An analogue of Jeu de taquin for Littelmann's crystal paths

Marc A. A. van Leeuwen

Université de Poitiers, Département de Mathématiques,
UFR Sciences SP2MI, Téléport 2, BP 30179, 86962 Futuroscope Chasseneuil Cedex, France
maavl@math.univ-poitiers.fr
URL: http://www-math.univ-poitiers.fr/~maavl/

ABSTRACT

Littelmann has given a combinatorial model for the characters of representations of semisimple Lie algebras, in terms of certain paths traced in the space of (rational) weights. From it, a description of the decomposition of tensor products can be derived that generalises the Littlewood-Richardson rule (the latter is valid in type A_n only). We present a new combinatorial construction that expresses in a bijective manner the symmetry of the tensor product in this path model. In type A_n, where there is a correspondence between paths and skew tableaux, this construction is equivalent to Schützenberger's *jeu de taquin*; in the general case the construction retains its most crucial properties of symmetry and confluence.

1991 Mathematics Subject Classification: 05E10.
Keywords and Phrases: Jeu de taquin, Littlewood-Richardson rule, crystal graphs, path crystals.

§1. Introduction.

In this note we wish to present a simple construction that appears to arise naturally in the context of Littelmann's paths. Indeed, we found it while trying to formulate an answer to a question asked (by Alain Lascoux) during a lecture by Littelmann on this subject at the Séminaire Lotharingien; the question was whether one could exhibit combinatorially the symmetry of the tensor product in the formula given, in terms of paths, for tensor product decompositions.

This paper is organised as follows. After recalling some of Littelmann's notions in §2, we analyse the symmetry of the traditional Littlewood-Richardson rule in §3, and translate the procedure that exhibits the symmetry (which is essentially jeu de taquin) into the language of paths. Then in §4 we extend this construction successively to two broader classes of paths, with instances for other types of groups than A_n, namely the classes of m-paths (built from the path models of minuscule representations) and of ψ-paths (incorporating also the path models of quasi-minuscule representations); the latter removes any restrictions on the type of the group or the representation. Finally in §5 we give a construction in the context of arbitrary piecewise linear paths that generalises the earlier constructions; however, we have not (yet) established the essential connection with Littelmann's root operations for this general construction.

§2. Notations used.

We shall assume without explicit reference the notations and results of [Litt2]. We mention in particular the following notations. We denote by X the weight lattice of a complex Lie algebra $\mathfrak{g}$, that for simplicity we shall assume to be finite dimensional and reductive, and by Π the set of piecewise linear paths in the space $X_{\mathbf{Q}} = X \otimes_{\mathbf{Z}} \mathbf{Q}$ of rational weights. All paths are parametrised by the interval $[0,1] \subseteq \mathbf{Q}$, and start at 0, so that $\pi(0) = 0$ for all $\pi \in \Pi$, while $\pi(1)$ denotes the end point of π. For $\mu \in X$, the straight path from 0 to μ is denoted by π_μ. The reverse or dual path of $\pi \in \Pi$ is denoted by π^*, and $\pi * \pi'$ denotes the concatenation of two paths. The set of paths π such that $\pi(t)$ is dominant for all $t \in [0,1]$, and such that $\pi(1) \in X$ (i.e., it is an integral dominant weight), is denoted by $\mathcal{P}^+$. The root operators e_α and f_α formally act on the free $\mathbf{Z}$-module $\mathbf{Z}\Pi$, but since the image of every generator $\pi \in \Pi$ of that $\mathbf{Z}$-module is either another such generator or 0, we shall consider the root operators as maps $\Pi \to \Pi \cup \{0\}$. The subset of Π reachable from some $\pi \in \mathcal{P}^+$ by repeated application of root operators is denoted by B_π. To avoid too deeply nested subscripts we shall write e_i and f_i for the root operators e_{α_i} and f_{α_i}. For a weight λ, a path π and a simple root α, we shall write $\lfloor \lambda + \pi \rfloor_\alpha$ for the number $\min_{t\in[0,1]} \langle \lambda + \pi(t), \alpha^\vee \rangle$, or simply $\lfloor \pi \rfloor_\alpha$ if $\lambda = 0$; the path π is called λ-dominant if $\lfloor \lambda + \pi \rfloor_\alpha \geq 0$ for all simple roots α.

§3. Paths in type A_{n-1}, tableaux, and jeu de taquin.

We shall first consider the correspondence between paths in type A_{n-1} and Young tableaux, and the connection between the Littlewood-Richardson rule in terms of paths and the classical one. Then we shall consider the question of symmetry of these rules with respect to the order of the tensorands.

3.1. *Paths and Littlewood-Richardson tableaux.*

Let us first recall the well known correspondence between partitions with at most n parts and dominant integral weights for $\mathfrak{g} = \mathfrak{gl}_n$. Let $\mathfrak{h} \subseteq \mathfrak{gl}_n$ be the Cartan subalgebra consisting of diagonal matrices, and $\mathfrak{b} \subseteq \mathfrak{gl}_n$ the Borel subalgebgra of upper triangular matrices. For $i = 1, \ldots, n$ let $\varepsilon_i \in \mathfrak{h}^*$ be the weight that takes the (i, i) diagonal entry; then the set of simple roots with respect to $\mathfrak{b}$ is $\{\, \alpha_i \mid i = 1, \ldots, n-1 \,\}$ where $\alpha_i = \varepsilon_i - \varepsilon_{i+1}$, and the set of fundamental weights is $\{\, \omega_i \mid i = 1, \ldots, n-1 \,\}$ where $\omega_i = \sum_{j \leq i} \varepsilon_j$. We shall identify any vector $\lambda = (\lambda_1, \ldots, \lambda_n) \in \mathbf{Z}^n$ with the weight $\sum_{i=1}^n \lambda_i \varepsilon_i$; since one has $\langle \lambda, \alpha_i^\vee \rangle = \lambda_i - \lambda_{i+1}$, it follows that this is a dominant integral weight if and only if λ is a partition.

Now we shall define a correspondence between the set Tab_μ of semistandard Young tableaux of shape μ, and the set B_π for a specific path $\pi = \pi_{\mathrm{c}}(\mu) \in \mathcal{P}^+$, which is determined as follows. Write μ as a sum of terms ω_i (so the term ω_i is repeated $\mu_i - \mu_{i+1}$ times) ordered by (weakly) increasing index i, and then replace each ω_i by its expression $\sum_{j \leq i} \varepsilon_j$ as a sum of terms ε_j, again ordered by increasing index. The path $\pi_{\mathrm{c}}(\mu)$ is obtained from the resulting sum by replacing each of the $|\mu|$ terms of the form ε_j by the corresponding path π_{ε_j}, and addition by concatenation. We shall call any path of the form $\pi_{\varepsilon_{j_1}} * \cdots * \pi_{\varepsilon_{j_l}}$ an ε-path of length l; such paths are characterised by the sequence $j_1, \ldots, j_l$ of indices. If f_i is applied to an ε-path and the result is not 0, then it changes one segment π_{ε_i} into $\pi_{\varepsilon_{i+1}}$. Therefore any path in $B_{\pi_{\mathrm{c}}(\mu)}$ is an ε-path of length $|\mu|$. Inserting its sequence of indices into a Young diagram of shape μ, proceeding by columns from right to left and from top to bottom within each column, one obtains a tableau, and it can be shown that this defines a bijection from $B_{\pi_{\mathrm{c}}(\mu)}$ to Tab_μ. Note that a subsequence of segments contributing to any one column of length i of the tableau stems from the sequence of segments π_{ε_j} corresponding (before application of the f_α) to one term ω_i in the sum for μ.

Remark. In fact one could have used instead of $\pi_{\mathrm{c}}(\mu)$ another path $\bar\pi_{\mathrm{c}}(\mu)$, formed by concatenating straight line paths π_{ω_i} corresponding to the terms ω_i in the first sum for μ. One then obtains a bijection between $B_{\bar\pi_{\mathrm{c}}(\mu)}$ and the *same* set Tab_μ of tableaux: every application of f_α that does not yield 0 transforms one path segment into another straight segment, of the form π_{ε_I} with I an i-element subset of $\{1, \ldots, n\}$ and $\varepsilon_I = \sum_{j \in I} \varepsilon_j$; that segment corresponds to a column with set of entries I in the Young tableau. Remarkably, Tab_μ can even be used to describe in a direct way many other sets $B_{\pi'}$, where $\pi' \in \mathcal{P}^+$ is an ε-path. For instance, if $\pi_{\mathrm{r}}(\mu)$ denotes the ε-path corresponding to the expression $\mu = \sum_{i=1}^n \mu_i \varepsilon_i$ (i.e., with weakly increasing indices), then each path in $B_{\pi_{\mathrm{r}}(\mu)}$ is an ε-path whose sequence of indices is obtained by listing the entries of a tableau in Tab_μ *by rows* from right to left. We have given prominence to $\pi_{\mathrm{c}}(\mu)$ rather than to $\pi_{\mathrm{r}}(\mu)$ because it seems to be the preferred choice for mathematical reasons: the proof that B_π corresponds to Tab_μ is the easiest for $\pi = \pi_{\mathrm{c}}(\mu)$, and the relation between $B_{\pi_{\mathrm{c}}(\mu)}$ and a set of tableaux has analogues in other classical types (see [KaNa]), which is not the case with $B_{\pi_{\mathrm{r}}(\mu)}$. Historically however it is (the set of sequences of indices corresponding to) $B_{\pi_{\mathrm{r}}(\mu)}$ that has received more attention: for instance, the fact that the operators e_α preserve compatibility with the tableau conditions is already assumed implicitly in [Rob], and proved in [Macd, I (9.6)]. Note also that our choice is not of crucial importance: while the relation between paths and tableaux was instrumental in finding the construction presented below, that construction itself will be formulated in terms of paths, and applicable regardless of any connection of those paths with tableaux.

By the decomposition formula of [Litt2], the multiplicity of the irreducible $\mathfrak{gl}_n$-module V_ν in the tensor product $V_\lambda \otimes V_\mu$ equals the number of λ-dominant paths in $B_{\pi_{\mathrm{c}}(\mu)}$ of weight $\nu - \lambda$, i.e., paths $\pi \in B_{\pi_{\mathrm{c}}(\mu)}$ for which the translated path $\lambda + \pi$ goes from λ to ν and lies entirely within the dominant chamber. Each segment of $\lambda + \pi$ goes from one dominant integral weight to another, so we obtain a sequence of partitions from λ to ν that we shall call the λ-chain of π. If π was derived, as shown above, from $T \in \mathrm{Tab}_\mu$, then its λ-chain can be formed, starting from λ, by traversing T in the order indicated, and for every entry i encountered forming a new partition by adding 1 to part i of the previous partition. If we extend this procedure slightly by filling at each step the square (in row i) added to the Young diagram of the partition with a particular number r, then this will associate to T a Littlewood-Richardson tableaux T' of shape ν/λ and weight μ (as used in the classical formulation of the Littlewood-Richardson rule). It suffices to specify r: it is the row number in T of the entry i encountered at the current step.

In fact this procedure defines a correspondence between the squares s of T and the squares t of T', i.e., a bijection between the squares of the Young diagram of μ and those of the skew diagram ν/λ. This bijection is such that in T the square s contains the row number the corresponding square t, while in T' the square t contains the row number of s. It follws that T can be reconstructed from T' by a quite similar procedure. In fact one may consider both T and T' merely as ways to represent the bijection between squares. The bijections so occurring can be characterised by geometric properties that are contained in the notion of *pictures* ([Zel1]); using this notion it becomes obvious that for T' one finds exactly the set of Littlewood-Richardson tableaux of shape ν/λ and weight μ. Pictures provide a very versatile means to study these tableaux, due to the fact that many operations can be defined directly in terms of pictures, see [vLee2] (in that paper the conditions in the definition of pictures are transposed; in citing results we shall adapt for this difference). Below we shall freely use constructions defined for pictures and their properties; for the convenience of those not acquainted with pictures, we shall also give the translations of those constructions in terms of tableaux. It is worth noting that if one takes T to represent a path $\pi' \in B_{\pi_{\mathrm{r}}(\mu)}$ rather than $\pi \in B_{\pi_{\mathrm{c}}(\mu)}$, then one not only obtains the same set of admissible T (i.e., π is λ-dominant if and only if π' is), but for each such T the two orders of traversal lead to the same picture (bijection between squares of μ and of ν/λ), and therefore construct the same Littlewood-Richardson tableau T'. The path π' moreover has the property that its λ-chain gives the standardisation of the semistandard tableau T', which is not the case with π.

3.2. *Symmetry of the Littlewood-Richardson rule.*

We now turn to the question of exhibiting the symmetry of the tensor product combinatorially in type A_{n-1}, using Littlewood-Richardson tableaux.

We are looking for a bijection between Littlewood-Richardson tableaux of shape ν/λ and weight μ on one side, and Littlewood-Richardson tableaux of shape ν/μ and weight λ on the other side. In terms of pictures such a bijection is fairly easy to construct. There is a unique picture $\lambda \to -\lambda$, which may be "glued" to any given picture $f: \nu/\lambda \to \mu$, to form a picture $\bar{f}: \nu \to \mu \uplus -\lambda$ (the operation '$\uplus$' is "concatenation" of skew diagrams in the anti-diagonal direction, or more precisely of classes of skew diagrams modulo translation in the plane). Then one can apply the Schützenberger algorithm for pictures to obtain a picture $S(\bar{f}): \nu \to \lambda \uplus -\mu$; the image under the inverse picture $S(\bar{f})^{-1}$ of the factor $-\mu$ of the image is necessarily the subdiagram μ of ν, so by restriction of $S(\bar{f})$ to the complement of this subdiagram we obtain the desired picture $\nu/\mu \to \lambda$. The construction is easily seen to be involutive, and hence it defines a bijection between the sets of pictures $\mathrm{Pic}(\nu/\lambda, \mu)$ and $\mathrm{Pic}(\nu/\mu, \lambda)$.

In terms of Littlewood-Richardson tableaux such as T', this construction amounts to the following. The skew tableau of shape ν/λ is extended to a tableau of shape ν by filling each square of λ with minus its distance to the bottom of its column in λ (so the lowest square in each column gets -1, the square above it -2, etc.); the important property of this subtableau of shape λ is that it corresponds under the Schützenberger involution to the "canonical" tableau of shape and weight λ, in which each row i is filled with entries i. One then applies the Schützenberger involution to the full semistandard tableau of shape ν to obtain another such tableau in which the multiset of entries is negated, so that for each of the negative entries in the original tableau (within the Young diagram of λ) one now has a positive entry (at some other place of course), and vice versa; the subtableau of positive entries of the new tableau is a Littlewood-Richardson tableau of shape ν/μ and weight λ.

This algorithm can be simplified, if one recalls that one way to compute the Schützenberger involution applied to a Young tableau Y, is by "inflation". This is done by traversing the entries i of Y in increasing order (as usual processing equal entries from left to right), using them to repeatedly modify a tableau Z, initially empty, as follows: one performs an outward jeu de taquin slide of Z into the square occupied by i in Y, after which the vacated top-left corner of Z is filled with the value $-i$. Applying this algorithm to the semistandard Young tableau extended from T', one sees that in a first stage of computation the Schützenberger involution is applied to the subtableau of shape λ; as remarked, the result is the canonical tableau of shape λ. In a second stage outward slides are then applied to this tableau according to T'; the (negative) entries added at the top-left during the second stage can be ignored, since they will be removed from the final result anyway. So the bijection expressing the symmetry of the Littlewood-Richardson rule with respect to the partitions λ and μ describing the tensorands is given by jeu de taquin, rather than by the full Schützenberger algorithm; it is described in the following proposition, whose proof is contained in the reasoning above.

3.2.1. Proposition. *A bijection between Littlewood-Richardson tableaux L of shape ν/λ and weight μ and Littlewood-Richardson tableaux M of shape ν/μ and weight λ is given by the following algorithm: the tableau M is obtained from the canonical tableau of shape and weight λ by applying a series of successive outward jeu de taquin slides into the squares of ν/λ, as ordered by increasing entries in L, where squares with equal entries are ordered from left to right. Applying the same algorithm to M (interchanging the values of λ and μ) will reconstruct L.* $\square$

3.3. *Jeu de taquin for chains of partitions, and for paths.*

We shall now translate the construction above back in terms of paths, which will result in a remarkably simple operation that can be generalised to other types than A_{n-1}. In associating paths with tableaux such as L and M in the proposition above, sequences of partitions are natural intermediate objects: on one hand L is used there only to obtain an ordering of the squares within its shape ν/λ, as represented by its standardisation, which corresponds to a saturated increasing chain of partitions from λ to ν; on the other hand such a chain of partitions is the λ-chain of some ε-path.

As was noted above, the λ-chain of a λ-dominant path $\pi \in B_{\pi_c(\mu)}$ does not correspond to the standardisation of the corresponding Littlewood-Richardson tableau T'. We can nevertheless interpret the jeu de taquin process as operating directly on paths, in two ways. One is to pragmatically choose to work with λ-dominant paths in $B_{\pi_r(\mu)}$ rather than in $B_{\pi_c(\mu)}$; as remarked above, the λ-chain of such a path does correspond to the standardisation the associated Littlewood-Richardson tableau. More fundamentally, one may observe that T' really represents a picture $\nu/\lambda \to \mu$, which has many specialisations (standard tableaux of shape ν/λ that can be associated with it according to some "reading" of μ); one of these is the standardisation of T', while the λ-chain of $\pi \in B_{\pi_c(\mu)}$ corresponds to another. Moreover, different specialisations of the same picture have the property that when used to determine sequences of jeu de taquin slides, the final effect of any of these sequences of slides on the same initial tableau is identical (this is more generally true for tableaux that are dual equivalent). Therefore, if in proposition 3.2.1 we take for L the Littlewood-Richardson tableau constructed from $\pi \in B_{\pi_c(\mu)}$, then the same tableau M will be computed as in the proposition if we apply slides according to the λ-chain of π, rather than according to the specialisation of L.

We arrive at describing jeu de taquin in terms of chains of partitions. It is not necessary that the initial tableau to which we apply outward slides is a Young tableau; therefore we shall admit chains that start in an arbitrary partition κ. Given a saturated increasing chain of partitions from κ to λ (for instance, with $\kappa = (0)$, the 0-chain of $\pi_c(\lambda)$) corresponding to a (semi)standard tableau C, and a similar chain from λ to ν (the λ-chain of a λ-dominant ε-path) corresponding to a tableau L, the question is to describe the chain of partitions corresponding to the skew tableau M resulting from the application of successive outward jeu de taquin slides to C into the squares added in the chain of L. This has been done in [vLee1, §2] (for the Schützenberger algorithm, but it applies also to jeu de taquin), see also [vLee4, §2.1]. The family of partitions $\lambda^{[i,j]}$, defined by the fact that $\lambda^{[i,0]}, \ldots, \lambda^{[i,l]}$ is the chain corresponding to the tableau obtained after applying i slides to C, satisfies a local condition that allows $\lambda^{[i+1,j]}$ to be determined when $\lambda^{[i,j]}$, $\lambda^{[i,j+1]}$, and $\lambda^{[i+1,j+1]}$ are given:

3.3.1. Rule. *One has $\lambda^{[i,j+1]} = \lambda^{[i+1,j]}$ if and only if the two squares of $\lambda^{[i+1,j+1]} \setminus \lambda^{[i,j]}$ are adjacent.*

Note that in the case that the mentioned squares are non-adjacent, $\lambda^{[i+1,j]}$ is the unique partition strictly between $\lambda^{[i,j]}$ and $\lambda^{[i+1,j+1]}$ that differs from $\lambda^{[i,j+1]}$. This rule allows all partitions $\lambda^{[i,j]}$ to be computed when they are initially given only for pairs (i,j) with $i = 0$ (by means of C) or $j = l$ (by means of L). The same rule also allows $\lambda^{[i,j+1]}$ to be determined when $\lambda^{[i,j]}$, $\lambda^{[i+1,j]}$, and $\lambda^{[i+1,j+1]}$ are given. This reaffirms that the construction of proposition 3.2.1 is its own inverse. Note that in the proposition we take C to be the canonical tableau of shape λ; this is possible since we know that M must be a Littlewood-Richardson tableau, and therefore be reducible by jeu de taquin to this canonical tableau. In order to define a similar construction in terms of paths however, it will be necessary to explicitly supply data corresponding to C: if we want to construct from a λ-dominant path $p \in B_\pi$ a corresponding μ-dominant path p' (where $\mu = \pi(1)$), then we must specify the path $\pi' \in \mathcal{P}^+$ such that $\mu \in B_{\pi'}$; similiarly the inverse operation requires π to be specified. Therefore the path analogue of the bijection of proposition 3.2.1 will be a bijective correspondence $(p, \pi') \leftrightarrow (p', \pi)$.

We shall now reformulate the rule above in terms of ε-paths. With respect to chains of partitions there are some minor differences. First of all, partitions are limited to those with at most n parts. Second, we are interested primarily in the vertical and horizontal difference vectors $v_{i,j} = \lambda^{[i+1,j]} - \lambda^{[i,j]}$ and

$h_{i,j} = \lambda^{[i,j+1]} - \lambda^{[i,j]}$, which lie in the set $\{\varepsilon_1, \ldots, \varepsilon_n\}$, and represent segments of ε-paths; indeed, they are equal to ε_r where r is the row number of the square $\lambda^{[i+1,j]} \setminus \lambda^{[i,j]}$ respectively of the square $\lambda^{[i,j+1]} \setminus \lambda^{[i,j]}$. Now in case the two squares mentioned in rule 3.3.1 are non-adjacent, the square $\lambda^{[i+1,j]} \setminus \lambda^{[i,j]}$ is equal to $\lambda^{[i+1,j+1]} \setminus \lambda^{[i,j+1]}$, and similarly the square $\lambda^{[i,j+1]} \setminus \lambda^{[i,j]}$ is equal to $\lambda^{[i+1,j+1]} \setminus \lambda^{[i+1,j]}$; in this case we therefore certainly have $v_{i,j} = v_{i,j+1}$ and (equivalently) $h_{i+1,j} = h_{i,j}$. These two equalities remain valid in case the squares mentioned in rule 3.3.1 are horizontally adjacent, i.e., they both lie in the same row r, since in that case $v_{i,j} = v_{i,j+1} = h_{i+1,j} = h_{i,j} = \varepsilon_r$. Therefore, the only case where $v_{i,j} \neq v_{i,j+1}$, and where $h_{i+1,j} \neq h_{i,j}$, is when the two squares of $\lambda^{[i+1,j+1]} \setminus \lambda^{[i,j]}$ are vertically adjacent; in that case, if the rows containing these squares are r and $r+1$, one has $h_{i,j} = \varepsilon_r = v_{i,j}$ and $v_{i,j+1} = \varepsilon_{r+1} = h_{i+1,j}$. Given these values of $h_{i,j}$ and $v_{i,j+1}$ (or of $v_{i,j}$ and $h_{i+1,j}$), the condition that there is indeed vertical adjacency can be expressed as $\lambda^{[i,j]}_r = \lambda^{[i,j]}_{r+1}$, or equivalently as $\langle \lambda^{[i,j]}, \alpha_r^\vee \rangle = 0$. Note that with this condition satisfied it would not even be possible to have $v_{i,j} = v_{i,j+1}$, since that would make $\langle \lambda^{[i,j+1]}, \alpha_r^\vee \rangle = -1$, contradicting the fact that $\lambda^{[i,j+1]}$ is a partition, and corresponds to a dominant weight. We arrive at the following rule that describes how $v_{i,j}$ and $h_{i+1,j}$ are determined by the values of $h_{i,j}$, $v_{i,j+1}$ and $\lambda^{[i,j]}$.

3.3.2. Rule. *One has $v_{i,j} = v_{i,j+1}$ and $h_{i+1,j} = h_{i,j}$, unless for some r one has $h_{i,j} = \varepsilon_r$, $v_{i,j+1} = \varepsilon_{r+1}$, and $\langle \lambda^{[i,j]}, \alpha_r^\vee \rangle = 0$, in which case $v_{i,j} = \varepsilon_r$ and $h_{i+1,j} = \varepsilon_{r+1}$.*

We can now reformulate jeu de taquin in terms of ε-paths. A strict translation of proposition 3.2.1 into this language would give a statement that only applies to ε-paths that correspond to semistandard Young tableaux, but as explained above we remove that restriction by supplying an extra parameter π'. To recover that proposition one should take $\kappa = (0)$, and π' equal to the ε-path corresponding to the canonical tableau of shape λ, i.e., to $\pi_{\mathrm{c}}(\lambda)$ or $\pi_{\mathrm{r}}(\lambda)$, depending on the chosen correspondence between paths and tableaux.

3.3.3. Construction (jeu de taquin for ε-paths). *Let κ, λ, ν be dominant integral weights for $\mathfrak{gl}_n$, π' a κ-dominant ε-path of length l with $\pi'(1) = \lambda - \kappa$, and p a λ-dominant ε-path of length k with $p(1) = \nu - \lambda$. We construct a dominant integral weight μ, a κ-dominant ε-path π of length k with $\pi(1) = \mu - \kappa$, and a μ-dominant ε-path p' of length l with $p'(1) = \nu - \mu$ in the following steps.*

- *Set $h_{0,0}, \ldots, h_{0,l-1}$ according to the sequence of segments of π', and $v_{0,l}, \ldots, v_{k-1,l}$ according to the sequence of segments of p;*
- *Set $\lambda^{[0,j]} := \kappa + \sum_{j'<j} h_{0,j'}$ for $0 \leq j \leq l$, and $\lambda^{[i,l]} := \lambda + \sum_{i'<i} v_{i',l}$ for $0 \leq i \leq k$;*
- *Determine the values $h_{i,j}$ for $0 < i \leq k$ and $l > j \geq 0$, as well as $v_{i,j}$ for $0 \leq i < k$ and $l > j \geq 0$ using rule 3.3.2, setting $\lambda^{[i,j+1]} := \lambda^{[i,j]} + v_{i,j} = \lambda^{[i+1,j+1]} - h_{i+1,j}$ after each application of the rule;*
- *Return $\pi = \pi_{v_{0,0}} * \cdots * \pi_{v_{k-1,0}}$ and $p' = \pi_{h_{k,0}} * \cdots * \pi_{h_{k,l-1}}$.*

Note that the relations between the parameters of the construction allow all of them to be deduced if κ, π', and p are given; we shall therefore consider the construction to be parametrised by (κ, π', p), and to return the pair (π, p'). The following theorem is obvious, both from the symmetry of the construction, and from the fact that that construction is just a reformulation of jeu de taquin.

3.3.4. Theorem (symmetry of jeu de taquin for ε-paths). *The construction 3.3.3 is its own inverse: if when applied to (κ, π', p) it returns (π, p'), then applied to (κ, π, p') it will return (π', p). Moreover, it is symmetric with respect to dualisation of paths: when applied to (ν, p^*, π'^*) it will return (p'^*, π^*).* $\square$

Despite its somewhat technical formulation, the following lemma is just an expression of the trivial fact that jeu de taquin consists of consecutive application of slides: performing inward slides according to a path π'_2 followed by performing inward slides according to π'_1 amounts to performing inward slides according to the concatenation $\pi'_1 * \pi'_2$.

3.3.5. Lemma. *Let construction 3.3.3 be applicable to (κ, π', p). If π' is of the form $\pi'_1 * \pi'_2$, then the construction is applicable to $(\kappa + \pi'_1(1), \pi'_2, p)$, and calling the result of this application (q, p'_2), it is also applicable to (κ, π'_1, q); calling the result of this second application (π, p'_1), the result of applying the construction to (κ, π', p) will be $(\pi, p'_1 * p'_2)$. A similar composition formula holds if p is of the form $p_1 * p_2$.* $\square$

The following theorem establishes the fundamental link between jeu de taquin and Littelmans's root operators e_α and f_α. It is not an entirely new result: the fact that the definition of e_α and f_α corresponds to jeu de taquin on tableaux of two rows is well known to experts; a discussion of this relation including a proof of a statement equivalent to the theorem can be found in in [vLee4, §3.1]. We shall give another proof here that is formulated in terms of paths, so that generalisation to other types will be straightforward.

3.3.6. Theorem. *In the situation of construction 3.3.3 the path π can be obtained from p by application of a sequence of operators e_i, and similarly the path p' can be obtained from π' by application of a sequence of operators f_i. In particular, if $\kappa = 0$, one has $p \in B_\pi$ and $p' \in B_{\pi'}$.*

Proof. By symmetry (theorem 3.3.4) it suffices to prove the first statement (about p and π). By lemma 3.3.5, it will suffice to prove the case where π' has length 1, which we therefore assume henceforth. In order to proceed by induction on the length of p, it is necessary to strengthen the statement being proved as follows: there exists a sequence of indices $i_1, \ldots, i_n$ (with $n \geq 0$), and a sequence of paths $p = p_n, p_{n-1}, \ldots, p_0 = \pi$, such that for $j = n, \ldots, 1$, one has $p_{j-1} = e_{i_j}(p_j)$ and $\lfloor \kappa + p_j \rfloor_{\alpha_{i_j}} = -1$ (which shows that p_j is *not* κ-dominant for $j > 0$). If p is of length 0 we take $n = 0$ and there is nothing to prove; assume therefore that p has positive length. Let $v = \lambda^{[1,1]} - \lambda^{[0,1]}$ in construction 3.3.3, so that we can write $p = \pi_v * q$; similarly put $v' = \lambda^{[1,0]} - \lambda^{[0,0]}$ and $\pi = \pi_{v'} * \rho$. With $\kappa' = \lambda^{[1,0]}$ we consider the construction applied to $(\kappa', \pi_{h_{1,0}}, q)$, which by lemma 3.3.5 returns (ρ, p'). By the induction hypothesis there exist indices $i_1, \ldots, i_m$ and paths $q = q_m, \ldots, q_0 = \rho$ with $q_{j-1} = e_{i_j}(q_j)$ and $\lfloor \kappa' + q_j \rfloor_{\alpha_{i_j}} = -1$ for $0 < j \leq m$. For $j \leq m$ we put $p_j = \pi_{v'} * q_j$ (so that in particular $p_0 = \pi$). Then for $0 < j \leq m$ one has $\lfloor \kappa + p_j \rfloor_{\alpha_{i_j}} = \lfloor \kappa' + q_j \rfloor_{\alpha_{i_j}} = -1$, since the path $\pi_{v'}$ is κ-dominant with $\pi_{v'}(1) = v' = \kappa' - \kappa$. We see moreover that the minimum taken in the first expression is attained only in the second part of the concatentation $p_j = \pi_{v'} * q_j$; from the definition of e_i we therefore have $e_i(p_j) = \pi_{v'} * e_j(q_j) = \pi_{v'} * q_{j-1} = p_{j-1}$. Now if $v' = v$, we have $p = \pi_v * q = \pi_{v'} * q_m = p_m$ so that we take $n = m$ and we are done. Otherwise we are in the exceptional case of rule 3.3.2 for $(i, j) = (0, 0)$, so that $\pi' = \pi_{v'} = \pi_{\varepsilon_r}$, $\pi_v = \pi_{\varepsilon_{r+1}}$, and $\langle \kappa, \alpha_r^\vee \rangle = \langle \lambda^{[1,1]}, \alpha_r^\vee \rangle = 0$ for some r. Since the path q is $\lambda^{[1,1]}$-dominant, this implies $\langle q(t), \alpha_r^\vee \rangle \geq 0$ for all t, so that $\lfloor \kappa + p \rfloor_{\alpha_r} = -1$, which minimum is first attained at the point of concatenation of $p = \pi_v * q$; consequently, we have by the definition of e_r that $e_r(p) = \pi_{\varepsilon_r} * q = \pi_{v'} * q_m = p_m$. In this case we therefore put $n = m + 1$ and $i_n = r$, and we have established all that needs to be proved. □

§4. Jeu de taquin for other types of groups.

Now the stage has been set in type A_{n-1}, we may consider possible generalisations to other types of groups. While the traditional planar form of jeu de taquin does not seem to be easily generalised, the situation is quite different for the formulation in terms of paths, since there are only a few points in the discussion above that are specific for type A_{n-1}, and need replacement for other types. Firstly, one needs a replacement for the set $\{ \pi_{\varepsilon_i} \mid i = 1, \ldots, n \}$ of elementary path segments, and hence for the class of ε-paths; secondly, rule 3.3.2 will need to be adapted to this new class of paths. Once this is done, a counterpart of the construction 3.3.3 can be defined with only the most obvious adaptations. Provided the replacement for rule 3.3.2 preserves its symmetry, the analogue of theorem 3.3.4 will be valid, with an equally simple proof; the analogue of lemma 3.3.5 remains a triviality. Having succeeded so far, we shall have a involutive construction that operates on pairs (p, π') of paths; then in order that we can use this construction to define, for any paths $\pi, \pi' \in \mathcal{P}^+$ (in the chosen class) with $\pi(1) = \mu$ and $\pi'(1) = \lambda$, a bijective correspondence between λ-dominant paths in B_π and μ-dominant paths in $B_{\pi'}$, it is essential that the analogue of theorem 3.3.6 holds. Most of its proof will remain valid without modification, but the final argument involving a single configuration governed by (an analogue of) rule 3.3.2 needs to be verified. Each time we establish these points, we obtain a combinatorial analogue of jeu de taquin that shares two of its most fundamental properties: symmetry (theorem 3.3.4) and confluence, i.e., the fact that, for $\kappa = 0$, the correspondence $p \to \pi$ is independent of the choice of $\pi' \in \mathcal{P}^+$ (this will be a consequence of the analogue of theorem 3.3.6 and the fact that from any path p at most one path $\pi \in \mathcal{P}^+$ can be obtained by applications of root operators e_α). The combinatorial constructions that we shall find have a common generalisation to arbitrary piecewise linear paths, and maybe even to continuous paths; however, in this generality the combinatorial nature of the construction will be lost.

4.1. *Minuscule representations and m-paths.*

Let W be the Weyl group of $\mathfrak{g}$. A non-trivial irreducible representation of $\mathfrak{g}$ is called *minuscule* if its set of weights forms a single W-orbit; these weights are called minuscule weights. The following representations are minuscule: all fundamental representations in type A_n, the natural (defining) representation in types C_n and D_n, the spin representation in type B_n, the half-spin representations in type D_n, the two 27-dimensional representations in type E_6, and the 56-dimensional representation in type E_7; there are no minuscule representations in types G_2, F_4 and E_8, since, as the weight lattice coincides with the root lattice in these types, all representations contain the weight 0. For any minuscule weight m and any root α one has $\langle m, \alpha^\vee\rangle \in \{-1, 0, 1\}$, since otherwise $m + \mathbf{Z}\alpha$ would intersect the weight system in more than two points. Therefore one has $B_{\pi_\lambda} = \{\, \pi_m \mid m \in W\lambda \,\}$ for any dominant minuscule weight λ. This makes the class of paths obtained by concatenation of segments π_m for m minuscule a good candidate to replace the class of ε-paths. We shall call any concatenation of l segments of the form π_m, with m minuscule, an m-path (pun not intended) of length l.

In order to formulate an analogue of rule 3.3.2 for m-paths, we are led to consider the following situation. Let κ, λ and μ be dominant integral weights such that $\pi_{\lambda-\kappa}$ and $\pi_{\nu-\lambda}$ are m-paths of length 1; the question is to find a dominant weight μ such that $\pi_{\mu-\kappa}$ can be obtained from $\pi_{\nu-\lambda}$ by a series of applications of operators e_α with $\langle\kappa, \alpha^\vee\rangle = 0$, and such that the argument π_{v_i} to which the operator is applied satisfies $\langle v_i, \alpha^\vee\rangle = -1$ (there is a similar condition for the transformation $\pi_{\lambda-\kappa} \to \pi_{\nu-\mu}$, which involves operators f_α). The main difference with the situation for ε-paths is that, whereas in that case at most one application of e_r suffices (which transforms $\pi_{\nu-\lambda} = \pi_{\varepsilon_{r+1}}$ into $\pi_{\mu-\kappa} = \pi_{\varepsilon_r}$), a series of applications may be needed for m-paths. This phenomeneon already occurs in type A_{n-1}: if one takes $\kappa = \nu = 0$, and $\lambda = \omega_i$, so that $\lambda - \kappa = \omega_i = \varepsilon_{\{1,\ldots,i\}}$, and $\nu - \lambda = -\omega_i = \varepsilon_{\{i+1,\ldots,n\}}$, then the only possibility is to have $\mu = \omega_{n-i}$, so that $\mu - \kappa = \omega_{n-i} = \varepsilon_{\{1,\ldots,n-i\}}$ and $\nu - \mu = -\omega_{n-i} = \varepsilon_{\{n-i+1,\ldots,n\}}$; the transformation of $\pi_{\varepsilon_{\{i+1,\ldots,n\}}}$ into $\pi_{\varepsilon_{\{1,\ldots,n-i\}}}$ requires a total of $i(n-i)$ applications of operators ε_α (some operators may be applied more than once, but never twice in succession).

We see in this example that the seqence of operators applied may not be unique (unlike in the proof of theorem 3.3.6), but the final result is. In fact it is not difficult to see that this is true in general. Let $S = \{\, i \mid \langle\kappa, \alpha_i^\vee\rangle = \langle\nu, \alpha_i^\vee\rangle = 0 \,\}$, and put $v_0 = \nu - \lambda$, $h_0 = \lambda - \kappa$; these are the initial candidates for $\mu - \kappa$ and $\nu - \mu$ (the sum of these candidates will always be $\nu - \kappa$). In order that μ be dominant, it is necessary that $\langle\mu - \kappa, \alpha_i^\vee\rangle = \langle\mu - \nu, \alpha_i^\vee\rangle \geq 0$ for all $i \in S$. Therefore, while there exists for the current candidates v_j and h_j for $\mu - \kappa$ and $\nu - \mu$ an $i \in S$ with $\langle v_j, \alpha_i^\vee\rangle = \langle -h_j, \alpha_i^\vee\rangle = -1$ we choose such an i and replace the candidates by $v_{j+1} = s_{\alpha_i}(v_j) = v_j + \alpha_i$ and $h_{j+1} = s_{\alpha_i}(h_j) = h_j - \alpha_i$. After a finite number of steps this process terminates, and we set $\mu = \kappa + v_l = \nu - h_l$ for the final values v_l, h_l. Writing $W_{\kappa,\nu}$ for the subgroup of W stabilising κ and ν (it is generated by $\{\, s_{\alpha_i} \mid i \in S \,\}$), and $\mathrm{dom}_{W_{\kappa,\nu}}$ for the map that sends any weight to the $W_{\kappa,\nu}$-dominant representative of its $W_{\kappa,\nu}$-orbit, we clearly have $\mu - \kappa = \mathrm{dom}_{W_{\kappa,\nu}}(\nu - \lambda)$ and $\mu - \nu = \mathrm{dom}_{W_{\kappa,\nu}}(\kappa - \lambda)$, which shows that these values are independent of the choices made of the indices i. We have achieved $\langle\mu, \alpha_i^\vee\rangle \geq 0$ for all $i \in S$; to prove that μ is dominant it suffices to establish the same for $i \notin S$. For such i we have $\langle\kappa, \alpha_i^\vee\rangle \geq 1$ or $\langle\nu, \alpha_i^\vee\rangle \geq 1$ (possible both); since $\langle\mu - \kappa, \alpha_i^\vee\rangle$ and $\langle\nu - \mu, \alpha_i^\vee\rangle$ lie in $\{-1, 0, 1\}$, either of these inequalities implies $\langle\mu, \alpha_i^\vee\rangle \geq 0$. From this description we see that in fact $\mu = \mathrm{dom}_W(\kappa + \nu - \lambda)$. We can therefore formulate a generalisation of rule 3.3.2 simply as follows.

4.1.1. Rule. $\lambda^{[i+1,j]} = \mathrm{dom}_W(\lambda^{[i,j]} + \lambda^{[i+1,j+1]} - \lambda^{[i,j+1]})$.

4.1.2. Lemma. *For fixed values of $\lambda^{[i,j]}$ and $\lambda^{[i+1,j+1]}$, the correspondence between $\lambda^{[i,j+1]}$ and $\lambda^{[i+1,j]}$ determined by rule 4.1.1 is symmetrical. Moreover the rule is symmetrical in $\lambda^{[i,j]}$ and $\lambda^{[i+1,j+1]}$.*

Proof. From the considerations above it follows that $\mu = w(\kappa + \nu - \lambda)$ for some $w \in W$ that fixes κ and ν; this implies $\lambda = w^{-1}(\kappa + \nu - \mu)$, and since λ is dominant, this establishes the first symmetry. The second symmetry is obvious. □

As the path segments $h_{i,j}$ and $v_{i,j}$ are absent from the formulation of the rule 4.1.1, we can formulate a construction that is in the spirit of the original formulation of jeu de taquin in term of chains of partitions, in that only a doubly indexed family of partitions is considered. It should be noted however that paths were used to find rule 4.1.1, and they will play a rôle in proofs concerning the construction as well.

4.1.3. Construction (jeu de taquin for m-paths). *Let κ, λ, ν be dominant integral weights for $\mathfrak{g}$, π' a κ-dominant m-path of length l with $\pi'(1) = \lambda - \kappa$, and p a λ-dominant m-path of length k with $p(1) = \nu - \lambda$; we assume for each of π' and p that their segments are traversed at equal speeds. We construct a dominant integral weight μ, a κ-dominant m-path π of length k with $\pi(1) = \mu - \kappa$, and a μ-dominant m-path p' of length l with $p'(1) = \nu - \mu$ in the following steps.*

- *Set $\lambda^{[0,j]} := \kappa + \pi'(j/l)$ for $0 \leq j \leq l$, and $\lambda^{[i,l]} := \lambda + p(i/k)$ for $0 \leq i \leq k$;*
- *Determine the values $\lambda^{[i,j]}$ for $0 < i \leq k$ and $l > j \geq 0$ using rule 4.1.1;*
- *Return $\pi = \pi_{v_1} * \cdots * \pi_{v_k}$ and $p' = \pi_{h_1} * \cdots * \pi_{h_l}$, where $v_i = \lambda^{[i,0]} - \lambda^{[i-1,0]}$ and $h_j = \lambda^{[k,j]} - \lambda^{[k,j-1]}$.*

4.1.4. Theorem (symmetry of jeu de taquin for m-paths). *The construction 4.1.3 is its own inverse: if when applied to (κ,π',p) it returns (π,p'), then applied to (κ,π,p') it will return (π',p). Moreover, it is symmetric with respect to dualisation of paths: when applied to (ν,p^*,π'^*) it will return (p'^*,π^*).*

Proof. This is immediate from lemma 4.1.2. □

4.1.5. Lemma. *Lemma 3.3.5 remains valid when construction 3.3.3 is replaced by construction 4.1.3.* □

4.1.6. Theorem. *In the situation of construction 4.1.3 the path π can be obtained from p by application of a sequence of operators e_i, and similarly the path p' can be obtained from π' by application of a sequence of operators f_i. In particular, if $\kappa=0$, one has $p\in B_\pi$ and $p'\in B_{\pi'}$.*

Proof. The proof of theorem 3.3.6 can be followed literally, with the obvious replacement of references by their counterparts for m-paths, up to and including the proof that $e_i(p_j)=p_{j-1}$ for $0<j\leq m$; after that we continue as follows. Let $v=v_0,\ldots,v_l=v'$ be the sequence of vectors in the discussion preceding the statement of rule 4.1.1; put $n=m+l$ and $p_{n-i}=\pi_{v_i}*q$ for $i=0,\ldots,l$ (this agrees with the previous definition of p_m, and we have $p_n=p$). It was established there that for all $i<l$ there exists a simple root α such that $\langle\kappa,\alpha^\vee\rangle=\langle\lambda^{[1,1]},\alpha^\vee\rangle=0$, $\langle v_i,\alpha^\vee\rangle=-1$, and $v_{i+1}=s_\alpha(v_i)$, so that $\pi_{v_{i+1}}=e_\alpha(\pi_{v_i})$. Since the path q is $\lambda^{[1,1]}$-dominant we have $\langle q(t),\alpha^\vee\rangle\geq 0$ for all t, so that $\lfloor\kappa+p_{n-i}\rfloor_\alpha=-1$, and $e_\alpha(p_{n-i})=e_\alpha(\pi_{v_i}*q)=\pi_{v_{i+1}}*q=p_{n-i-1}$; this establishes all that needs to be proved. □

4.2. *Quasi-minuscule weights and ψ-paths.*

We can extend the class of paths for which our construction works beyond that of m-paths by allowing path segments that correspond to the weights of representations slightly larger than the minuscule ones, which will in particular allow us to treat paths for groups of the types G_2, F_4, and E_8 that do not possess minuscule representations. As we shall see, the extra freedom will lead to a considerable increase in the number of situations that need to be treated.

An irreducible representation is called quasi-minuscule if its set of weights consists of two W-orbits, one of which is $\{0\}$; the weights in the other orbit are called quasi-minuscule weights. The orbit of quasi-minuscule weights is contained in the root lattice, and its dominant representative is a minimal non-zero element of the intersection of the dominant chamber with the root lattice; in particular quasi-minuscule weights are roots. One checks easily that for every simple type there is a unique quasi-minuscule representation: this is the adjoint representation for the simply laced types A_n, D_n and E_n, and the representation whose non-zero weights are the short roots for the other types B_n, C_n, F_4, and G_2 (for type B_n this is the natural representation). For similar reasons as mentioned for minuscule weights, one has for any root α and any quasi-minuscule weight $m\notin\{-\alpha,\alpha\}$ that $\langle m,\alpha^\vee\rangle\in\{-1,0,1\}$. It follows that if λ is a dominant quasi-minuscule weight, then the only paths in B_{π_λ} that are not of the form π_m for $m\in W\lambda$ are those of the form $e_\alpha(\pi_{-\alpha})=f_\alpha(\pi_\alpha)=\pi_{-\alpha/2}*\pi_{\alpha/2}$ for the simple roots α occurring in $W\lambda$; we shall denote such a path by ψ_α. We define a ψ-path of length l to be a concatenation of l segments that occur in the union of the sets B_{π_λ} for the dominant weights λ that are either minuscule or quasi-minuscule.

Now we consider the question of extending rule 4.1.1 to deal with any pair of ψ-paths of length 1. As before we consider dominant integral weights κ, λ, ν, and we put $S=\{\,i\mid\langle\kappa,\alpha_i^\vee\rangle=\langle\nu,\alpha_i^\vee\rangle=0\,\}$. Since the paths involved are no longer necessarily linear, it does not suffice to consider just the differences $\lambda-\kappa$ and $\nu-\lambda$; therefore let p and q be ψ-paths of length 1 such that $p(1)=\nu-\lambda$ and $q(1)=\lambda-\kappa$. Our goal is to find paths p' and q', obtained from p and q respectively by applications of operators e_α and f_α, such that $\mu=\kappa+p'(1)=\nu-q'(1)$ is dominant, and moreover p' is κ-dominant and q' is μ-dominant (this is an extra condition only when p' or q' is of the form ψ_α).

We consider first the case that both p and q are linear (i.e., not of the form ψ_α), which includes the case of m-paths treated above. Like in that case we perform an iteration, but now on a pair of paths, which we initialise by $p_0=p,q_0=q$. The iteration is the following one:

($*$) As long as there exists for the current pair p_j,q_j some $i\in S$ with $\langle p_j(1),\alpha_i^\vee\rangle=\langle -q_j(1),\alpha_i^\vee\rangle=-1$, we choose such an i and put $p_{j+1}=e_{\alpha_i}(p_j)$ and $q_{j+1}=f_{\alpha_i}(q_j)$.

All the paths so obtained are linear, since we could only have $p_{j+1}=\psi_\alpha$ if $\alpha_i=\alpha$ and $p_j=\pi_{-\alpha}$, which would violate $\langle p_j(1),\alpha_i^\vee\rangle=-1$. It follows also that this iteration cannot make a transition from negative to positive roots: if any p_j is of the form π_β for a negative root β, then the same is true for all p_j's; a similar statement holds when $p_j=\pi_\beta$ for a positive root β, and also for the q_j's in place of the p_j's.

If for the final paths p_k, q_k obtained after the iteration, the weight $\tilde\mu = \lambda + p_k(1) = \nu - q_k(1)$ is dominant, then we put $p' = p_k$ and $q' = q_k$. Otherwise, let i be such that $\langle \tilde\mu, \alpha_i^\vee \rangle < 0$. Suppose first that $i \notin S$; then since either $\langle \kappa, \alpha_i^\vee \rangle > 0$ or $\langle \lambda, \alpha_i^\vee \rangle > 0$, we must have $p_k = \pi_{-\alpha_i}$ or $q_k = \pi_{\alpha_i}$ (possibly both). We put $p' = p_{k+1} = e_{\alpha_i}(p_k)$ and $q' = q_{k+1} = f_{\alpha_i}(q_k)$; since $\mu = \lambda + p'(1) = \nu - q'(1)$ is equal to κ or ν, it is dominant, and one has $\langle \mu, \alpha_i^\vee \rangle = 1$, from which it follows that p' is κ-dominant and q' is μ-dominant. Now suppose $i \in S$; since the iteration $(*)$ has terminated, it must be that $\langle p_k(1), \alpha_i^\vee \rangle = -2$, so $p_k = \pi_{-\alpha_i}$ and $q_k = \pi_{\alpha_i}$. We then put $p_{k+1} = q_{k+1} = \psi_{\alpha_i}$ and $p_{k+2} = \pi_\alpha$, $q_{k+2} = \pi_{-\alpha}$ (so that $p_{j+1} = e_{\alpha_i}(p_j)$ and $q_{j+1} = f_{\alpha_i}(q_j)$ for $j = k, k+1$), after which we resume the iteration $(*)$, and set p' and q' respectively to the final paths p_l, q_l so obtained. From the fact that the iteration cannot make a transition between negative and positive roots it follows that this time $\mu = \lambda + p'(1) = \nu - q'(1)$ must be dominant.

Next we consider the case that $p = \psi_\alpha$ and $q = \psi_\beta$ for simple roots α, β. If $\alpha \neq \beta$ or if $\alpha = \beta$ and $\langle \lambda, \alpha^\vee \rangle > 1$ then we put $p' = p$ and $q' = q$, so that $\mu = \kappa = \lambda = \nu$. If $\alpha = \beta$ and $\langle \kappa, \alpha^\vee \rangle = 1$ we put $p_1 = e_\alpha(p) = \pi_\alpha$ and $q_1 = f_\alpha(q) = \pi_{-\alpha}$ and then perform iteration $(*)$. We set p' and q' respectively to the final paths p_l, q_l obtained; as in the case above we see that $\mu = \lambda + p'(1) = \nu - q'(1)$ is dominant.

We are left with the possibility that exactly one of p and q is linear. We shall only treat the case that this is q, as the other case is symmetric (by interchange of κ and ν and dualisation of the paths); let $p = \psi_\alpha$. If $\langle \kappa, \alpha^\vee \rangle > 0$, then we put $p' = p$ and $q' = q$, so that $\mu = \kappa$ is dominant and p' is κ-dominant; we assume henceforth that $\langle \kappa, \alpha^\vee \rangle = 0$. If $q = \pi_\alpha$ then we put $p' = e_\alpha(p) = \pi_\alpha$ and $q' = f_\alpha(q) = \psi_\alpha$. Otherwise we put $p_1 = e_\alpha(p) = \pi_\alpha$ and $q_1 = f_\alpha(q)$, after which perform iteration $(*)$, calling the final paths obtained p_k, q_k. As in the case of linear p and q it is possible that for $\tilde\mu = \lambda + p_k(1) = \nu - q_k(1)$ there is some i for which $\langle \tilde\mu, \alpha_i^\vee \rangle < 0$, but this time only for $i \notin S$ and $q_k = \pi_{\alpha_i}$, since $p_k \neq \pi_{-\alpha_i}$. If this is the case we put $p' = p_{k+1} = e_{\alpha_i}(p_k)$ and $q' = q_{k+1} = f_{\alpha_i}(q_k) = \psi_{\alpha_i}$, and otherwise ($\tilde\mu$ is dominant) we put $p' = p_k$ and $q' = q_k$; both cases are just like the corresponding ones for linear p and q.

This concludes the description of the determination of p' and q'. We shall now try to formulate the result as concisely as possible. To this end we shall use the fact that μ determines either of p' and q' if the path in question is linear, i.e., if μ differs from κ respectively from ν; if not, then the path is of the form ψ_α, and it suffices to specify in addition to μ the simple root α. We also simplify the formulation by using the fact that if λ equals either κ or ν, then the other one can be expressed as $\kappa + \nu - \lambda$. Since like before the rule stated will be used in a larger construction, we give the weights and path segments in the construction as elements of doubly indexed families. We leave it to the reader to verify that the results computed above satisfy the description below.

4.2.1. Rule. *Let $\kappa = \lambda^{[i,j]}$, $\lambda = \lambda^{[i,j+1]}$, $\nu = \lambda^{[i+1,j+1]}$, $p = v_{i,j+1}$, and $q = h_{i,j}$; the weight $\mu = \lambda^{[i+1,j]}$ and the paths $p' = v_{i,j}$ and $q' = h_{i+1,j}$ are determined according to the following cases.*

(a) If p and q are linear, then $\mu = \mathrm{dom}_W(\kappa + \nu - \lambda)$; in case μ equals κ or ν, the corresponding path is equal to ψ_α, where $\alpha = \mu - \mathrm{dom}_{W_{\kappa,\nu}}(\kappa + \nu - \lambda)$.

(b) If $\langle \lambda, \alpha^\vee \rangle = 1$ and $\psi_\alpha \in \{p, q\}$ for some simple root α, and either $p = q$ or $\langle \kappa + \nu - \lambda, \alpha^\vee \rangle = 0$, then $\mu = \mathrm{dom}_W(\kappa + \nu - \lambda + \alpha)$; in case μ equals κ or ν, the corresponding path is equal to $\psi_{\alpha'}$, where $\alpha' = \mu - \mathrm{dom}_{W_{\kappa,\nu}}(\kappa + \nu - \lambda + \alpha)$.

(c) If $\psi_\alpha \in \{p, q\}$ for some simple root α, with $\langle \lambda, \alpha^\vee \rangle = 2$ and $\langle \kappa + \nu - \lambda, \alpha^\vee \rangle = 0$, then $p' = q$, $q' = p$, and $\mu = \lambda$.

(d) If either $\{p, q\} = \{\psi_\alpha, \psi_\beta\}$ for simple roots $\alpha \neq \beta$, or $\psi_\alpha \in \{p, q\}$ for some simple root α and $\langle \kappa + \nu - \lambda, \alpha^\vee \rangle > 0$, then $p' = p$, $q' = q$, and $\mu = \kappa + \nu - \lambda$.

4.2.2. Lemma. *For fixed values of $\lambda^{[i,j]}$ and $\lambda^{[i+1,j+1]}$, the correspondence determined by rule 4.2.1 between $\lambda^{[i,j+1]}$ and $\lambda^{[i+1,j]}$, and between $(h_{i,j}, v_{i,j+1})$ and $(v_{i,j}, h_{i+1,j})$, is symmetrical.*

Proof. This follows from a careful analysis of the different cases that can arise. In cases (c) and (d) of rule 4.2.1, replacement of λ by the indicated value of μ leads to the same case, and gives back the original value of λ for μ. We may therefore assume that one of cases (a) and (b) applies. Define weights $\mu_0, \mu_1, \mu_2, \mu_3$ by $\mu_0 = \kappa + \nu - \lambda$, $\mu_1 = \mu_0$ in case (a) and $\mu_1 = \mu_0 + \alpha$ in case (b), $\mu_2 = \mathrm{dom}_{W_{\kappa,\nu}}(\mu_1)$, and $\mu_3 = \mu$; replacing λ by μ, call the corresponding weights $\lambda_0, \lambda_1, \lambda_2, \lambda_3$. One then proves successively that $\lambda_i = \kappa + \nu - \mu_{3-i}$ for $i = 0, 1, 2, 3$ in a straightforward manner in all the cases, using the details that were given before the statement of rule 4.2.1. □

4.2.3. Construction (jeu de taquin for ψ-paths). *Let κ, λ, ν be dominant integral weights for $\mathfrak{g}$, π' a κ-dominant ψ-path of length l with $\pi'(1) = \lambda - \kappa$, and p a λ-dominant ψ-path of length k with $p(1) = \nu - \lambda$. We construct a dominant integral weight μ, a κ-dominant ψ-path π of length k with*

$\pi(1) = \mu - \kappa$, and a μ-dominant ψ-path p' of length l with $p'(1) = \nu - \mu$, in the following steps.

- *Let $\pi' = h_{0,0} * \cdots * h_{0,l-1}$ and $p = v_{0,l} * \cdots * v_{k-1,l}$, where the $h_{i,j}$ and $v_{i,j}$ are ψ-paths of length 1;*
- *Set $\lambda^{[0,j]} := \kappa + \sum_{j'<j} h_{0,j'}(1)$ for $0 \leq j \leq l$, and $\lambda^{[i,l]} := \lambda + \sum_{i'<i} v_{i',l}(1)$ for $0 \leq i \leq k$;*
- *Determine the weights $\lambda^{[i+1,j]}$ and the paths $h_{i+1,j}$ and $v_{i,j}$ for $0 \leq i < k$ and $l > j \geq 0$, using rule 4.2.1;*
- *Return $\pi = v_{0,0} * \cdots * v_{k-1,0}$ and $p' = h_{k,0} * \cdots * h_{k,l-1}$.*

4.2.4. Theorem (symmetry of jeu de taquin for ψ-paths). *The construction 4.2.3 is its own inverse: if when applied to (κ, π', p) it returns (π, p'), then applied to (κ, π, p') it will return (π', p). Moreover, it is symmetric with respect to dualisation of paths: when applied to (ν, p^*, π'^*) it will return (p'^*, π^*).*

Proof. This is immediate from lemma 4.2.2. □

4.2.5. Lemma. *Lemma 3.3.5 remains valid when construction 3.3.3 is replaced by construction 4.2.3.* □

4.2.6. Theorem. *In the situation of construction 4.2.3 the path π can be obtained from p by application of a sequence of operators e_i, and similarly the path p' can be obtained from π' by application of a sequence of operators f_i. In particular, if $\kappa = 0$, one has $p \in B_\pi$ and $p' \in B_{\pi'}$.*

Proof. By symmetry (theorem 4.2.4) it suffices to prove the first statement (about p and π). By lemma 4.2.5, it will suffice to prove the case where π' has length 1, which we therefore assume henceforth. In order to proceed by induction on the length of p, it is necessary to strengthen the statement being proved as follows. There exists a sequence of indices $i_1, \ldots, i_n$ (with $n \geq 0$) and a sequence of paths $p = p_n, p_{n-1}, \ldots, p_0 = \pi$, such that for $j = n, \ldots, 1$, one has $e_{i_j}(p_j) = p_{j-1}$, and moreover $\lfloor \kappa + p_j \rfloor_{\alpha_{i_j}} < 0$ except when $\pi' = \psi_\alpha$ for some simple root α and $j = n$, in which case one has $\alpha_{i_n} = \alpha$ and $\lfloor \kappa + p_j \rfloor_\alpha = 0$. To this we add one more detail: if $\pi' = \psi_\alpha$ and $\lfloor \kappa + p \rfloor_\alpha = 0$, then $n > 0$.

If p is of length 0 we take $n = 0$ and there is nothing to prove; assume therefore that p has positive length. Let $v = v_{0,1}$ in construction 3.3.3, so that we can write $p = v * q$; similarly put $v' = v_{0,0}$ and $\pi = v' * \rho$. Put $\kappa' = \lambda^{[1,0]}$ and $\lambda' = \lambda^{[1,1]}$; we consider the construction applied to $(\kappa', h_{1,0}, q)$, which by lemma 4.2.5 returns (ρ, p'). By the induction hypothesis there exist indices $i_m, \ldots, i_1$ and paths $q = q_m, \ldots, q_0 = \rho$ with $q_{j-1} = e_{i_j}(q_j)$ for $j = m, \ldots, 1$. Let $v = v_0, \ldots, v_l = v'$ be the sequence of paths called $p_0, \ldots, p_l$ in the discussion preceding the statement of rule 4.2.1; put $n = m + l$. It was established there that for all $j < l$ there exists a simple root α such that $v_{j+1} = e_\alpha(v_j)$; let the index of this root be i_{n-j}, thus extending our sequence of indices to $i_n, \ldots, i_1$. We shall say that we are in the exceptional case if $m > 0$, $l > 0$, and $h_{1,0} = \psi_\alpha$ for some simple root α; otherwise we are in the regular case. Define a sequence of paths $p = p_n, \ldots, p_0 = \pi$ as follows: set $p_j = v_{n-j} * q$ for $n \geq j > m$ and $p_j = v' * q_j$ for $m > j \geq 0$; finally set $p_m = v' * q$ in the regular case, and $p_m = v_{l-1} * q_{m-1}$ in the exceptional case.

We shall first show that the only possibility to have $f_{i_j}(v' * q_{j-1}) \neq v' * q_i$ for $0 < j \leq m$ occurs for $j = m$ in the exceptional case, and that we then have $f_{i_m}(v' * q_{m-1}) = v_{l-1} * q_{m-1}$; this will establish $e_{i_j}(p_j) = p_{j-1}$ for $j \leq m$. Putting $\alpha = \alpha_{i_j}$, the operator f_α will only apply to the left factor of $v' * q_{j-1}$ if

$$\lfloor \kappa' + v'^* \rfloor_\alpha < \lfloor \kappa' + q_{j-1} \rfloor_\alpha. \tag{1}$$

Since v'^* is κ'-dominant, the left hand side is non-negative, so (1) can only hold if its right hand side, which equals $\lfloor \kappa' + q_j \rfloor_\alpha + 1$ since $q_{j-1} = e_\alpha(q_j)$, is strictly positive; by the induction hypothesis this only happens when $h_{1,0} = \psi_\alpha$ and $j = m$, and the right hand side then equals 1. Therefore the left hand side must be 0, which excludes case (d) of rule 4.2.1, so that we have $l > 0$ and are in the exceptional case. It can be seen from rule 4.2.1 that $h_{1,0} = \psi_\alpha$ and $l > 0$ imply that $\alpha_{i_{m+1}} = \alpha$ and $\lfloor \kappa + v_l \rfloor_\alpha = 0$; therefore the left hand side of (1) is indeed 0 in this case and $f_\alpha(v') = v_{l-1}$, so that we have $f_\alpha(v' * q_{m-1}) = v_{l-1} * q_{m-1}$ as claimed.

We proceed to show similarly that the only possibility to have $e_{i_{n-j}}(v_j * q) \neq v_{j+1} * q$ for $0 \leq j < l$ occurs for $j = l - 1$ in the exceptional case, and that we then have $e_{i_{m+1}}(v_{l-1} * q) = v_{l-1} * q_{m-1}$; this will establish $e_{i_k}(p_k) = p_{k-1}$ for $k > m$. Putting $\alpha = \alpha_{i_{n-j}}$, the operator e_α will only apply to the right factor of $v_j * q$ if

$$\lfloor \lambda' + v_j^* \rfloor_\alpha > \lfloor \lambda' + q \rfloor_\alpha. \tag{2}$$

Since q is λ'-dominant, the right hand side is non-negative, so (2) can only hold if its left hand side is strictly positive, and in particular $\langle \lambda', \alpha^\vee \rangle > 0$. Since we have $e_\alpha(v_j) = v_{j+1}$, it can be seen from rule 4.2.1 that we must have $j = 0$ or $j = l - 1$; however if $j = 0 \neq l - 1$, we would be in case (b) of that rule with $v_0 = v = \psi_\alpha$, and the left hand side of (2) would be 0, which allows us to conclude $j = l - 1$. Now whichever of the cases (a), (b), or (c) gives $v_l = e_\alpha(v_{l-1})$ with $\langle \lambda', \alpha^\vee \rangle > 0$, it also gives $h_{1,0} = \psi_\alpha$, and makes the left hand side of (2) equal to 1. By the induction hypothesis (including the detail added) the right hand side of (2) will now be 0 if and only if $m > 0$, which means we are in the exceptional case; we then have moreover $\alpha_{i_m} = \alpha$, so that indeed $e_\alpha(v_{l-1} * q) = v_{l-1} * e_\alpha(q) = v_{l-1} * q_{m-1}$, as claimed.

It remains to establish the statements involving $\lfloor \kappa + p_j \rfloor_{\alpha_{i_j}}$ needed for the induction. If $0 < j < m$, or if $j = m$ in the regular case, one has $p_j = v' * q_j$; as v' is κ-dominant with $v'(1) = \kappa' - \kappa$ this implies $\lfloor \kappa + p_j \rfloor_{\alpha_{i_j}} = \lfloor \kappa' + q_j \rfloor_{\alpha_{i_j}}$. Everything then follows immediately from the corresponding part of the induction hypothesis (if $l = 0$ one uses $h_{1,0} = \pi'$). This covers all cases with $l = 0$, so from now on assume $l > 0$. If $m < j < n$, or if $j = n$ and π' is linear, then we have $\lfloor \lambda + v_{n-j} \rfloor_{\alpha_{i_j}} < 0$ by the construction of the sequence $v_0, \ldots, v_l$, and since $p_j = v_{n-j} * q$ this implies $\lfloor \lambda + p_j \rfloor_{\alpha_{i_j}} < 0$. If on the other hand $\pi' = \psi_\alpha$ (so that $\kappa = \lambda$), then we see from cases (b) and (c) of rule 4.2.1 that $\alpha_{i_n} = \alpha$ and $\lfloor \kappa + v \rfloor_\alpha = 0$, which implies $\lfloor \kappa + p \rfloor_\alpha = 0$ since p is λ-dominant. The only case left is $j = m$ in the exceptional case; put $\alpha = \alpha_{i_{m+1}} = \alpha_{i_m}$, so that $h_{1,0} = \psi_\alpha$. One can show $\lfloor \kappa + p_m \rfloor_\alpha = -1$ in various ways, as the minimum is attained at both sides of the concatenation $p_m = v_{l-1} * q_{m-1}$. For instance, we have seen that $\lfloor \kappa + v_l \rfloor_\alpha = 0$ in this case, which implies $\lfloor \kappa + v_{l-1} \rfloor_\alpha = -1$ since $v_l = e_\alpha(v_{l-1})$. This completes our proof. □

§5. Generalisation of jeu de taquin to piecewise linear paths.

We shall now generalise the constructions considered so far to a much larger class of paths than that of the ψ-paths, namely for the entire class Π of piecewise linear paths in the space of rational weigths. The rule that describes the construction in the elementary cases will become simpler than rule 4.2.1, and in fact resembles rule 4.1.1, yet we shall see that the global construction contains construction 4.2.3 as a special case. Given this circumstance, it may seem silly that we went through all the complications of the preceding subsection. There is however an important price that we pay for the simplicity and generality of the new construction: it gives us no direct control over integrality, and therefore does not allow a direct connection to be made with the root operators e_α and f_α.

It turns out that the simplest way to describe the jeu de taquin construction for piecewise linear paths is not using doubly indexed families of paths, or collections of "horizontal" and "vertical" path segments, but using "2-dimensional" paths, that is to say, piecewise linear maps $f \colon [0,1] \times [0,1] \to X_{\mathbf{Q}}$ (here piecewise linear means there is a finite triangulation of $[0,1] \times [0,1]$ such that the restriction of f to each of the triangles is linear). For these maps we do not require (as was done for paths) that they must always "start at 0", but we shall require that their image is contained in the dominant chamber. Then instead of conditions like rule 4.2.1, we shall impose the following somewhat curious functional equation.

5.1. Rule. *For every pair of intervals $[s_0, s_1], [t_0, t_1] \subseteq [0,1]$ such that f is linear on each of the line segments $\{s_0\} \times [t_0, t_1]$ and $[s_0, s_1] \times \{t_1\}$, one has $f(s,t) = \mathrm{dom}_W\bigl(f_0(s,t)\bigr)$ for $(s,t) \in [s_0, s_1] \times [t_0, t_1]$, where f_0 is the linear function given by $f_0(s,t) = f(s_0,t) + f(s,t_1) - f(s_0,t_1)$.*

If we prescribe f on each of the segments $\{s_0\} \times [t_0, t_1]$ and $[s_0, s_1] \times \{t_1\}$ by functions that are linear and everywhere dominant, then $f(s,t) = \mathrm{dom}_W\bigl(f_0(s,t)\bigr)$ (with f_0 as in the rule) defines an extension of f to $[s_0, s_1] \times [t_0, t_1]$ that is piecewise linear and everywhere dominant. In particular this extension determines piecewise linear paths on the edges $[s_0, s_1] \times \{t_0\}$ and $\{s_1\} \times [t_0, t_1]$ of the rectangle $[s_0, s_1] \times [t_0, t_1]$ opposite to those on which f was prescribed. We shall call this operation of extending f across a rectangle $[s_0, s_1] \times [t_0, t_1]$ an elementary extension of f. We still need to show that the condition of rule 5.1 is satisfied for any applicable subintervals of $[s_0, s_1]$ and $[t_0, t_1]$, so let $[s'_0, s'_1] \subseteq [s_0, s_1]$ and $[t'_0, t'_1] \subseteq [t_0, t_1]$ be such that f is linear on $L_1 = \{s'_0\} \times [t'_0, t'_1]$ and on $L_2 = [s'_0, s'_1] \times \{t'_1\}$. We first show that the weights $f_0(s'_0, t'_0)$ and $f_0(s'_1, t'_1)$ are not separated by any wall, i.e., that there is no root β (positive or negative) for which the linear functional $\phi(x,y) = \langle f_0(x,y), \beta^\vee \rangle$ has $\phi(s'_0, t'_0) < 0$ and $\phi(s'_1, t'_1) > 0$. If there were such a root β, then $\phi(s'_0, t'_1) \neq 0$ would contradict the linearity of f either on L_1 or on L_2, whereas $\phi(s'_0, t'_1) = 0$ would imply by linearity that $\phi(s_0, t_0) < 0$ and $\phi(s_1, t_1) > 0$, contradicting the fact that both $f_0(s_0, t_0)$ and $f_0(s_1, t_1)$ are dominant. Therefore, there exists a $w \in W$ such that $f(s,t) = w(f_0(s,t))$ on $L_1 \cup L_2$. Being linear, f_0 satisfies $f_0(s,t) = f_0(s'_0, t) + f_0(s, t'_1) - f_0(s'_0, t'_1)$; hence the validity of rule 5.1 is established by the following computation for $(s,t) \in [s'_0, s'_1] \times [t'_0, t'_1]$:

$$\begin{aligned} f(s,t) = \mathrm{dom}_W\bigl(w(f_0(s,t))\bigr) &= \mathrm{dom}_W\bigl(w\bigl(f_0(s'_0,t) + f_0(s,t'_1) - f_0(s'_0,t'_1)\bigr)\bigr) \\ &= \mathrm{dom}_W\bigl(f(s'_0,t) + f(s,t'_1) - f(s'_0,t'_1)\bigr). \end{aligned}$$

Because of the way the rule is formulated, there is no need for a counterpart of construction 4.2.3: a piecewise linear function f defined on $[0,1]\times[0,1]$ that satisfies the rule must match any function constructed by repeated elementary extensions from the restriction of f to the edges $\{0\}\times[0,1]$ and $[0,1]\times\{1\}$ of the unit square. However, it is not immediately obvious that repeated elementary extensions suffice to cover all of the unit square. To see the difficulty, imagine that *every* elementary extension would result at each of the opposite edges of the rectangle in a path consisting of two different linear parts; then infinitely many elementary extensions could be applied, but they would fail to define f beyond a certain subset with fractal boundary. We shall show that this cannot happen; to do so we need to consider the directions of the segments of the paths obtained by elementary extension. For a dominant weight λ and weight μ in its orbit $W\lambda$, the set $\{\, w\in W \mid \mu=w(\lambda)\,\}$ is a coset in W/W_λ, and it does not change when μ is multiplied by a positive scalar. We define this coset to be the direction of π_μ, or of any translate of a positive multiple of π_μ, and endow W/W_λ with the Bruhat order, and the associated length function l. Then the following lemma is an immediate consequence of the definition of elementary extensions.

5.2. Lemma. *Let f satisfy rule 5.1 on $[s_0,s_1]\times[t_0,t_1]$ and be linear on $\{s_0\}\times[t_0,t_1]$ and $[s_0,s_1]\times\{t_1\}$; let the direction of the path defined by f along $[s_0,s_1]\times\{t_1\}$ be τ, and let the path defined by f along $[s_0,s_1]\times\{t_0\}$ be $\pi_0*\cdots*\pi_n$ where the π_i are linear paths with differenct directions τ_i. Then $\tau\le\tau_0<\cdots<\tau_n$, and consequently $n\le l(\tau)$. Similar statements hold for the other two edges.* □

5.3. Theorem/construction (jeu de taquin for piecewise linear paths). *Let κ,λ,ν be dominant integral weights for $\mathfrak{g}$, π' a κ-dominant piecewise linear path with $\pi'(1)=\lambda-\kappa$, and p a λ-dominant piecewise linear path with $p(1)=\nu-\lambda$. There is a unique piecewise linear function $f\colon[0,1]\times[0,1]\to X_{\mathbf{Q}}$ with $f(0,t)=\kappa+\pi'(t)$ and $f(t,1)=\lambda+p(t)$ for $t\in[0,1]$ that satisfies rule 5.1. Putting $\mu=f(1,0)$, we may define a κ-dominant piecewise linear path π with $\pi(1)=\mu-\kappa$ by $\pi(t)=f(t,0)-\kappa$, and a μ-dominant piecewise linear path p' with $p'(1)=\nu-\mu$ by $p(t)=f(1,t)-\mu$.*

Proof. We shall show that by repeatedly applying elementary extensions to f we succeed after a finite number of step in finding an extension of f to all of $[0,1]\times[0,1]$, which is then automatically unique. By a trivial induction on the number of linear segments from which p is concatenated, we may reduce to the case that p is linear. We cannot continue with a similar induction on the number of segments of π' however; instead we apply induction on the length $l(\tau)$ of the direction τ of the path p. For $l(\tau)=0$, which is equivalent to $p\in\mathcal{P}^+$, the function f given by $f(s,t)=\kappa+\pi'(t)+p(s)$ is everywhere dominant, and is therefore the unique function satisfying the conditions in the theorem. Now suppose $l(\tau)=l>0$. For each fixed value of l we apply induction on the number of linear segments from which π' is concatenated. If π' is linear, elementary extension suffices to define f uniquely on $[0,1]\times[0,1]$. Otherwise we apply elementary extension to p and the final linear segment of π'. Let the piecewise linear path obtained at the side opposite to p be $p_0*\cdots*p_n$; let the segment p_i give the values of f on $[s_i,s_{i+1}]\times\{t_1\}$ and have direcion τ_i ($i=0,\ldots,n$). Because of lemma 5.2 we have $l(\tau_0)\le l$ and $l(\tau_i)<l$ for $i>0$. In case $l(\tau_0)=l$ we can extend f to $[s_0,s_1]\times[0,t_1]$ by induction on the number of segments of π'; for the remaining segments p_i (including p_0 if $l(\tau_0)<l$) we can apply the hypothesis of induction with respect to $l(\tau)$, and conclude that we can extend f successively to the rectangles $[s_i,s_{i+1}]\times[0,t_1]$; this defines f uniquely on $[0,1]\times[0,1]$. □

It is not difficult to see that the rules 4.1.1 and 4.2.1 can obtained as instances of construction 5.3; as a consequence, that construction generalises the constructions 4.1.3 and 4.2.3. The symmetries of those constructions are preserved, since one can show with similar arguments as we gave to show that the function obtained by elementary extension satisfies rule 5.1, that the transpose function $f'(s,t)=f(t,s)$ also satisfies that rule.

5.4. Theorem (symmetry of jeu de taquin for piecewise linear paths). *If a piecewise linear function f satisfies rule 5.1, then the transpose function f' defined by $f'(s,t)=f(t,s)$ satisfies that rule as well. In particular, the construction 5.3 is its own inverse: if when applied to (κ,π',p) it returns (π,p'), then applied to (κ,π,p') it will return (π',p). Moreover, construction 5.3 is symmetrical with respect to dualisation of paths: when applied to (ν,p^*,π'^*) it will return (p'^*,π^*).* □

Since construction 5.3 is does not refer to any integrality condition at all, it is not easy to relate it directly to the root operators e_α and f_α. It is for instance not true that the paths π and p are related by a sequence of applications of such operators, not even in the case of an elementary extension. What one gets instead is a sequence of applications of fractional powers e_α^x or f_α^x of root operators with $x\in\mathbf{Q}$;

here e_α^x maps a path p to a non-zero value if and only if $\lfloor p \rfloor_\alpha \leq -x$, in which case one has

$$e_\alpha^x(p)(s) = p(s) + \max\left(0, x + \lfloor p \rfloor_\alpha - \min_{t\in[0,s]} \langle p(t), \alpha^\vee \rangle\right) \alpha.$$

A relation with root operators which seems plausible, and which we hope to establish in further work, can be formulated as follows. Let us call a path $p \in \Pi$ of integral shape if $p \in B_\pi$ for some $\pi \in \mathcal{P}^+$. Then if $\kappa \in X$, and if π' and p are of integral shape, then so are the paths π and p' obtained from κ, π', and p by construction 5.3; moreover π can be obtained from p by a sequence of applications of operators e_α, and p' from π' by applications of operators f_α. This would imply in particular that if $\kappa = 0$ then $p \in B_\pi$ and $p' \in B_{\pi'}$, and that the construction defines, for any fixed pair of paths $\pi, \pi' \in \mathcal{P}^+$ with $\pi'(1) = \lambda$ and $\pi(1) = \mu$, a bijection between the λ-dominant paths $p \in B_\pi$ and the μ-dominant paths $p' \in B_{\pi'}$, with moreover $\lambda + p(1) = \mu + p'(1)$. It would also imply that any class of paths of integral shape, that is closed under the root operators (such as for instance the class of Lakshmibai-Seshadri paths), would also be closed under construction 5.3, which would therefore give rise to a special instance of it, similar to constructions 4.1.3 and 4.2.3.

References.

[KaNa] M. Kashiwara and T. Nakashima, "Crystal graphs for representations of the q-analogue of classical Lie algebras", *Journal of Algebra* **165**, (1994), 295–345.

[vLee1] M. A. A. van Leeuwen, "The Robinson-Schensted and Schützenberger algorithms, an elementary approach", *Electronic J. of Combinatorics* **3 (no. 2)**, R15, (1996), 32 pp.

[vLee2] M. A. A. van Leeuwen, "Tableau algorithms defined naturally for pictures", *Discrete Mathematics* **157**, (1996), 321–362.

[vLee3] M. A. A. van Leeuwen, *Edge sequences, ribbon tableaux, and an action of affine permutations*, Report MAS-R9707, CWI, Amsterdam. To appear in *European Journal of Combinatorics*.

[vLee4] M. A. A. van Leeuwen, *Some bijective correspondences involving domino tableaux*, Report MAS-R9708, CWI, Amsterdam.

[Litt1] P. Littelmann, "A Littlewood-Richardson rule for symmetrizable Kac-Moody algebras", *Inventiones mathematicae* **116**, (1994), 329–346.

[Litt2] P. Littelmann, "Paths and root operators in representation theory", *Annals of Mathematics* **142**, (1995), 499–525.

[Litt3] P. Littelmann, "A plactic algebra for semisimple Lie algebras", *Advances in Math.* **124 (2)**, (Dec. 1996), 312–331.

[Macd] I. G. Macdonald, *Symmetric Functions and Hall Polynomials*, Oxford Mathematical Monographs, Clarendon press, Oxford, 1979.

[Rob] G. de B. Robinson, "On the representations of the symmetric group", *American Journal of Math.* **60**, (1938), 745–760.

[Zel1] A. V. Zelevinsky, "A Generalisation of the Littlewood-Richardson Rule and the Robinson-Schensted-Knuth Correspondence", *Journal of Algebra* **69**, (1981), 82–94.