Corollary 6.4(i) corresponds to a composite of Carter-Payne homomorphisms

$$\nabla(\widetilde{\lambda^1}'') \to \nabla(\lambda^u[\widetilde{\lambda^1}]'')$$

which has already been determined by induction.

We will take advantage of the notation used for elements in $\max(\lambda)$ to talk about maps starting from differing weights. Given two weights λ and τ and a composite of Carter-Payne maps $\phi = \phi_{\dagger}^{*,u} \cdots \phi_{\dagger}^{*,1}$ from $\nabla(\lambda)$, we shall call the map $\theta_{\dagger}^{*,u} \cdots \theta_{\dagger}^{*,1}$ from $\nabla(\tau)$ with $\theta_{\dagger}^{*,i} = \phi_{\dagger}^{*,i}$ for all $1 \leq i \leq u$ the map from $\nabla(\tau)$ *corresponding to* ϕ.

7. Determining homomorphisms I: interior weights

Suppose that λ is not in the lower closure of a just dominant (up or down) alcove. If λ is on a vertex then all homomorphisms are known inductively by (1), so we shall henceforth assume that this is not the case. In Section 4 and Theorem 5.1 we observed that there was an obvious upper bound Υ_λ for the set of weights μ for which there exist homomorphisms from $\nabla(\lambda)$ to $\nabla(\mu)$. In this section we will refine these sets to give a precise description of when homomorphisms occur, in most cases. The remaining, exceptional cases will be dealt with in Section 9. Those λ not of the form above will be considered in Section 8.

We will hereafter identify weights with the facets in which they lie — and (when we have a particular linkage class in mind) vice versa — and simple modules with their labels. Before continuing further, we note a convenient property of the vertex weights that simplifies certain verifications. We will wish to argue inductively from weights at vertices in our facet diagrams. For any such weight θ of the form $\theta = p\theta'' + (p-1)\rho$ we will need to consider the set of weights τ such that there is a non-zero homomorphism from $\nabla(\theta)$ to $\nabla(\tau)$. In principle to calculate such homomorphisms we need to consider θ'' and its images under the dot action of W_p (and thus reflect our weights in a different shifted set of hyperplanes). However, as this corresponds to translation by ρ (and we have rescaled) we can determine the weights τ by using the ordinary dot action of W_{p^2} on θ itself.

It will be convenient in what follows to have the following easy consequence of the definition of $\max(\lambda)$ and the associated Corollary 6.4.

Lemma 7.1. *Let λ, μ, and ν be dominant weights such that there exist Carter-Payne maps $\phi_L^a : \nabla(\lambda) \to \nabla(\mu)$ and $\phi_R^b : \nabla(\mu) \to \nabla(\nu)$ for some $a, b > 0$. Then the composite map*

$$\phi_R^b \phi_L^a : \nabla(\lambda) \to \nabla(\nu)$$

is non-zero. A similar result holds with the roles of left and right reversed.

Proof: It is enough to consider the case where λ does not lie on a vertex. Regardless of the type of the facet containing λ, the map $\phi_R^b \phi_L^a$ occurs as a subcomposition of an element in $\max(\lambda)$, by inspection (and the dominance of μ). Hence by Corollary 6.4 it is enough to show that the map from $\nabla_p(\tilde{\lambda})$ to $\nabla_p(\nu[\tilde{\lambda}])$ is non-zero. But $\tilde{\lambda}''$ and $\nu[\tilde{\lambda}]''$ are also related by a map of the form $\phi_R^* \phi_L^*$, and hence the result follows by induction on $a + b$. $\square$

Suppose $\theta \in X^+$ is a vertex weight. We refer to the closure of the hexagon surrounding θ, as the *neighbourhood* of θ (see Figure 9(a)). We refer to the open star surrounding θ as the *extended neighbourhood* of θ (see Figure 9(b)).

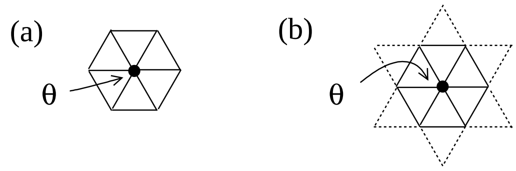

FIGURE 9.

The *local composition pattern* associated to a weight μ is the collection of elements of $W_p.\mu$ labelling terms in a p-filtration of $\nabla(\mu)$. If $\mu \in W_{p^2}.\lambda$ and $L(\mu)$ is a composition factor of $\nabla(\lambda)$, then these weights are generically composition factors of $\nabla(\lambda)$ which may be seen be using the p-filtration of $\nabla(\lambda)$.

For a non-vertex weight $\lambda \in X^+$ and a vertex weight $\theta \in X^+$, we say a dominant weight is a θ-*translate* for $\nabla(\lambda)$ if it corresponds to a composition factor of $\nabla(\lambda)$ and it lies in the neighbourhood of θ, if λ lies in the closure of a down alcove, and the extended neighbourhood of θ if λ lies in an up alcove.

Let λ be in the closure of an interior down alcove and let η the vertex below λ. We wish to define a set of eligible θ-translates for $\nabla(\lambda)$. Unfortunately the definition is complicated slightly when η is on a p^2-wall or p^2-vertex, so we first consider the remaining (i.e. p^2-regular) cases.

Let θ be another vertex dominated by η and with $\theta = w.\eta$ where $w \in W_{p^2}$. The *eligible θ-translates for* $\nabla(\lambda)$ are those highest weights corresponding to composition factors which are θ-translates for $\nabla(\lambda)$, and which occur in the local composition pattern for $\mu = w.\lambda$. Note that each of these factors lies in the same ∇_p-class as some weight labelling the p-filtration of $\nabla(\lambda)$.

When η is not p^2-regular the element w in the definition above is no longer unique. Let w be the element of W_{p^2} corresponding to the shortest sequence of reflections taking η to θ. Now the *eligible θ-translates for* $\nabla(\lambda)$ are those highest weights corresponding to composition factors which are θ-translates for $\nabla(\lambda)$, and which occur in the local composition pattern either for $\mu = w.\lambda$. or for some μ^- where $\mu^- \leq \mu$ is obtained from μ by a sequence of reflections in W_{p^2} all of which fix θ.

We have illustrated the six configurations of eligible θ-translates that will be of interest to us in Figure 10 for the case of λ in the interior of the alcove or on a left-hand diagonal wall. In both these cases the case of η non-p^2-regular gives the same result as the regular case. In both cases the vertex θ is indicated by a dot. For the alcove case the weight denoted μ in the preceding paragraph is in the highest alcove of the cluster of alcoves drawn (which represents the corresponding local composition pattern), and the unshaded alcoves in each configuration are those containing eligible θ-translates. For the wall case μ lies on the highest marked wall (the markings represent the corresponding local composition pattern), and the marked walls contained in the unshaded region are those containing eligible θ-translates. The right-hand diagonal wall case is symmetric, while the eligible θ-translates in the horizontal wall case are illustrated with bold walls in Figure 13 for η p^2-regular. For η non-p^2-regular the only modifications occur in cases (a) (b) and (c), which are illustrated in Figure 14 for μ a p^2-vertex; the p^2-wall cases are similar.

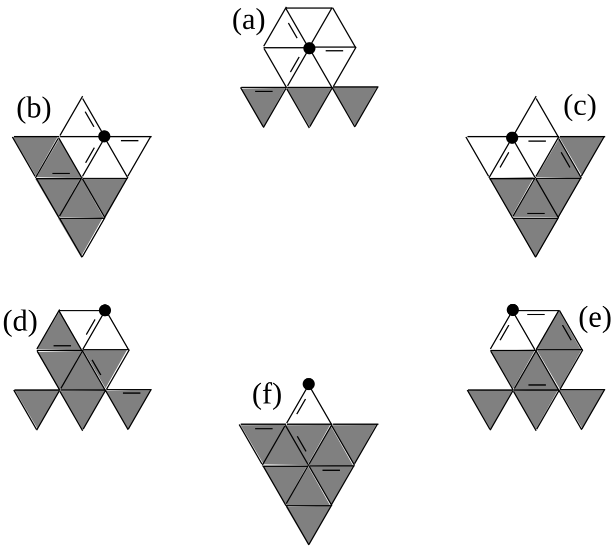

FIGURE 10.

Note that if ν is the lowest vertex in the closure of the alcove containing λ, with a homomorphism from $\nabla(\nu)$ to $\nabla(\theta)$, then the set of eligible θ-translates in $\nabla(\lambda)$ is a subset of Υ_λ. The next Theorem shows that the sets of eligible θ-translates provide the appropriate refinement of this set.

Theorem 7.2. *Suppose τ is a vertex and we have a composite of Carter-Payne maps $\nabla(\tau) \to \nabla(\theta)$. If λ lies in the closure of the down alcove with lower vertex τ, and $\mu \in W.\lambda$ lies in the neighbourhood of θ then*
$$\mathrm{Hom}(\nabla(\lambda), \nabla(\mu)) \neq 0$$
if and only if μ is an eligible θ-translate in $\nabla(\lambda)$. Further, if there is a non-zero homomorphism then it is either a composite of Carter-Payne maps or a twisted map.

Proof: We consider the case where η is p^2-regular; the remaining cases are easy modifications which are left to the reader. We will also only consider the case where λ lies in the interior of the down alcove; the remaining wall cases are similar, and are also left as an exercise for the reader. Note that λ and θ are such that we are in one of the 6 configurations illustrated in Figure 10 (a) to (f) (but it is not necessarily the case that all alcoves occurring in the local composition pattern of μ correspond to composition factors of $\nabla(\lambda)$). We begin by showing that there exist suitable maps for each eligible θ-translate.

As λ is in a down alcove, it is easy to verify that if μ is the lowest eligible θ-translate (i.e. the one immediately below θ) then it is a composition factor of $\nabla_p(\tilde{\lambda})$. Further, the existence of a map from $\nabla(\tau)$ to $\nabla(\theta)$ immediately implies that there is a map from $\nabla_p(\tilde{\lambda})$ to $\nabla_p(\mu)$ (as they both come from taking the map from $\nabla(\tau'')$ to $\nabla(\theta'')$, twisting by F, and tensoring both sides with an appropriate simple module). Thus if μ is the lowest eligible θ-translate then there is a twisted map from $\nabla(\lambda)$ to $\nabla(\mu)$.

Next suppose that μ is the highest eligible θ-translate. We will show that the desired map can be constructed as a composite of Carter-Payne maps using Corollary 6.4. In this case μ is obtained from λ by the same sequence of reflections used to obtain θ from τ. Write $\tau = (p^i-1)\rho + p^i\tau^{(i)}$ where $\tau^{(i)}$ does not lie on a vertex. Then the sequence of reflections used must label a subcomposition of a composition in $\max(\tau^{(i)})$ by Corollary 6.4 (though the maps will have different superscripts). If the same subcomposition also occurs as a subcomposition of a composition in $\max(\lambda)$ then the first condition in Corollary 6.4 is satisfied; by the argument in the previous paragraph, the second condition is also satisfied, and hence there is a non-zero composite of Carter-Payne maps from $\nabla(\lambda)$ to $\nabla(\mu)$.

Thus it only remains to consider the cases where a subcomposition ϕ of an element in $\max(\tau^{(i)})$ does not also occur as a subcomposition of an element in $\max(\lambda)$. By inspection we see that the following cases need to be considered:

(i) $\tau^{(i)}$ lies in a down alcove and ϕ is one of $\phi_R^*\phi_B^1\phi_L^*$, $\phi_B^1\phi_L^*$, $\phi_L^*\phi_B^1\phi_R^*$, or $\phi_B^1\phi_R^*$.
(ii) $\tau^{(i)}$ lies in an up alcove and $\phi = \phi_R^*\phi_B^1$ or $\phi_L^*\phi_B^1$.
(iii) $\tau^{(i)}$ lies on a right-hand diagonal wall and $\phi = \phi_L^*\phi_B^*$, $\phi_L^*\phi_B^*\phi_R^*$, or $\phi_B^*\phi_R^*$.
(iv) $\tau^{(i)}$ lies on a left-hand diagonal wall and $\phi = \phi_R^*\phi_B^*$, $\phi_R^*\phi_B^*\phi_L^*$, or $\phi_B^*\phi_L^*$.

In each case we will show that there is a second non-zero composite of Carter-Payne maps from τ to θ which does occur as a subcomposition of an element of $\max(\lambda)$. As we know that there is (up to scalars) a unique map from $\nabla(\tau)$ to $\nabla(\theta)$ this new map must coincide with the original one, and the arguments of the preceding paragraph will then apply to this new map to give the desired composite map from λ to μ. In what follows all equalities of maps should be interpreted as being up to some non-zero scalar multiple.

Consider first case (i), with $\phi = \phi_R^a\phi_B^1\phi_L^b$ for some $a, b > 0$. By Lemma 7.1 and the uniqueness of maps we have that $\phi_B^1\phi_L^b = \phi_L^b\phi_R^1$, and hence $\phi = \phi_R^a\phi_L^b\phi_R^1$, which does occur as a subcomposition of an element in $\max(\lambda)$. (Note that as ϕ is non-zero it follows that the conditions of Lemma 7.1 are satisfied.) Similar arguments hold for the other maps in case (i). Case (ii) is similar as, for example, we have that $\phi_R^a\phi_B^1 = \phi_L^1\phi_R^a$.

Cases (iii) and (iv) are symmetric, so we only need consider (iii). We would like to argue as in the preceding paragraph, but it is no longer obvious that in applications of Lemma 7.1 the intermediate weight is dominant. Consider the case where $\phi = \phi_L^a\phi_B^b$ (the other cases are similar). Consider the original map ϕ from $\nabla(\tau)$ to $\nabla(\theta)$. By Corollary 6.4 there must be a corresponding composite map from $\nabla_p(\tilde{\tau})$, and hence from $\nabla(\tau'')$. Repeating this argument for $\nabla(\tau'')$ (and so on), we deduce that the map must ultimately have come from some corresponding composite with a or b equal to 1. By considering the possible cases (up alcove or RH diagonal wall) where such a composite can arise, we see that we must have $b \leq a$, and further that the map from λ of the form ϕ_L^a must be to a dominant weight. We have $\phi_R^b\phi_L^a = \phi_L^a\phi_B^b$ from τ, and hence $\phi = \phi_R^b\phi_L^a$. This is an element of $\max(\lambda)$, and hence the argument above produces the desired map.

The only remaining cases of eligible θ-translates occur in cases (a), (b) and (c) in Figure 10. As above, the sequence of reflections relating τ and θ must be a subcomposition of some element in $\max(\tau)$. By considering the various possible cases we see that a configuration as in case (b) (respectively (c)) can only occur when θ is obtained from τ by a single reflection to the left (respectively right), while for (a) we must have $\theta = \tau$. It is now easy to construct an explicit composite of Carter-Payne maps for each of the remaining weights μ.

For λ not too close to a p^2-wall we have shown that of the six alcoves in each cluster of possible candidates for homomorphisms, either we have a morphism of the desired form or the alcove contains no composition factor of $\nabla(\lambda)$. Thus in that case we are done. When λ is close to a p^2-wall it is possible that some of the six alcoves in a cluster are generated as composition factors of $\nabla(\lambda)$ by a weight on a vertex other than the one at the centre of the cluster. As long as we can eliminate such alcoves as candidates for a homomorphism the argument will proceed for the remaining alcoves just as above. We will consider the case where λ is close to a single p^2-wall; the case where it is close to two such walls is very similar.

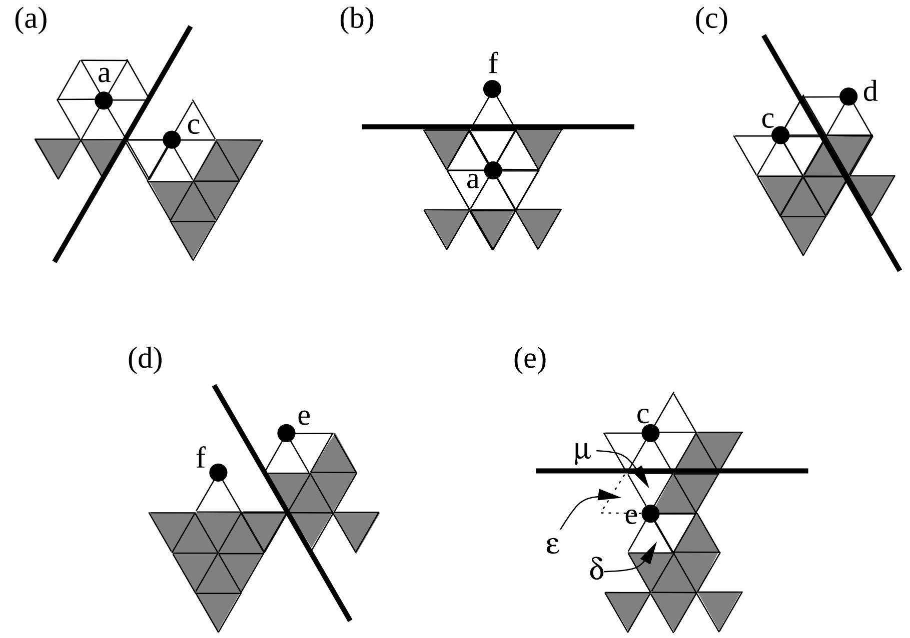

FIGURE 11.

If λ lies close to a p^2-wall, we see that the problem case can occur if there exist vertices θ_1 and θ_2 separated by a single wall such that $\theta_1 > \theta_2$ and there is a non-zero homomorphism from $\nabla(\tau)$ to $\nabla(\theta_2)$. In this case it may be that there is a local composition pattern clustered around θ_1 formed from composition factors of $\nabla(\lambda)$, which intersects the neighbourhood of θ_2. The various configurations that can occur (after using symmetry considerations to reduce the number of cases) are shown in Figures 11 and 12. Here we label each of the local composition patterns by the corresponding label in Figure 10 and shade in all but the θ_1- and θ_2-translates. Figure 11 corresponding to τ'' in a down alcove, and Figure 12 to τ'' in an up alcove (the other cases are similar). By inspection, we see that the only cases where additional homomorphisms from $\nabla(\lambda)$ to $\nabla(\mu)$ might arise are when μ is as indicated in Figure 11(e) or Figure 12(b) or (e).

In these cases we need to show that there is no homomorphism $\nabla(\lambda) \to \nabla(\mu)$: once we have done this we will have completed the proof for λ in the interior of a down alcove. So suppose there is a homomorphism $\nabla(\lambda) \to \nabla(\mu)$ with μ as shown. An easy calculation shows that the factor in the p-filtration of $\nabla(\mu)$ in the same ∇_p-class as $\nabla_p(\tilde{\lambda})$ is $\nabla_p(\delta)$ where δ lies in the alcove shown. As $\nabla(\lambda) \to \nabla(\mu)$ is non-zero, the head of $\nabla(\lambda)$ (i.e. the head of $\nabla_p(\tilde{\lambda})$) must survive in the image as a composition factor of $\nabla_p(\delta)$. But we have local homomorphisms $\nabla(\mu) \to \nabla(\epsilon)$ (where ϵ lies in the dotted alcove shown) where all of $\nabla_p(\delta)$ survives. Therefore we obtain a non-zero homomorphism $\nabla(\lambda) \to \nabla(\mu) \to \nabla(\epsilon)$, which is impossible as $L(\epsilon)$ is not a composition factor of $\nabla(\lambda)$. □

We next want to similarly refine Υ_λ for λ in an up alcove using sets of eligible θ-translates.

Let λ lie in an up alcove and η the vertex directly below it.(That is, η will be the centre of the star with λ in the topmost alcove of the star.) We wish to define the set of eligible θ-translates for $\nabla(\lambda)$; just as in the down alcove case we begin by supposing that η is p^2-regular.

Let θ be another vertex dominated by η with $\theta = w.\eta$ and $w \in W_{p^2}$. The *eligible θ-translates for* $\nabla(\lambda)$ are those weights corresponding to composition factors lying in the extended neighbourhood of θ which occur in the local composition pattern for $\mu = w.\lambda$. When η is not p^2-regular the modifications are exactly as in the down alcove case.

As in the down alcove case there are six configurations of eligible θ-translates that will be of interest to us, and these are illustrated in Figure 13 for η p^2-regular. (These six configurations may be superimposed if we are near a p^2-wall.) As before the vertex θ is indicated by a dot,

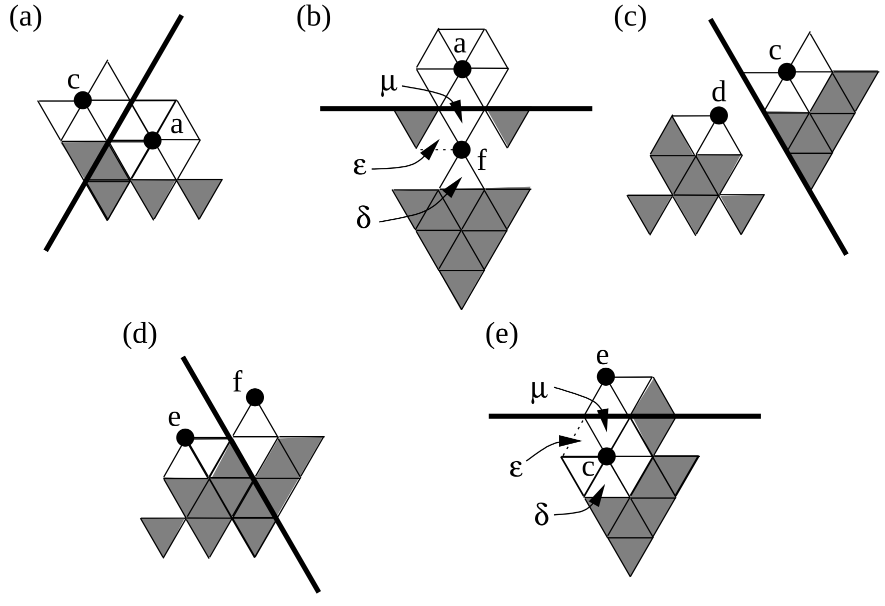

FIGURE 12.

the weight denoted μ in the preceding paragraph is in the highest alcove of the cluster of alcoves drawn (which represents the corresponding local composition pattern), and the unshaded alcoves in each configuration are those containing eligible θ-translates. The thick lines indicate the eligible θ-translates for λ in the horizontal wall case, and are illustrated here for convenience later. For η non-p^2-regular the only modifications occur in cases (a) (b) and (c), which are illustrated in Figure 14 for η a p^2-vertex (the p^2-wall cases being similar).

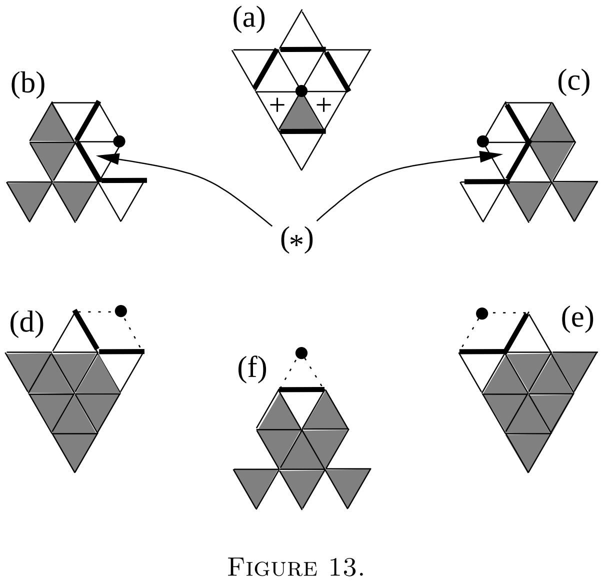

FIGURE 13.

Unfortunately the up alcove analogue of Theorem 7.2 is a little more complicated than in the down alcove case. We define an exceptional alcove for the pair (θ, λ) to be an alcove $(*)$ as in Figures 13 or 14, when the relative positions of θ and the θ-eligible translates in $\nabla(\lambda)$ are as shown. We now have

Theorem 7.3. *Suppose τ is a vertex and we have a composite of Carter-Payne maps $\nabla(\tau) \to \nabla(\hat{\theta})$. Let λ lie in an up alcove, with ν on the horizontal wall below λ and τ the lower vertex in the closure*

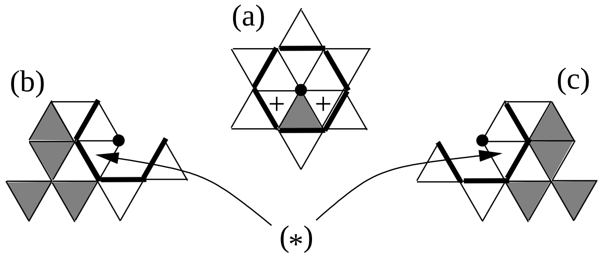

FIGURE 14.

of the down alcove below ν. If we have a weight $\mu \in W.\lambda$ with μ in the extended neighbourhood of $\hat{\theta}$ then

$$\mathrm{Hom}(\nabla(\lambda), \nabla(\mu)) \neq 0$$

if and only if there exists some weight θ satisfying the following four conditions:

1. *There exists a non-zero composite of Carter-Payne maps $\nabla(\tau) \to \nabla(\theta)$.*
2. *The weight μ is an eligible θ-translate in $\nabla(\lambda)$.*
3. *There is a non-zero homomorphism from $\nabla(\nu)$ to $\nabla(\delta)$ with δ on the wall in the closure of the alcove containing μ which lies in the boundary of the neighbourhood of θ.*
4. *The weight μ is not an exceptional weight for the pair (θ, λ).*

Further, if there is a non-zero homomorphism then it is either a composite of Carter-Payne maps or a twisted map.

Remark 7.4. (i) As we will see in the following proof, we must have $\theta = \hat{\theta}$ unless λ is too close to a p^2-wall (i.e. unless there are two vertices θ_1 and θ_2 with overlapping extended neighbourhoods such that θ_1 and θ_2 are both in $W_{p^2}.\tau$).
(ii) Condition (3) eliminates the possibility of maps into the alcoves marked (+) in Figure 13 in the generic case.

Proof: We proceed as in the proof of the last Theorem, and as in that proof we will only consider the case where η is p^2-regular; the remaining cases are easy modifications which are left to the reader. We begin by assuming that λ is not close to a p^2-wall, so that λ and θ are such that we are in one of the six configurations illustrated in Figure 13 (a) to (f). We now consider the eligible θ-translates.

As in the preceding theorem, it is easy to verify that if μ is the lowest eligible θ-translate (i.e. the down alcove below θ) then it is a composition factor of $\nabla_p(\tilde{\lambda})$. Just as before, we deduce that there exists a twisted map from $\nabla(\lambda)$ to $\nabla(\mu)$.

Similarly, if μ is the highest eligible θ-translate, we wish to argue exactly as in the preceding proof to show that there is a composite of Carter-Payne maps from $\nabla(\lambda)$ to $\nabla(\mu)$. As there we see that the only problem occurs when the composite of Carter-Payne maps from τ to θ does not correspond to a subcomposition of an element of $\max(\lambda)$. With the notation used in the preceding proof, it remains to consider the following cases:

(i) $\tau^{(i)}$ lies in a down alcove and $\phi = \phi_R^*\phi_B^1\phi_L^*$ or $\phi_L^*\phi_B^1\phi_R^*$.
(ii) $\tau^{(i)}$ lies in an up alcove and $\phi = \phi_R^*\phi_B^1$ or $\phi_L^*\phi_B^1$.
(iii) $\tau^{(i)}$ lies on a right-hand diagonal wall and $\phi = \phi_L^*\phi_B^*$ or $\phi_L^*\phi_B^*\phi_R^*$.
(iv) $\tau^{(i)}$ lies on a left-hand diagonal wall and $\phi = \phi_R^*\phi_B^*$ or $\phi_R^*\phi_B^*\phi_L^*$.

As before, in each case we will construct a second non-zero composite of Carter-Payne maps from τ to θ which does correspond to a subcomposite of an element from $\max(\lambda)$. The desired maps then exist by the arguments in the preceding proof.

Case (i) is straightforward, as there is only one eligible θ-translate. Thus the highest eligible θ-translate is also the lowest, and we have already constructed the required map. Case (ii) is also easy, as we have that $\phi_R^a\phi_B^1 = \phi_L^1\phi_R^a$ (and similarly with the roles of left and right reversed), and it is easy to see that this composite satisfies the conditions of Lemma 7.1.

As the right-hand and left-hand wall cases are symmetric, we only need consider case (iii). If $\phi = \phi_L^a\phi_B^1$ we argue as in case (ii). If $\phi = \phi_L^a\phi_B^b\phi_R^c$ we have to be careful (as in the proof of case (iii) in Theorem 7.2) that in applying Lemma 7.1 we have that the intermediate weight is dominant. However, as before, we can deduce that the map must ultimately come from a map from a right-hand diagonal or internal down alcove, and just as in the corresponding case in Theorem 7.2 we deduce that $\phi_B^b\phi_R^c = \phi_R^c\phi_L^b$, and hence that $\phi = \phi_L^a\phi_R^c\phi_L^b$ which corresponds to an element in $\max(\lambda)$ as required.

The only remaining cases of eligible θ-translates occur in cases (a), (b) and (c) in Figure 13. As in the preceding proof, we see that a configuration as in case (b) (respectively (c)) can only occur when θ is obtained from τ by a single reflection to the left (respectively right), while for (a) we must have $\theta = \tau$.

In case (a) we already know that there exist maps of the desired form into all alcoves except those labelled (+). For each of these latter alcoves no map exists by translation arguments, as there is no map from $\nabla(\nu)$ to $\nabla(\delta)$, with δ as in condition 3. (This case is the reason for this condition in the Theorem, which eliminates these two possibilities.)

For cases (b) and (c) it is now easy to construct an explicit composite of Carter-Payne maps for each of the remaining weights μ, except for the exceptional weights labelled $(*)$. As these two cases are symmetric we consider case (c), and must show that there is no non-zero map from $\nabla(\lambda)$ to $\nabla(\mu)$ when μ is the weight $(*)$. Consider the p-filtration for $\nabla(\mu)$, using the labelling of factors given in Figure 1a(i). If there were a non-zero map then the head of $\nabla(\lambda)$ must survive, and by earlier arguments (and our assumption that λ is not near a p^2-wall) it must be a composition factor in $\nabla_p(\mu_8)$. As there is a map from $\nabla(\mu)$ to $\nabla_p(\mu_4)$ which does not kill any of $\nabla_p(\mu_8)$, the composite of this pair of maps would give a non-zero map from $\nabla(\lambda)$ to $\nabla(\mu_4)$. But $L(\mu_4)$ is not a composition factor of $\nabla(\lambda)$, which gives the desired contradiction.

By our assumption on λ, the remaining weights in the extended neighbourhood of θ for which we have not yet constructed maps are not composition factors of $\nabla(\lambda)$. Thus we are done in the case where λ is not too close to a p^2-wall.

It remains to consider the case where λ is close to a p^2-wall. If λ is close to only one p^2-wall then the new configurations that can occur are illustrated in Figure 15 (for τ'' in a down alcove, the wall cases are similar) and Figure 16 (for τ'' in an up alcove). Here the vertices are labelled by the labels of the corresponding configurations in Figure 13, and all alcoves not in the extended neighbourhood of one of the two marked vertices are shaded. When λ is close to two p^2-walls it is easy to verify that the obvious modification of the arguments below (considering each p^2-wall case separately, then superimposing them) holds.

The construction of maps proceeds just as in the case considered above, with the understanding that there may be two 'highest' eligible θ-translates: one for each local composition diagram when a pair of such overlap. However, it remains to show that there are no additional maps into the remaining unshaded alcoves in Figures 15 and 16.

If there is a map from $\nabla(\nu)$ to $\nabla(\delta)$ in Figure 13 with δ and ν as in Condition 3 then the alcove marked with a + must coincide with an exceptional alcove (marked $(*)$) and so this map is eliminated as in cases (b) and (c). For the alcove marked $(*)$ in Figure 13 the argument given above only breaks down in the case shown in Figures 15(c) and 16(a). Consider first the former case. The argument given above breaks down because μ_4 (here labelled by δ_B) is a composition factor of $\nabla(\lambda)$. However, it will be enough to show that there is no map from $\nabla(\lambda)$ to $\nabla(\delta_B)$, which is proved below.

Next consider the case shown in Figure 16(a). For convenience we have redrawn this case in Figure 17 with the labelling of the p-filtration for $\nabla(\mu)$ shown and all other alcoves shaded. Here we may no longer deduce that the head of $\nabla(\lambda)$ must be a composition factor of $\nabla_p(\mu_8)$, as it could also occur in $\nabla_p(\mu_3)$. Suppose that this is the case. Then our map must induce a morphism

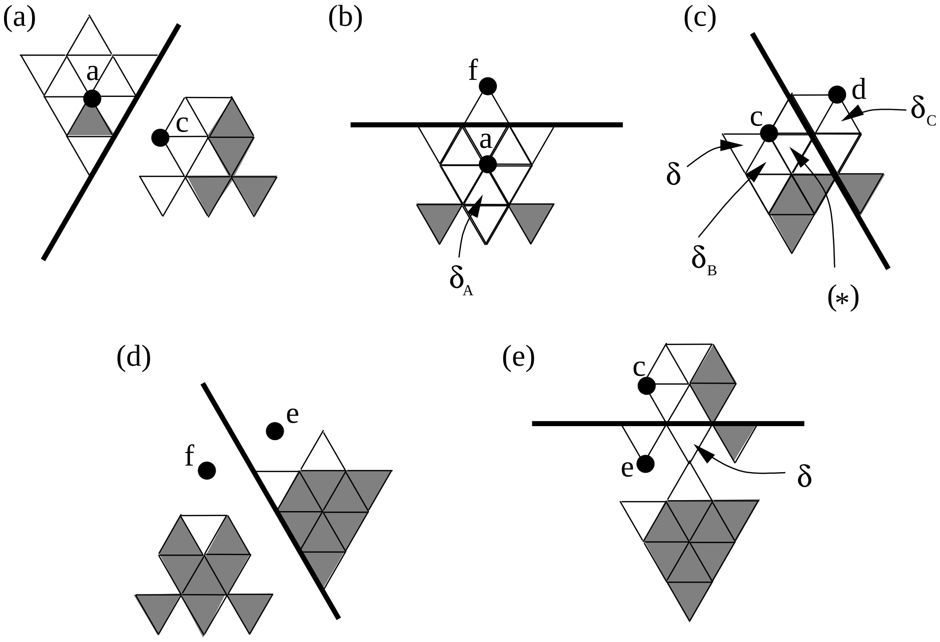

FIGURE 15.

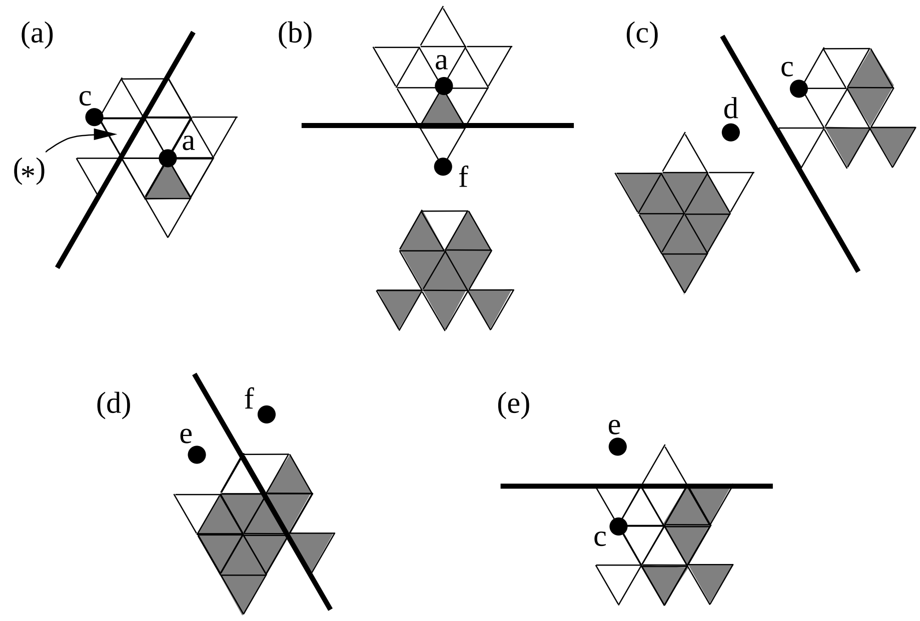

FIGURE 16.

from $\nabla_p(\lambda_3)$ to $\nabla_p(\mu_3)$, and hence there exists a map from $\nabla(\lambda_3'')$ to $\nabla(\mu_3'')$. But these two weights differ by a multiple of a single root and are not reflections of each other about a wall. Thus this map cannot exist by [11, (3.6) Theorem and (4.3) Corollary]. Thus we see that the head of $\nabla(\lambda)$ must be a composition factor of $\nabla_p(\mu_8)$, and the argument given in the generic case goes through unchanged.

It remains to show that the unshaded alcoves in Figures 15 and 16 for which we have not yet constructed maps do indeed have no maps into them. First suppose that there are maps from $\nabla(\tau)$

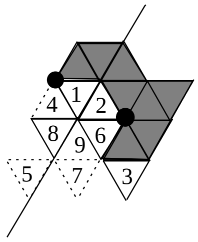

FIGURE 17.

into *both* labelled vertices in a given configuration. Then the remaining cases to consider are the alcoves labelled by Greek letters in Figures 15(b), (c) and (e).

Those labelled simply by a δ clearly have no maps, as they fail condition (3) of the Theorem, which is a necessary condition for a map to exist by the translation arguments in the proof of Theorem 5.1. Next consider $\mu = \delta_A$, δ_B, or δ_C and suppose there is a non-zero map from $\nabla(\lambda)$ to $\nabla(\mu)$. In each case μ must be a composition factor of $\nabla_p(\lambda_4)$. Hence the head of $\nabla_p(\lambda_9)$ must survive under the map, as the submodule of $\nabla(\lambda)$ with ∇_p-factors $\nabla_p(\lambda_4)$, $\nabla_p(\lambda_5)$, $\nabla_p(\lambda_6)$, $\nabla_p(\lambda_7)$, $\nabla_p(\lambda_8)$ and $\nabla_p(\lambda_9)$ has simple head. But there is no composition factor of $\nabla_p(\lambda_9)$ in $\nabla(\mu)$ as $\nabla(\mu)$ does not contain this ∇_p-class.

Finally, we have to consider the case where there is a map from τ to just *one* of the two labelled vertices in a given configuration. As for the cases labelled by δ's above, an appeal to the necessary condition (3) in the Theorem is enough to eliminate the remaining alcoves in most cases.

However, two cases cannot be eliminated in this manner. These are the configuration shown in Figure 15(b) when there is a map from τ only to vertex f, and the configuration shown in Figure 15(c) when there is map from τ only to vertex d. In each case we must eliminate the possibility of a map to each of the alcoves adjacent to vertex f (respectively d).

In case (b) it is easy to verify that the alcove is in the same ∇_p-class as λ_4. Then arguing as above we see that the head of $\nabla_p(\lambda_9)$ must survive in the image of any map, and thus the composite of this map followed by a local reflection to the right would be non-zero. However, no such map exists by our assumption.

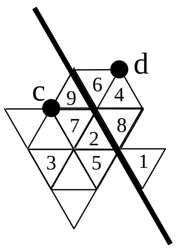

FIGURE 18.

In case (c) the lower of the two alcoves has already been eliminated above. The remaining alcove is marked 6 in Figure 18, where the numbering indicates the ∇_p-class using the labelling for λ. The structure of the p-filtrations of $\nabla(\lambda)$ and $\nabla(\mu)$ are very similar to those shown on the left (respectively right) sides of Figure 5 — except that in the latter case there is an extension of 3 by 8, as these are in the same ∇_p-class. Any map from $\nabla(\lambda)$ must preserve the head of $\nabla_p(\lambda_3)$, and its image must lie in the term in the p-filtration of $\nabla(\mu)$ corresponding to the alcove marked 3 or 8.

If this composition factor lies in $\nabla_p(\mu[\lambda_8])$ then composing the map with the local reflection to the right gives a non-zero map to $\nabla(\mu[\lambda_4])$, which contradicts the arguments above. Hence the image of $\nabla(\lambda)$ under our map must be in the submodule of $\nabla(\mu)$ with p-filtration labelled by $\mu[\lambda_3]$,

$\mu[\lambda_6]$, and $\mu[\lambda_9]$. But then the composite of our map with the obvious quotient map to $\nabla_p(\mu[\lambda_3])$ must be non-zero, which implies the existence of a non-zero map from $\nabla(\lambda)$ to $\nabla(\mu[\lambda_3])$. However, no such map exists by our assumption. □

For twisted maps, the analogue of Theorems 7.2 and 7.3 is in most cases much more straightforward.

Theorem 7.5. *Suppose τ is a vertex and we have a twisted map $\nabla(\tau) \to \nabla(\theta)$. Let λ lie in the closure of the down alcove with lower vertex τ, or in an up alcove with ν on the horizontal wall below λ and τ the lower vertex in the closure of the down alcove below ν. Write $\tau = (p^i - 1)\rho + \tau^+$ where τ^+ is not a vertex weight. If $\widetilde{\tau^+} = p((\tau^+)'' - 1)\rho + w_a((\tau^+)')$ then*

$$\mathrm{Hom}(\nabla(\lambda), \nabla(\mu)) \neq 0$$

if and only if μ is an eligible θ-translate in $\nabla(\lambda)$. Further, if there is a non-zero homomorphism then it is twisted map.

Proof: By our assumption on τ, the existence of a twisted map from $\nabla(\tau)$ to $\nabla(\theta)$ implies that the θ-eligible factors of $\nabla(\lambda)$ are in one of the configurations (f) in Figures 10 and 13. As in both these cases there is only one eligible factor μ (which is thus the lowest such) the proof of the existence of a (necessarily twisted) map follows just as in the corresponding case in Theorems 7.2 and 7.3. □

The assumption in the preceding Theorem that τ can be written in the form $\tau = (p^i - 1)\rho + \tau^+$ where τ^+ is not a vertex weight only fails to hold when τ^+ is close to the boundary of the dominant region. For such weights a twisted map from $\nabla(\tau)$ to $\nabla(\theta)$ will give rise to the first of our exceptional maps, as well as the twisted maps in Theorem 7.5. We return to this problem in Section 9, after we have considered those weights λ not of the form considered in Theorems 7.2–7.5.

8. Determining homomorphisms II: just dominant weights

We next turn our attention to weights λ for which the methods of the previous section do not apply. We first consider the case where λ is in the closure of a down alcove whose lower vertex is non-dominant. By the translation principle, we may assume that $\lambda = (a, 0)$ or $(0, b)$. It is easy to show (cf. [17, II, 2.16]) that $\nabla(a, 0) \cong S^a(V)$, where V is the natural module for $\mathrm{SL}_3(k)$, and the submodule structure of these has been determined by Doty [10]. We will use (and refine) the reinterpretation of these results in [5].

Let $a = \sum_{i=0}^{m} a_i p^i$ where $0 \leq a_i \leq p-1$ for all i and $a_m \neq 0$. Doty gives a procedure for defining a set $C(a)$ of m-tuples of nonnegative integers (called carry patterns) which is in one-to-one correspondence with the set of composition factors of $S^a(V)$. Further, if we impose a poset structure on $C(a)$ by setting $\mathbf{c} \leq \mathbf{d}$ if $c_i \leq d_i$ for all i then this correspondence induces a lattice isomorphism from the lattice of order closed subsets of $C(a)$ to the lattice of submodules of $S^d(V)$. Doty's construction is purely arithmetic; we will use the alcove theoretic version given in [5].

Recall that in [5, Figure 5] each type of alcove that can contain a composition factor of $S^a(V)$ was associated to either 0, 1 or 2. We will extend this procedure to walls: we reproduce the correspondence between alcoves and integers in Figure 19, and extend it to the wall cases as shown. Note that in the new cases the numbering depends on the type of wall containing λ.

We now recall the recursive procedure for determining composition factors of $\nabla(\lambda)$ given in [5], in the special case where $\nabla(\lambda)$ is a symmetric power. (This allows certain simplifications to be made.) As this is a little complicated, an example will be given below. The algorithm starts by finding the largest i such that λ does not lie in the closure of the lowest p^i-alcove. By regarding the facets shown in Figure 19 as p^i-facets, we let $\mathrm{cf}(\lambda, i)$ be the set of weights obtained from λ by intersecting $W_{p^i}.\lambda$ with the configuration of p^i-facets corresponding to the position of λ. (The weight λ will always occur in this set.) In this first iteration of the algorithm the configuration containing λ will be one with a shaded region to the right of it.

For the inductive step, we construct a new set $\mathrm{cf}(\lambda, i-1)$ from $\mathrm{cf}(\lambda, i)$. For each weight μ in $\mathrm{cf}(\lambda, i)$ we consider the unique translate of the set of p^{i-1}-restricted weights which contains μ and

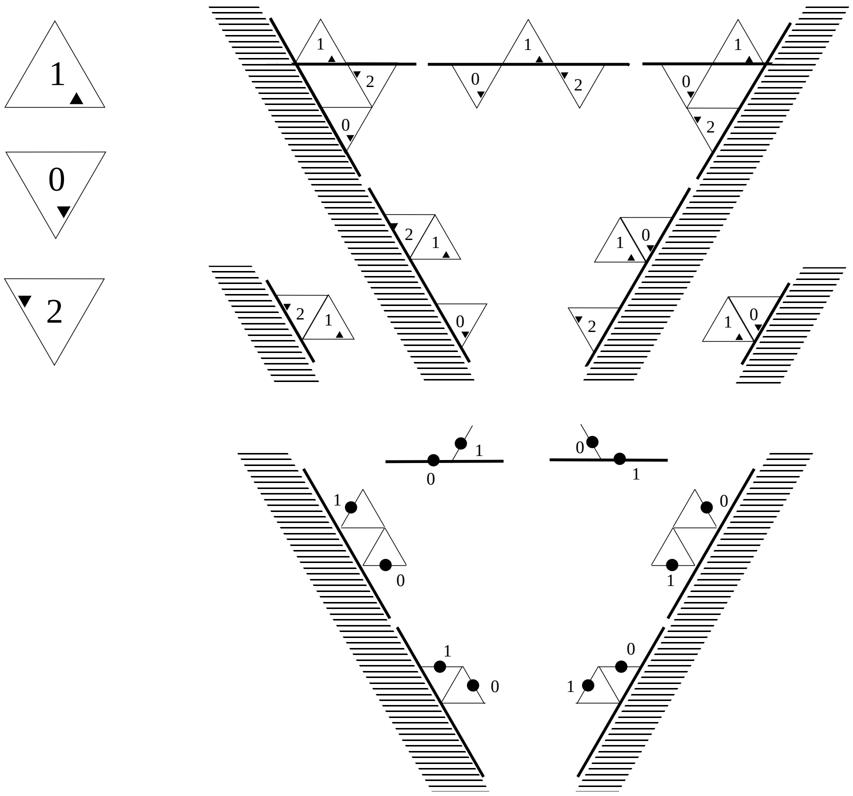

FIGURE 19.

determine the configuration of p^{i-1}-facets given in Figure 19 which corresponds to the position of μ *in this set*. (Thus the shaded regions here represent weights outside of the translate considered.) Then the intersection of $W_{p^{i-1}}.\mu$ with these p^{i-1}-facets gives the (immediate) descendants of μ in $\mathrm{cf}(\lambda, i-1)$. All elements of $\mathrm{cf}(\lambda, i-1)$ arise in this manner. In [5] it was proved that $\mathrm{cf}(\lambda, 1)$ is precisely the set of composition factors of $\nabla(\lambda)$.

As $S^a(V)$ has no repeated composition factors, each weight μ that labels such a factor has a unique collection of ancestors; i.e. a series $\mu^i = (a,0), \mu^{i-1} \ldots, \mu^1 = \mu$ such that μ^j occurs as an immediate decedent of μ^{j+1} in $\mathrm{cf}(\lambda, j)$. (Note that not all of these weights are necessarily distinct.) Each such μ^j comes from a configuration of facets given in Figure 19, and hence has an associated integer $a_j(\mu)$ from $\{0,1,2\}$. This is the series of integers which we assign to μ. We can now prove the following refinement of [5, Proposition 5.4].

Proposition 8.1. *Given a composition factor $L(\mu)$ of $S^a(V)$, the associated carry pattern is the m-tuple of integers whose jth entry equals $a_j(\mu)$*

Proof: Argue as in the proof of [5, Proposition 5.4], noting that in the extra cases (where $(a,0)$ lies on a wall), the p^i-facets containing $(a,0)$ are all the same type until the first such p^j-facet which is a alcove, after which the remaining p^i-facets (for $i \geq j$) are also all of the same type. $\square$

As an example of this Proposition consider $S^{184}(V)$ wit $p = 5$. The composition factors and their associated carry patterns are illustrated in Figure 20, together with the first two iterations of the composition factor algorithm and the complete submodule lattice.

At the first iteration of the algorithm, the configuration of alcoves used is as shown in Iteration 1, and hence the final digit of each carry pattern is either 0 or 1 by Figure 19. Using Iteration 2 we can determine the penultimate digit of each factor; for example those arising from Iteration 2(c) will have final pair 21. For the final iteration we use the wall cases given in Figure 19 to give the initial digit of each carry pattern.

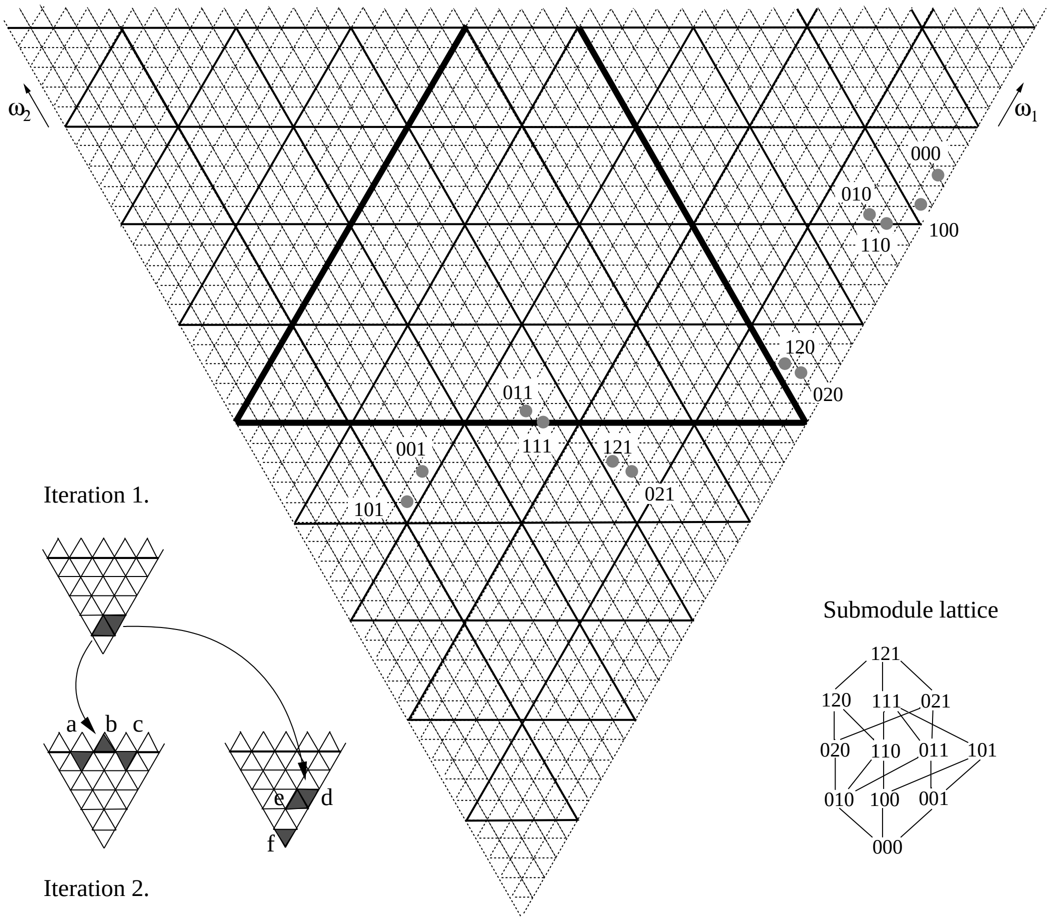

FIGURE 20.

Another, slightly more detailed, example of the mechanics of the composition factor algorithm (in the alcove case, with $\lambda = (181, 0)$, again with $p = 5$) is given in [5, Figure 2]. The corresponding submodule lattice is given in [5, Figure 6]. (Note that there is an edge missing in the latter diagram from the node 120 to the node 121.)

Remark 8.2. Note that the first digits of the carry patterns distinguish the terms in the p-filtration of $S^a(V)$ in which the corresponding composition factors lie. Similarly for a fixed choice of the first i digits, the $(i+1)$st digits distinguish the terms in the p^{i+1}-filtration of an appropriate $\nabla_{p^i}(\mu)$.

We are now in a position to prove

Theorem 8.3. *Suppose that $\lambda = (a, 0)$ or $(0, b)$. All non-zero homomorphisms $\nabla(\lambda) \to \nabla(\mu)$ are either composites of Carter-Payne maps (as described in the preceding section) or induced from non-zero homomorphisms $\nabla_p(\tilde{\lambda}) \to \nabla_p(\mu)$.*

In particular, for λ of the form $(a, 0)$ all non-twisted maps are of the form ϕ_L^* or $\phi_R^*\phi_L^*$, and twisted maps correspond (after translation and twisting) to homomorphisms from some $\nabla(c, 0)$ where $(c, 0)$ is any weight lying in the same facet as $(\tilde{\lambda})''$. The case $\lambda = (0, b)$ is similar.

Proof: As the contravariant duals of $\nabla(a, b)^*$ and $L(a, b)^*$ are $\nabla(b, a)$ and $L(b, a)$ respectively, we may assume that $\lambda = (a, 0)$. For there to be a non-zero homomorphism it is clearly necessary that the quotient of $\nabla(\lambda)$ with simple socle $L(\mu)$ must only contain composition factors $L(\tau)$ with $\tau \leq \mu$. (As $\nabla(\lambda)$ is a symmetric power, and hence is multiplicity free, this quotient is well defined.) By the universal property of $\nabla(\mu)$ this condition is also sufficient; we will determine precisely the set of such μ (and show in doing so that the maps are of the form described in the statement of the Theorem).

Suppose that λ lies in the interior of an alcove. For μ to be a possible candidate for a homomorphism, the series of steps in the composition factor algorithm which leads to μ must be such that, at every stage, the immediate descendants of μ^{j+1} occurring above μ^j in the relevant configuration from Figure 19 must be labelled by smaller integers in Figure 19. (Otherwise it is easy to verify that there exists a weight $\tau \geq \mu$ whose carry pattern is above that of μ.) By considering the various configurations given in Figure 19, we see that the only candidates for homomorphisms are those weights for which maps are already known to exist by the results in the preceding section. That completes the proof when λ lies in the interior of an alcove; the two wall cases are similar. $\square$

The reader may find it helpful to consider the example in Figure 20. In this case there are homomorphisms into the induced modules labelled by 000 (identity), 011 (ϕ_L^3), 121 ($\phi_R^2\phi_L^3$), 111 ($\phi_R^1\phi_L^3$), 010 (ϕ_L^2), 110 ($\phi_R^1\phi_L^2$), 100 (twist of identity map from $100''$), 110 (twist of ϕ_L^1 from $100''$), and 120 (twist of twist of identity from $(120'')''$ in $100''$). (Here we are abusing notation by combining carry patterns with our conventions on weights with primes on.) Note that these are precisely the vertices v in the submodule lattice such that all vertices above them in the lattice label weights occurring below the weight labelled by v.

The only class of facets which we have not yet considered are just dominant up alcoves. The argument here is an easier form of Theorem 7.3. Although there is no longer a vertex below the weight λ under consideration, we will use our knowledge of homomorphisms from a weight ν on the horizontal wall below it instead (which was already required as condition (3) in the statement of Theorem 7.3). By the translation arguments in the proof of Theorem 5.1 we know that a necessary condition for the existence of a map from $\nabla(\lambda)$ to $\nabla(\mu)$ is that there is a map from $\nabla(\nu)$ to $\nabla(\delta)$ for some δ on one of the walls in the closure of the alcove containing μ.

Theorem 8.4. *Suppose that ν lies on a just dominant horizontal wall, and we have a composite of Carter-Payne maps from $\nabla(\nu)$ to $\nabla(\delta)$. Let λ lie in an up alcove with ν on the wall below λ and τ the lower vertex in the closure of the down alcove below ν. If we have a weight $\mu \in W.\lambda$ with μ in either of the alcoves adjacent to the wall containing δ then*

$$\mathrm{Hom}(\nabla(\lambda), \nabla(\mu)) \neq 0.$$

When λ is a right-hand just dominant weight the set of such maps is precisely those maps in

$$\{\phi_R^*\phi_L^*, \phi_L^*, \phi_B^1, \phi_L^1\phi_B^1\}$$

which only involve reflections between dominant weights. The case of λ a left-hand just dominant weight is similar.

Proof: We consider the right-hand just dominant case; the other is similar. By Theorem 8.3 the map from $\nabla(\nu)$ to $\nabla(\delta)$ must be of the form ϕ_L^* or $\phi_R^*\phi_L^*$, and hence the various configurations that must be considered are as shown in Figure 21.

By Lemma 7.1 and the local data in Section 3 we see that each of the composites given in the statement of the Theorem is non-zero, and by considering Figure 21 we see that no further maps are required. $\square$

In Theorem 7.5 we constructed a second class of maps for interior weights, coming from twisted maps. Thus we also need to consider the corresponding case for λ in a just dominant up alcove. The untwisted form of the top ∇_p-factor in $\nabla(\lambda)$ will be a symmetric power. For any map from such a symmetric power which is a composite of Carter-Payne maps, it is easy to verify using Figure 21 and the classification of such maps in Theorem 8.3 that the twisted version is actually one of the Carter-Payne maps of the original module considered in the Theorem above. Thus the only remaining problems are the cases arising from twists of maps which are not Carter-Payne composites. These are the subject of the following section.

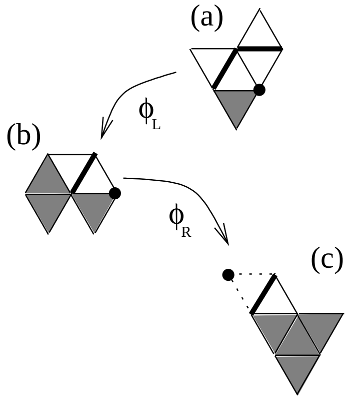

FIGURE 21.

9. Determining homomorphisms III: exceptional maps

Let us review the cases which have been covered so far. If λ is just dominant then all possible homomorphisms have been constructed in Section 8.

If λ is a vertex then $\nabla(\lambda) = \nabla_p(\lambda)$ and the result is clear from (1). For all remaining λ we have used translation arguments to classify all homomorphisms induced from vertex homomorphisms which are either composites of Carter-Payne maps or twisted maps for certain weights. (Note that this includes all weights in the lowest p^2-alcove, where all translations are of the identity map.)

There remains one extra class of possible maps, which we call *exceptional maps*. These maps are those which are scaled versions of the exceptional p-good maps. That is, they are obtained either by translating a twisted map from a vertex which is a smaller exceptional map, or by translating an exceptional map from a vertex. In what follows we consider the case of exceptional maps coming from weights near the right-hand edge of the dominant region — the left-hand case is entirely symmetric. Since the exceptional p-good maps are generically not composites of Carter-Payne maps neither are the scaled versions of these maps. So to construct those maps which cannot be found by factoring through the G_1-head of $\nabla(\lambda)$ we must use some other method.

We first define the notion of two weights λ and ξ being in *exceptional configuration* inductively. A map from $\nabla(\lambda)$ to $\nabla(\xi)$ is then an *exceptional map* if the two weights λ and ξ are in exceptional configuration.

First, the weights labelled by 1 and 3 in Figures 2a(i) and 1 and 2 in 2c(i) are in an *exceptional configuration*. (These are the weights that give the exceptional p-good maps.)

We say that two Steinberg weights λ and ξ are in an *exceptional configuration* if they are in the same G-block and they come from two smaller weights in exceptional configuration. That is we have $\lambda = p^d\lambda_1 + (p^d - 1)\rho$ and $\xi = p^d\xi_1 + (p^d - 1)\rho$ and λ_1 and ξ_1 are in exceptional configuration with $d \in \mathbb{N}$ and λ_1 and ξ_1 both non-Steinberg weights.

We say that two non-Steinberg weights λ and ξ are in an *exceptional configuration* if the Steinberg weight σ_1 directly below λ and σ_2 with $\sigma_2 < \sigma_1$ are in exceptional configuration and ξ is an eligible σ_2 translate. This is illustrated for various cases in Figures 22, 23, 24 and 25 where the weights λ and α and λ and β are in exceptional configuration. These figures represent the generic situation and the weights are related by large p^d-reflections. It is possible for σ_1'' to lie on a wall or be a Steinberg weight. If σ_1'' is a Steinberg weight then it must lie in the interior of larger p^d-alcove and then σ_1 and σ_2 are related by p^d-reflections for some $d > 1$. If σ_1'' is on a wall then it is possible for the p-filtration pattern of $\nabla(\lambda)$ to go into the other p^d-alcove.

If there is no Steinberg weight below λ then we say λ and ξ are in an *exceptional configuration* if there is a homomorphism from $\nabla(\lambda)$ to $\nabla(\xi)$ (these were determined in the previous section) but this is not constructed by a composite of Carter-Payne maps. That is, it can only be obtained by factoring through the G_1-head of $\nabla(\lambda)$.

Thus the exceptional configurations are just scaled versions of the exceptional p-good maps and hence always have the same basic configuration as in Figures 22, 23, 24 and 25 (and the dual versions for the other edge of the dominant region).

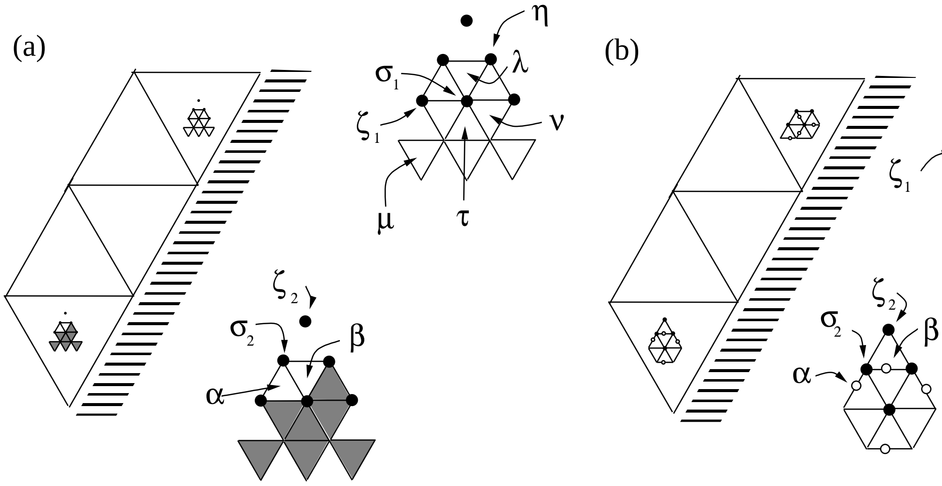

FIGURE 22.

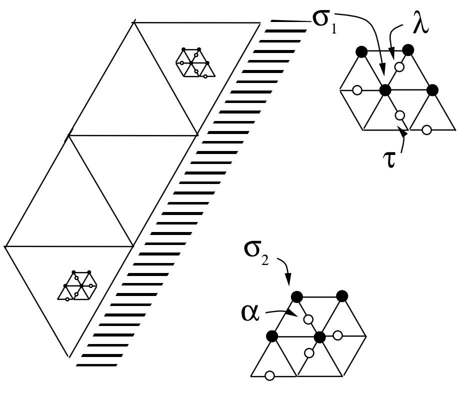

FIGURE 23.

Before we construct these maps, we note

Corollary 9.1. *Assumption 5.2 holds.*

Proof: This follows from Theorems 7.2 and 7.5, and the possible configurations coming from exceptional maps as considered above. □

Thus we have completed the proof of Theorem 5.1, and it only remains to construct the exceptional maps.

Theorem 9.2. *If λ and ξ are in an exceptional configuration then there is a homomorphism from $\nabla(\lambda)$ to $\nabla(\xi)$.*

Proof: If λ (and hence ξ) is p^2-restricted then the only possible exceptional configuration is exactly that which gives the exceptional p-good homomorphisms. Thus the theorem is true in this case. We have also proved this theorem (essentially by definition) if $\nabla(\lambda)$ is as in Section 8.

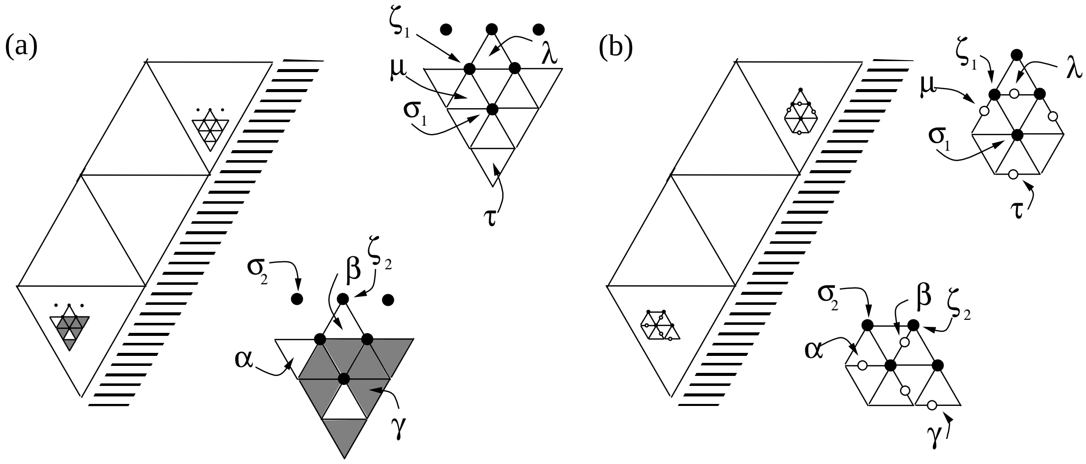

FIGURE 24.

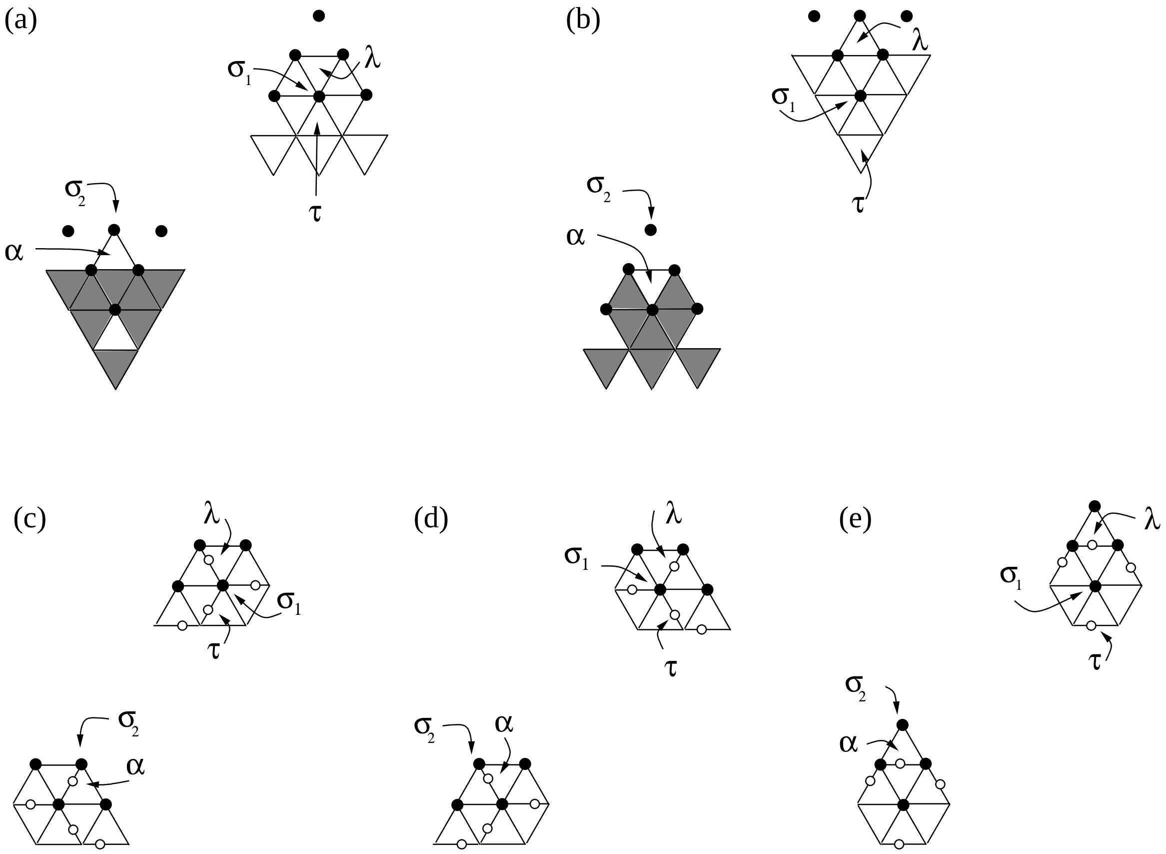

FIGURE 25.

We now take the statement of the theorem as our inductive hypothesis for the rest of this section.

By left/right symmetry and the results from Section 8, it remains to consider the case where λ and ξ are in an exceptional configuration with σ_1 below λ and we have a homomorphism from $\nabla(\sigma_1)$ to $\nabla(\sigma_2)$ with eligible σ_2-translates α and β as shown. (Note that for the remainder of this section we have adopted a non-standard labelling of the weights; in particular τ is the weight usually denoted $\tilde{\lambda}$.)

By inspection we see that there are several possible type of cases that can occur, as illustrated in the figures. The lower alcove and lefthand wall case have a similar weight configuration and the upper alcove and horizontal wall case also has a similar configuration.

In all the cases there is an obvious map from $\nabla(\lambda)$ to $\nabla(\alpha)$. It is the one induced by the map from $\nabla_p(\tau)$ to $\nabla_p(\alpha)$. It remains to construct the map from $\nabla(\lambda)$ to $\nabla(\beta)$ as in Figures 22 and 24.

Lemma 9.3. *Assume* σ_1, σ_2 , ζ_1 *and* ζ_2 *are as in Figures 22 and 24. Then* $\mathrm{Hom}(\nabla(\sigma_1), \nabla(\sigma_2)) \neq 0$ *implies that* $\mathrm{Hom}(\nabla(\zeta_1), \nabla(\zeta_2)) \neq 0$.

Proof: Clearly, by (1), this is equivalent to proving that $\mathrm{Hom}(\nabla(\sigma_1''), \nabla(\sigma_2'')) \neq 0$ implies that $\mathrm{Hom}(\nabla(\zeta_1''), \nabla(\zeta_2'')) \neq 0$.

Now ζ_1'' and σ_1'' are so close together that one must be in the closure of the facet containing the other. Most of time they will be in the same facet. Firstly if ζ_1'' is in the closure of the facet containing σ_1'' then the translation principle gives the desired result. So we may assume that σ_1'' is in the closure of the facet containing ζ_1''. We now claim that if σ_1'' and σ_2'' are in an exceptional configuration then so are ζ_1'' and ζ_2'' except for one case (where we will construct the map directly). Then our inductive hypothesis gives the desired result.

To see the claim we observe that if ζ_1'' is not in the same facet as σ_1'' then σ_1'' is either a Steinberg weight and ζ_1'' is on a wall or σ_1''' is on a wall and ζ_1'' is in an up alcove. We now look at Figures 22 and 24 and we see that ζ_1'' and ζ_2'' are in an exceptional configuration. The only thing that can go wrong is if σ_1 is on a left-hand wall (is α_1-singular) and this is the same wall used to get the exceptional configuration. Although ζ_1'' and ζ_2'' are not then in an exceptional configuration, we can construct the map from $\nabla(\zeta_1'')$ to $\nabla(\zeta_2'')$ directly: it is the map $\phi_L^d \phi_R^d$ which is non-zero using Lemma 7.1. □

Before returning to the proof of Theorem 9.2, we review the general theory that we will require. Suppose we have some module category with non-trivial maps $\phi : A \to B$ and $\psi : C \to D$ and extensions

$$0 \to A \to E \to C \to 0$$
$$0 \to B \to F \to D \to 0.$$

We want the maps ϕ and ψ to glue together to give a map from $E \to F$. This happens exactly when the push out of E is the pull back of F. That is we need $H := \phi E \cong F\psi$ (using the notation of MacLane[20, chapter 3]). Then the following diagram commutes

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \longrightarrow & E & \longrightarrow & C & \longrightarrow & 0 \\ & & \phi\big\downarrow & & \bar{\phi}\big\downarrow & & Id\big\downarrow & & \\ 0 & \longrightarrow & B & \longrightarrow & H & \longrightarrow & C & \longrightarrow & 0 \\ & & Id\big\downarrow & & \bar{\psi}\big\downarrow & & \psi\big\downarrow & & \\ 0 & \longrightarrow & B & \longrightarrow & F & \longrightarrow & D & \longrightarrow & 0 \end{array}$$

and the composite $\theta : E \to H \to F$ is non-zero. Moreover its image has the following short exact sequence

$$0 \to \operatorname{im} \phi \to \operatorname{im} \theta \to \operatorname{im} \psi \to 0. \tag{10}$$

We first consider the down alcove and left hand wall case (depicted in Figure 22) and start off by assuming that μ and ν (in $W.\lambda$) in the figure are in different ∇_p-classes.

In this case we take $A = \nabla_p(\mu)$, $B = \nabla_p(\beta)$, $C = \nabla_p(\tau)$ and $D = \nabla_p(\alpha)$. We know that there is an extension of A by C at the top of the p-filtration of $\nabla(\lambda)$ (by our assumption on μ and ν), and of B by D at the bottom of the p-filtration for $\nabla(\beta)$. We denote these extensions by E and F respectively. Thus it remains to show that $\phi E \cong F\psi$.

First consider the diagram defining the pushout E' of E along ϕ:

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \nabla_p(\mu) & \longrightarrow & E & \longrightarrow & \nabla_p(\tau) & \longrightarrow & 0\\
 & & \phi\big\downarrow & & \bar{\phi}\big\downarrow & & Id\big\downarrow & & \\
0 & \longrightarrow & \nabla_p(\beta) & \longrightarrow & E' & \longrightarrow & \nabla_p(\tau) & \longrightarrow & 0.
\end{array}$$

This is exactly the definition of the connecting homomorphism ∂ in the long exact sequence obtained by applying $\mathrm{Hom}(-, \nabla_p(\beta))$ to the defining short exact sequence for E:

$$0 \to \mathrm{Hom}(\nabla_p(\tau), \nabla_p(\beta)) \to \mathrm{Hom}(E, \nabla_p(\beta)) \to \mathrm{Hom}(\nabla_p(\mu), \nabla_p(\beta)) \xrightarrow{\partial} \mathrm{Ext}^1(\nabla_p(\tau), \nabla_p(\beta)).$$

As $\nabla_p(\tau)$ and $\nabla_p(\beta)$ are in different ∇_p-classes, the first Hom-space is zero. Further, E sits at the top of a filtration of $\nabla(\lambda)$, and hence has simple head in the ∇_p-class of $\nabla_p(\tau)$. Therefore the second Hom-space is also zero, and ∂ is an embedding. By our inductive hypothesis and (1) we have $\mathrm{Hom}(\nabla_p(\mu), \nabla_p(\beta)) \cong k$.

Consider $\mathrm{Ext}^1(\nabla_p(\tau), \nabla_p(\beta))$. By the results in [21, Lemma 4.2] and [22, Proposition 3.3.2] (reproduced in [21, Proposition 4.1]) we have

$$\begin{aligned}
\mathrm{Ext}^1(\nabla_p(\tau), \nabla_p(\beta)) &\cong \mathrm{Hom}(\nabla(\tau''), \nabla(\beta'') \otimes \nabla(0,1))\\
&\cong \mathrm{Hom}(\nabla(\sigma_1''), T_{\beta''}^{\sigma_2''} \nabla(\beta'')).
\end{aligned} \tag{11}$$

This final Hom-space is one-dimensional (by our inductive hypothesis) either because $T_{\beta''}^{\sigma_2''} \nabla(\beta'') \cong \nabla(\sigma_2'')$ (if β'' is not on a wall) or because

$$\mathrm{Hom}(\nabla(\sigma_1''), T_{\beta''}^{\sigma_2''} \nabla(\beta'')) \cong \mathrm{Hom}(T_{\sigma_2''}^{\beta''} \nabla(\sigma_1''), \nabla(\beta'')) \cong \mathrm{Hom}(\nabla(\epsilon), \nabla(\beta''))$$

for some ϵ on the wall containing μ''.

Thus any non-zero homomorphism from $\nabla_p(\mu)$ to $\nabla_p(\beta)$ pushes out to the unique non-zero extension

$$0 \to \nabla_p(\beta) \to E' \to \nabla_p(\tau) \to 0$$

as ∂ is an isomorphism.

The case for the pullback is similar. We have the diagram

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \nabla_p(\beta) & \longrightarrow & F' & \longrightarrow & \nabla_p(\tau) & \longrightarrow & 0\\
 & & Id\big\downarrow & & \bar{\psi}\big\downarrow & & \psi\big\downarrow & & \\
0 & \longrightarrow & \nabla_p(\beta) & \longrightarrow & F & \longrightarrow & \nabla_p(\alpha) & \longrightarrow & 0
\end{array}$$

and corresponding connecting homomorphism

$$0 \to \mathrm{Hom}(\nabla_p(\tau), \nabla_p(\beta)) \to \mathrm{Hom}(\nabla_p(\tau), F) \to \mathrm{Hom}(\nabla_p(\tau), \nabla_p(\alpha)) \xrightarrow{\partial} \mathrm{Ext}^1(\nabla_p(\tau), \nabla_p(\beta)).$$

By arguments as above the first pair of Hom-spaces are both zero, and ∂ becomes an isomorphism from k to k. Thus any non-zero homomorphism from $\nabla_p(\tau)$ to $\nabla_p(\alpha)$ pulls back to the unique non-zero extension $F' \cong E'$ and the composite map $\bar{\psi} \circ \bar{\phi}$ from E to F is non-zero by (10). This gives the required homomorphism

$$\nabla(\lambda) \twoheadrightarrow E \xrightarrow{\bar{\psi}\circ\bar{\phi}} F \hookrightarrow \nabla(\beta).$$

When μ and ν are in the same ∇_p-class we have to modify the above argument, as the extension E above does not sit at the top of the p-filtration of $\nabla(\lambda)$. This problem can be rectified by taking A to be the extension of $\nabla_p(\mu)$ by $\nabla_p(\nu)$ occurring in the p-filtration for $\nabla(\lambda)$, and E to be the extension of A by $\nabla_p(\tau)$ occurring at the top of this p-filtration. Note that A has been described in Lemma 6.2; using this we deduce that there exists a weight η on the same wall as σ_1'' such that

$$\begin{aligned}
\mathrm{Hom}(A, \nabla_p(\beta)) &\cong \mathrm{Hom}(T_\eta^{\mu''} \nabla(\eta), \nabla(\beta''))\\
&\cong \mathrm{Hom}(\nabla(\eta), T_{\mu''}^{\eta} \nabla(\beta''))\\
&\cong \mathrm{Hom}(\nabla(\eta), \nabla(\rho))
\end{aligned}$$

where ρ is a weight on the same wall as σ_2''. By the translation principle this latter Hom-space in one-dimensional. The argument now goes through as before, with the new choices for A and E.

The argument for Figures 23 and 24 is the same in the generic case where α and γ are in different ∇_p classes. When α and γ are in the same ∇_p-class then we need to construct a map from the extension $\nabla_p(\mu)$ by $\nabla_p(\tau)$ to the submodule of $\nabla(\beta)$ which has uniserial p-filtration by $\nabla_p(\beta)$, $\nabla_p(\gamma)$ and $\nabla_p(\alpha)$. This can be constructed in the similar way as for the down alcove case, using Lemma 6.2 as before.

As every exceptional configuration is of one of the forms depicted in Figures 22, 23, 24 or 25 (or their left-handed analogues) we have constructed all remaining maps, which completes our induction for this section and our classification of homomorphisms.

10. Examples

In this section we give some examples for $p = 3$ of how we may apply the various theorems in the paper to give all the possible homomorphisms starting at a particular induced module.

In the diagrams an alcove is shaded if it contains a weight that is a composition factor of the initial induced module, $\nabla(\lambda)$. This shading is dark if there is a homomorphism from $\nabla(\lambda)$ to the induced module corresponding to the weight in the alcove. That is $\mathrm{Hom}(\nabla(\lambda), \nabla(\mu)) \cong k$ where μ is in the W_p-orbit of λ and inside the shaded alcove. The lightly shaded alcoves are those which contain weights $\mu \in W.\lambda$ which label composition factors but for which there is no non-zero homomorphism from $\nabla(\lambda) \to \nabla(\mu)$.

In Figure 26 we determine all non-zero maps from $\nabla(10,1)$, $\nabla(33,6)$ and $\nabla(102,21)$. We start with $\nabla(10,1)$. The weight $(10,1)$ is in the interior of a just dominant up alcove, and $(10,0)$ lies on the horizontal wall just below it. Thus we will apply Theorem 8.4 with this pair of weights playing the role of λ and ν respectively. Using Theorem 8.3, we only get two non-zero homomorphisms starting from $\nabla(10,0)$: the identity map and the map $\nabla(10,0) \to \nabla(6,1)$. We then apply Theorem 8.4 to obtain the four darkly shaded alcoves shown. (Alternatively it is not hard to see that the only homomorphisms we get in this case are the p-good maps.) Thus we have (after transforming the alcoves back into weights)

$$\Theta_{(10,1)} = \{(10,1), (9,0), (6,3), (7,1)\}.$$

We can now determine the result for $\nabla(33,6)$. The vertex immediately below $(33,6)$ is $3(10,1) + (2,2)$. So we follow the procedure in Section 7 and use Theorem 7.2 to obtain the dark shaded alcoves as shown. That is we take the set $p\Theta_{(10,1)} + (p-1, p-1)$ and then shade in the alcoves according to Figure 10. We get

$$\Theta_{(33,6)} = \{(33,6), (34,4), (33,3), (31,7), (30,6), (31,4), (28,1), (19,13), (18,12), (19,10), (21,9), (22,4), (24,4)\}.$$

We then can determine the result for $\nabla(102,21)$ using Theorem 7.2 and Theorem 7.5. Rather than give the complete list of weights here, we depict them graphically in Figure 26. In these cases the only types of maps we get are composites of Carter–Payne maps and twisted maps.

To illustrate a case involving walls we include Figure 27, where we determine all maps from $\nabla(11,0)$, $\nabla(37,4)$ and $\nabla(119,16)$. The reader may verify that the procedure in the preceding sections gives rise to the set of maps shown. The case of $\nabla(37,4)$ is an example where the vertex is not p^2-regular. These examples give all the different types of maps. Note that the exceptional p-good map $\nabla(11,0) \to \nabla(6,1)$ gives rise to the exceptional configurations $((37,4),(18,3))$ and $((37,4),(22,1))$ which in turn give rise to the exceptional configurations $((119,16),(54,9))$, $((119,16),(66,3))$ and $((119,16),(70,1))$.

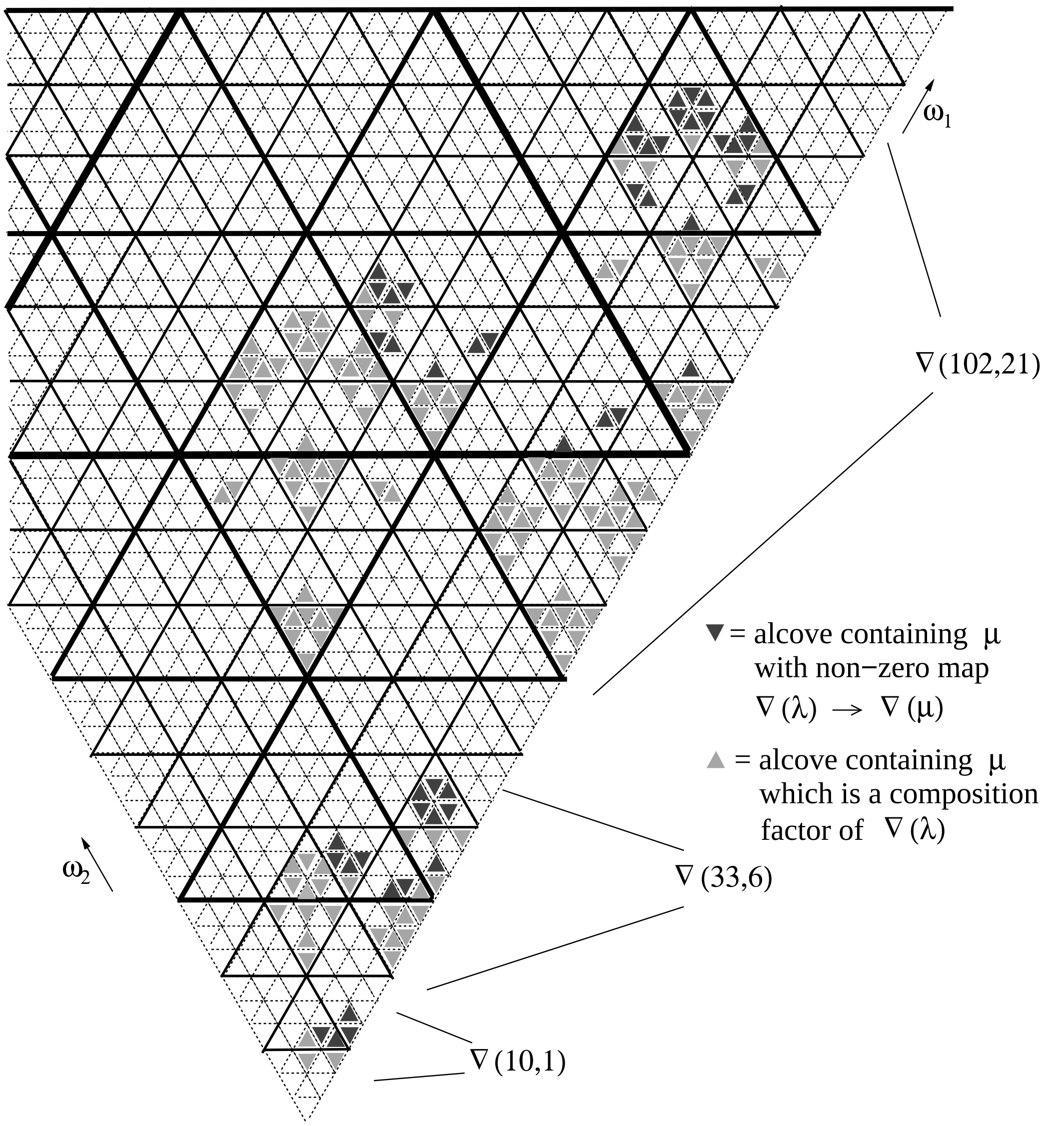

FIGURE 26.

11. HOMOMORPHISMS FOR THE SYMMETRIC GROUP

In this section we will show how the results for SL_3 can be used to classify homomorphisms between certain Specht modules for the symmetric group.

Let Σ_d be the symmetric group on d letters. For each partition λ of d we can explicitly define a Specht module S^λ for $k\Sigma_d$. (See for example [15, Chapter 3].)

Partitions with at most n parts can also be used to label representations of GL_n. The representation theory of this group is essentially the same as that for SL_n; in particular given n part partitions λ and μ of d we have that

$$\mathrm{Hom}_{\mathrm{GL}_n}(\nabla(\lambda), \nabla(\mu)) \cong \mathrm{Hom}_{\mathrm{SL}_n}(\nabla(\bar{\lambda}), \nabla(\bar{\mu}))$$

where for a partition $\tau = (\tau_1, \tau_2, \ldots, \tau_n)$ we set $\bar{\tau} = (\tau_1 - \tau_2, \tau_2 - \tau_3, \ldots, \tau_{n-1} - \tau_n)$. Those representations of GL_n labelled by partitions of d are also representations of the corresponding Schur algebra $S(n, d)$, and again the Hom-spaces are unchanged.

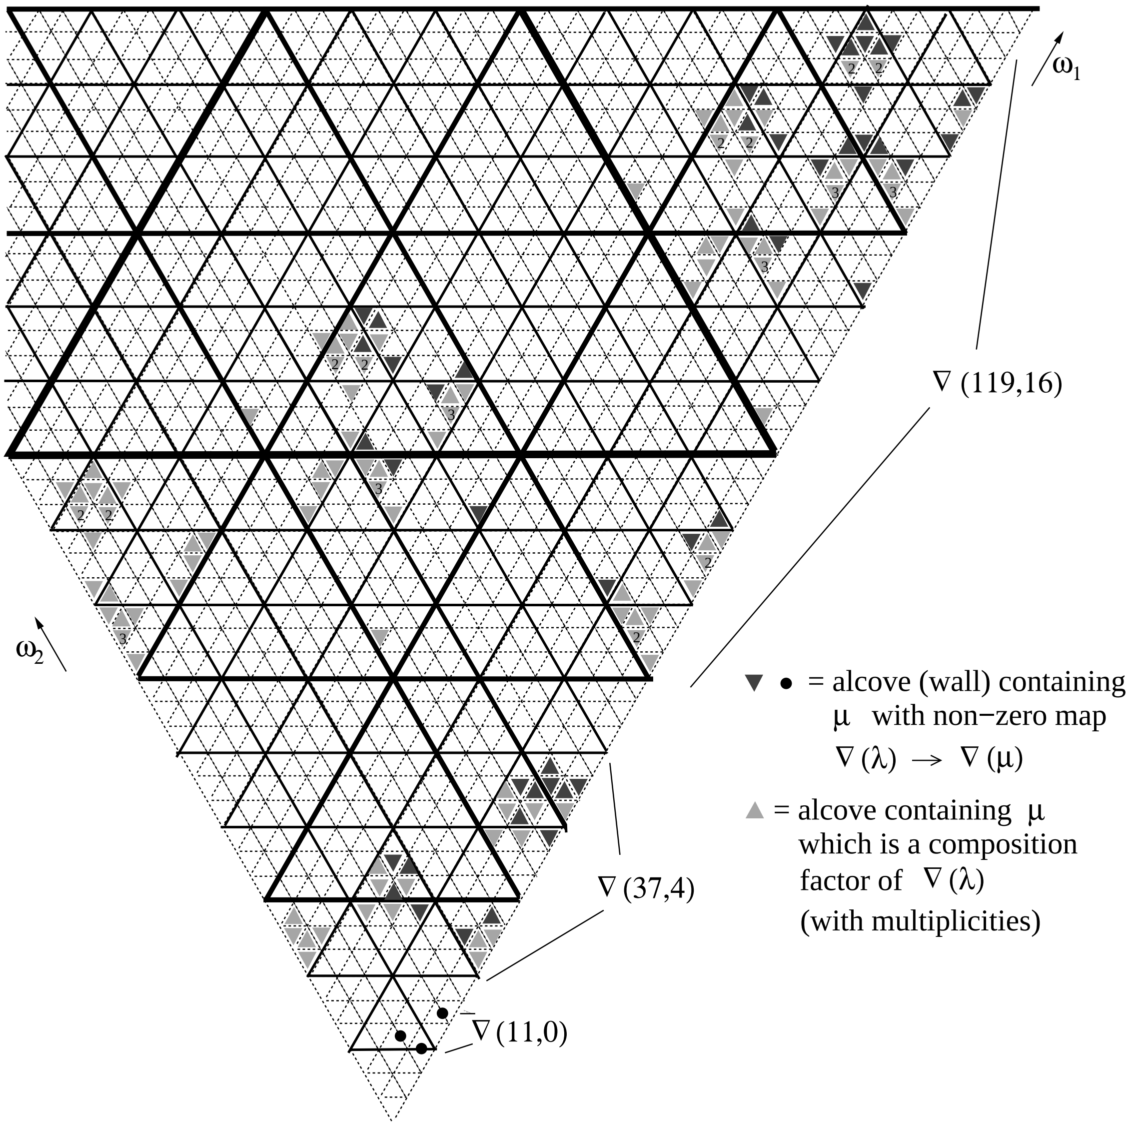

FIGURE 27.

The representation theory of the general linear group and the symmetric group are related. In particular we have the following theorem of Carter and Lusztig

Theorem 11.1 ([3, Theorem 3.7]). *If $p \geq 3$ and λ and μ are partitions of d with at most n parts then*

$$\mathrm{Hom}_{\mathrm{GL}_n}(\nabla(\lambda), \nabla(\mu)) \cong \mathrm{Hom}_{k\Sigma_d}(S^\lambda, S^\mu).$$

We thus have

$$\mathrm{Hom}_{k\Sigma_d}(S^\lambda, S^\mu) \cong \mathrm{Hom}_{\mathrm{SL}_n}(\nabla(\bar{\lambda}), \nabla(\bar{\mu})).$$

Combining this with our earlier results gives

Corollary 11.2. *Suppose that $p \geq 3$, $d > 0$ and λ and μ are three part partitions of d. Then $\mathrm{Hom}_{k\Sigma_d}(S^\lambda, S^\mu)$ can be determined from the results in the preceding sections.*

12. Levi factors and a theorem of Fayers and Lyle

In this final section we will show how our SL_3 results can be applied to determine certain Hom-spaces in higher rank cases, and show how the result of Carter and Lusztig relates work of Donkin to a tensor product theorem for Hom-spaces due to Fayers and Lyle [13]. In each case, our main tool will be a theorem of Donkin concerning certain Hom-spaces for Levi factors of a reductive group G, which we begin by recalling.

Let G be a reductive algebraic group as in Section 2. Given a subset I of the simple roots S, we may also consider representations of the corresponding Levi factor G_I of G. To distinguish such modules from those for the original group we shall use the subscript I. We have the following result of Donkin (see [11, (4.3) Corollary]).

Theorem 12.1. *If λ and $\mu \in X^+$ are such that $\lambda - \mu \in \mathbb{Z}I$ then*

$$\mathrm{Ext}^i_G(\nabla(\lambda), \nabla(\mu)) \cong \mathrm{Ext}^i_{G_I}(\nabla_I(\lambda), \nabla_I(\mu))$$

for all $i \geq 0$.

Now let G be GL_n with the usual choice of simple roots as in [17, II, 1.21]. Suppose that λ and μ are two partitions of d with at most n parts (which we regard as elements of X^+) and that we can write $\lambda = (\lambda(1), \ldots, \lambda(t))$ and $\mu = (\mu(1), \ldots, \mu(t))$, where for each i the elements $\lambda(i)$ and $\mu(i)$ are both partitions of d_i into at most n_i parts. Then it is easy to verify that $\lambda - \mu$ is an element of $\mathbb{Z}I$, where I is the subset of S giving rise to the Levi factor $G_I \cong \mathrm{GL}_{n_1} \times \cdots \times \mathrm{GL}_{n_t}$.

As noted in [7, (4)] we have that $\nabla_I(\lambda) \cong \nabla_1(\lambda(1)) \boxtimes \cdots \boxtimes \nabla_t(\lambda(t))$ (and similarly for $\nabla_I(\mu)$), where $\nabla_i(\lambda(i))$ is the induced module of weight $\lambda(i)$ for GL_{n_i}. Combining this with Theorem 12.1 and the Künneth formula we obtain

Corollary 12.2. *Let $m \in \mathbb{N}$, $\lambda = (\lambda(1), \ldots, \lambda(t))$ and $\mu = (\mu(1), \ldots, \mu(t))$ be as above. Then we have*

$$\begin{aligned}&\mathrm{Ext}^m_{G_I}(\nabla(\lambda), \nabla(\mu))\\ &\qquad\cong \bigoplus_{m = i_1 + \cdots + i_t} \mathrm{Ext}^{i_1}_{\mathrm{GL}_{n_1}}(\nabla_1(\lambda(1)), \nabla_1(\mu(1))) \otimes \cdots \otimes \mathrm{Ext}^{i_t}_{\mathrm{GL}_{n_t}}(\nabla_t(\lambda(t)), \nabla_t(\mu(t))).\end{aligned}$$

Corollary 12.3. *If λ and $\mu \in X^+$ differ only by an element of $\mathbb{Z}I$, I a subset of the simple roots and I can be realised as a root system whose largest connected component is of type A_2, then the results of the preceding sections (together with the known result for type A_1) can be applied to determine $\mathrm{Hom}_G(\nabla(\lambda), \nabla(\mu))$. In particular, the Hom-space is at most one-dimensional.*

Combining Theorem 11.1 and Corollary 12.2 we obtain

Corollary 12.4. *Let λ and μ be as above and $p \geq 3$. We have*

$$\mathrm{Hom}_{k\Sigma_d}(S^\lambda, S^\mu) \cong \mathrm{Hom}\, k\Sigma_{d_1}(S^{\lambda(1)}, S^{\mu(1)}) \otimes \cdots \otimes \mathrm{Hom}\, k\Sigma_{d_t}(S^{\lambda(t)}, S^{\mu(t)}).$$

This is the theorem of Fayers and Lyle [13, Theorem 2.2]. We also note here that Donkin [9, Section 10, Proposition 4] uses Schur functors to prove a generalised version of the above result for Ext-groups (in small degree) for $p \geq 3$.

Example 12.5. Let $p = 5$, $G = \mathrm{GL}_7(k)$, $\lambda = (90, 63, 40, 8, 4, 4, 4)$ and $\mu = (90, 54, 39, 18, 4, 4, 4)$. We denote the simple roots for G by $\alpha_i = \epsilon_i - \epsilon_{i+1}$, where $\{\epsilon_i\}$ is the usual basis for X. We have $\lambda - \mu = (0, 9, 1, -10, 0, 0, 0) = 9\alpha_2 + 11\alpha_3$. Thus

$$\begin{aligned}\mathrm{Hom}_G(\nabla(\lambda), \nabla(\mu)) &\cong \mathrm{Hom}_{\mathrm{GL}_3(k)}(\nabla(63, 40, 8), \nabla(54, 39, 18))\\ &\cong \mathrm{Hom}_{\mathrm{SL}_3(k)}(\nabla(23, 32), \nabla(15, 21)) \cong k.\end{aligned}$$

We also get $\mathrm{Hom}_{k\Sigma_{213}}(S^\lambda, S^\mu) \cong k$

Acknowledgements

The authors would like to thank Øyvind Solberg for some helpful suggestions used in Section 9. We would also like to thank the organisers of the LMS Durham Symposium: Representations of Finite Groups and Related Algebras, and of the Conference on Representation theory for Algebraic Groups at Aarhus University, where part of this work was carried out.

References

1. H. H. Andersen, *On the structure of Weyl modules*, Math. Z. **170** (1980), 1–14.
2. ———, *p-filtrations and the Steinberg module*, J. Algebra **244** (2001), 664–683.
3. R. W. Carter and G. Lusztig, *On the modular representations of the general linear and symmetric groups*, Math. Z. **136** (1974), 193–242.
4. R. W. Carter and M. T. J. Payne, *On homomorphisms between Weyl modules and Specht modules*, Math. Proc. Camb. Phil. Soc. **87** (1980), 419–425.
5. A. G. Cox, *Decomposition numbers for distant Weyl modules*, J. Algebra **243** (2001), 448–472.
6. A. G. Cox and K. Erdmann, *On* Ext^2 *between Weyl modules for quantum* GL_n, Math. Proc. Camb. Phil. Soc. **128** (2000), 441–463.
7. S. Donkin, *A note on decomposition numbers of general linear groups and symmetric groups*, Math. Proc. Camb. Phil. Soc **97** (1985), 473–488.
8. ———, *On tilting modules for algebraic groups*, Math. Z. **212** (1993), 39–60.
9. ———, *Tilting modules for algebraic groups and finite dimensional algebras*, preprint, 2003.
10. S. Doty, *Submodules of symmetric powers of the natural module for* GL_n, Contemp. Math. **88** (1989), 185–191.
11. K. Erdmann, Ext^1 *for Weyl modules of* $\mathrm{SL}_2(K)$, Math. Z. **218** (1995), 447–459.
12. K. Erdmann and A. Henke, *On Schur algebras, Ringel duality and symmetric groups*, J. Pure Appl. Algebra **169** (2002), 175–199.
13. M. Fayers and S. Lyle, *Row and column removal theorems for homomorphisms between Specht modules*, J. Pure Appl. Algebra **185** (2003), 147–164.
14. J. W. Franklin, *Homomorphisms between Verma modules and Weyl modules in characteristic* p, Ph.D. thesis, Warwick, 1981.
15. G. D. James and A. Kerber, *The representation theory of the Symmetric group*, Encyclopedia of Mathematics and its Applications, vol. 16, Addison-Wesley, 1981.
16. J. C. Jantzen, *Darstellungen halbeinfacher Gruppen und ihrer Frobenius-Kerne*, J. reine angew. Math. **317** (1980), 157–199.
17. ———, *Representations of algebraic groups*, Academic Press, 1987.
18. M. Koppinen, *Homomorphisms between neighbouring Weyl modules*, J. Algebra **103** (1986), 302–319.
19. K. Kühne-Hausmann, *Zur Untermodulstruktur der Weylmoduln für* SL_3, Bonner Math. Schriften **162** (1985).
20. S. MacLane, *Homology*, Springer, 1963.
21. A. E. Parker, *The global dimension of Schur algebras for* GL_2 *and* GL_3, J. Algebra **241** (2001), 340–378.
22. S. el B. Yehia, *On a question of Verma about the indecomposable representations of algebraic groups and of their Lie algebras*, Ph.D. thesis, London, 1992.

Centre for Mathematical Science, City University, Northampton Square, London, EC1V 0HB, England.

E-mail address: A.G.Cox@city.ac.uk

School of Mathematics and Statistics F07, University of Sydney, NSW 2006, Australia.

E-mail address: alisonp@maths.usyd.edu.au