The Littlewood-Richardson rule, and related combinatorics

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

An introduction is given to the Littlewood-Richardson rule, and various combinatorial constructions related to it. We present a proof based on tableau switching, dual equivalence, and coplactic operations. We conclude with a section relating these fairly modern techniques to earlier work on the Littlewood-Richardson rule.

1991 Mathematics Subject Classification: 05E05, 05E10.

Keywords and Phrases: Schur functions, Littlewood-Richardson rule, jeu de taquin, tableau switching, dual equivalence, coplactic operations, pictures, Robinson-Schensted correspondence.

§0. Introduction.

The Littlewood-Richardson rule is a combinatorial rule describing the multiplication of Schur polynomials; it was first formulated in [LiRi], but its general validity remained unproved for several decades. The various proofs that have been given since have created a rich combinatorial theory, with many interrelated constructions, including the Robinson-Schensted correspondence, and *jeu de taquin*. We describe several such constructions, and use them to prove the rule; we are not however narrowly focused on this proof, and discuss several topics that are not used in it. We make use of certain “modern” (post-1980) notions, while we do not treat some other notions that figure prominently in many other proofs of the Littlewood-Richardson rule, and are well documented elsewhere (e.g., [Fult], [Fom]); specifically, we focus on tableau switching (which includes jeu de taquin), dual equivalence, and coplactic operations*, but do not introduce Schensted’s algorithm, Knuth equivalence, or Greene’s poset invariant. We do not really define or use Zelevinsky’s pictures, but they are mentioned in several places, and have inspired much of our work. At the end, we give an overview of earlier work on the subject, to help clarify its relation to our approach.

The proof we give does not require many technical verifications; moreover, we believe that our main theorem 3.3.1 makes the correspondences considered more transparent. We do prove in detail some properties of the constructions that are so basic that they might have been left to the reader; this is because we feel that it is often these “low-level” properties that really explain why the more significant theorems work. Most facts presented in this paper are known (at least to experts), although for some it is hard to find a published reference. Nonetheless the global structure of our proof, and theorem 3.3.1, seem to be new, even if the latter is related to the known fact that “coplactic operations are compatible with plactic equivalence” (which in fact motivates their name).

The remainder of this paper is organised in four sections: in §1 we formulate the Littlewood-Richardson rule; in §2 we define tableau switching and derive an expression for Littlewood-Richardson coefficients in terms of jeu de taquin; in §3 we define coplactic operations and establish our main theorem, which implies the Littlewood-Richardson rule; finally in §4 we comment on earlier work.

A word on notation: we always start indexing at 0; in particular this applies to parts of partitions, and rows, columns, and entries of tableaux. We define $[n] = \{ i \in \mathbf{N} \mid i < n \}$ for $n \in \mathbf{N}$.

* Arguably coplactic operations are not at all new: even if not considered in isolation, they do occur in various forms and contexts, as a part of larger constructions; see [LiRi], [Rob], [BrTeKr], [GrKl], [Thom2], [LaSch].

§1. Formulation of the Littlewood-Richardson rule.

1.1. *Symmetric polynomials and partitions.*

We fix some $n \in \mathbf{N}$, and let $\{ X_i \mid i \in [n] \}$ be a set of n indeterminates. The symmetric group $\mathbf{S}_n$ acts on this set, and hence on the ring $\mathbf{Z}[X_0, \ldots, X_{n-1}]$, by permuting the indeterminates. A polynomial $P \in \mathbf{Z}[X_0, \ldots, X_{n-1}]$ is called symmetric if it is fixed by every $\pi \in \mathbf{S}_n$. We shall denote by Λ_n the set of symmetric polynomials, which is a subring of $\mathbf{Z}[X_0, \ldots, X_{n-1}]$ (the ring of invariants for the action of $\mathbf{S}_n$); since the the action of $\mathbf{S}_n$ preserves the natural grading of $\mathbf{Z}[X_0, \ldots, X_{n-1}]$ by total degree, Λ_n is a graded ring, and we shall denote by Λ_n^d the set of homogeneous symmetric polynomials of degree d. Our interest will be in an explicit description of the multiplicative structure of this ring, expressed on a particular $\mathbf{Z}$-basis, that of the so-called Schur polynomials. But let us first consider some other $\mathbf{Z}$-bases.

The simplest symmetric polynomials are those which are formed as the sum over an orbit of a monomial in $\mathbf{Z}[X_0, \ldots, X_{n-1}]$; we shall call these minimal symmetric polynomials. The monomials in $\mathbf{Z}[X_0, \ldots, X_{n-1}]$ are of the form $X_0^{\alpha_0} \cdots X_{n-1}^{\alpha_{n-1}}$, which we shall abbreviate to X^α, where $\alpha = (\alpha_0, \ldots, \alpha_{n-1}) \in \mathbf{N}^n$. We shall write $|\alpha| = \deg X^\alpha = \sum_{i \in [n]} \alpha_i$. In order to select a specific representative within each orbit of monomials, we define a partial ordering '$\leq$' on the set of monomials X^α; or equivalently on the set $\mathbf{N}^n$ of multi-exponents α. This ordering is generated by the relations $X_i > X_{i+1}$ and the condition that $X^\alpha < X^\beta$ implies $MX^\alpha < MX^\beta$ for any monomial M (as a consequence, monomials of distinct degrees are always incomparable); explicitly, one has $X^\alpha \leq X^\beta$ if and only if $\sum_{i<k} \alpha_i \leq \sum_{i<k} \beta_i$ for $k \in [n]$, and $|\alpha| = |\beta|$. Every $\mathbf{S}_n$-orbit of monomials contains a maximum for '$\leq$', which is a monomial X^λ with $\lambda_0 \geq \cdots \geq \lambda_{n-1} \geq 0$; such a $\lambda \in \mathbf{N}^n$ is called a *partition* of $d = |\lambda|$ into n parts, written $\lambda \in \mathcal{P}_{d,n}$. We define the minimal symmetric polynomial $m_\lambda(n) = \sum_{\alpha \in \mathbf{S}_n \cdot \lambda} X^\alpha$, where $\mathbf{S}_n \cdot \lambda$ denotes the $\mathbf{S}_n$-orbit of λ in $\mathbf{N}^n$. The $m_\lambda(n)$, for $\lambda \in \mathcal{P}_{d,n}$, form a $\mathbf{Z}$-basis of Λ_n^d. A partition λ of d (without qualification, written $\lambda \in \mathcal{P}_d$) is defined as a sequence $(\lambda_i)_{i \in \mathbf{N}}$ of natural numbers (called parts) with $\lambda_i \geq \lambda_{i+1}$ for all i and $\lambda_m = 0$ for some $m \in \mathbf{N}$, and with $\sum_{i \in [m]} \lambda_i = d$. A partition is denoted by the parenthesised list of its parts, with trailing zeros omitted; $\mathcal{P}_{d,n}$ is identified with the subset of $\mathcal{P}_d$ of partitions that have at most n non-zero parts. Finally, we put $\mathcal{P} = \bigcup_{d \in \mathbf{N}} \mathcal{P}_d$.

In the special case that λ is the partition of $d \leq n$ for which all non-zero parts are 1 (so there are d such parts), the minimal symmetric polynomial $m_\lambda(n)$ is called the d-th elementary symmetric polynomial, and written $e_d(n)$. The "fundamental theorem on symmetric functions" states that the polynomials $e_i(n)$, for $1 \leq i \leq n$, generate Λ_n as a ring, and are algebraically independent; in other words, Λ_n is isomorphic as a graded ring to $\mathbf{Z}[Y_1, \ldots, Y_n]$ with $\deg Y_i = i$, the isomorphism sending Y_i to $e_i(n)$. We shall however not make use of this fact (we refer to [Macd] for many more facts about symmetric polynomials and functions). We may also express $e_d(n)$ as $\sum_{0 \leq i_0 < \cdots < i_{d-1} < n} X_{i_0} \cdots X_{i_{d-1}}$. If in this expression we replace the strict inequalities between the indices by weak ones, then we obtain another symmetric polynomial $h_d(n) = \sum_{0 \leq i_0 \leq \cdots \leq i_{d-1} < n} X_{i_0} \cdots X_{i_{d-1}}$, called the d-th complete symmetric polynomial. The fact that this is indeed a symmetric polynomial follows from the fact that it contains every monomial X^α with $|\alpha| = d$ exactly once, whence $h_d(n) = \sum_{\lambda \in \mathcal{P}_{d,n}} m_\lambda(n)$. Note that if we would replace only some strict inequalities by weak ones, the result would not be a symmetric polynomial. Like the $e_i(n)$, the $h_i(n)$ for $1 \leq i \leq n$ form a polynomial basis for Λ_n. We saw that $\{ m_\lambda(n) \mid \lambda \in \mathcal{P}_{d,n} \}$ is a $\mathbf{Z}$-basis of Λ_n^d; since monomials of degree d in the generators $e_i(n)$ or $h_i(n)$ are naturally parametrised by $\mathcal{P}_{d,n}$, each part i representing a factor $e_i(n)$ or $h_i(n)$, we have two more such bases parametrised by the same set.

1.2. *Semistandard tableaux and Schur polynomials.*

The Schur polynomials of degree d form yet another $\mathbf{Z}$-basis of Λ_n parametrised by $\mathcal{P}_{d,n}$. Although their significance is not immediately obvious from a purely ring-theoretic perspective, they are of fundamental importance in many situations where the ring Λ_n is encountered. For instance, Λ_n occurs as the character ring of polynomial $\mathbf{GL}_n(\mathbf{C})$ representations, and the Schur polynomials are the irreducible characters. They can be defined as quotients of alternating polynomials, or be expressed in terms of power sums $(m_d(n))$ using symmetric group characters; for our purposes however, a purely combinatorial description will serve best, and it is in that way that we shall define Schur polynomials. For those who wish some motivation for this definition, we refer to places where it is respectively deduced from an algebraic definition [Macd, I (5.12)], from an explicit construction of irreducible $\mathbf{GL}_n(\mathbf{C})$ representations [Fult], and even in an axiomatic approach [Zel2].

Our definition of Schur polynomials is rather similar to the description of $e_d(n)$ and $h_d(n)$ as the sum of a collection of monomials $X_{i_0} \cdots X_{i_{d-1}}$; indeed $e_i(n)$ and $h_i(n)$ are instances of Schur polynomials. The linear sequence of strict respectively weak inequalities relating the indices $i_0, \ldots, i_{n-1}$ for $e_i(n)$ and $h_i(n)$ are replaced for general Schur polynomials by a more complicated mixture of strict and weak inequalities. As we remarked above, this does not always give rise to a symmetric polynomial; we shall see however that it will do so when the inequalities follow a specific pattern associated to certain d-element subsets of $\mathbf{N}^2$ called diagrams. For the Schur polynomial $s_\lambda(n)$ (with $\lambda \in \mathcal{P}_d$), this will be the *Young diagram* $Y(\lambda)$ of λ, defined as $\{\, (i,j) \in \mathbf{N}^2 \mid j \in [\lambda_i] \,\}$. We display Young diagrams as sets of squares in the plane, arranged like matrix entries: for a square (i,j), the row index is i (increasing downwards) and the column index is j (increasing to the right); the reader is warned however that other display conventions can be found in the literature. One obtains a sequence of left-justified rows of successive lengths $\lambda_0, \lambda_1, \ldots$; e.g., for $\lambda = (4,2,1)$ we display $Y(\lambda)$ as $\begin{smallmatrix}\square&\square&\square&\square\\ \square&\square&&\\ \square&&&\end{smallmatrix}$. For future reference we mention the notion of *diagonals* in $\mathbf{N}^2$, defined as sets $\{\, (i,j) \in \mathbf{N}^2 \mid j - i = k \,\}$ for some constant k. The value k is called the *index* of a diagonal, and we consider the set of all diagonals as totally ordered by their indices.

For defining $s_\lambda(n)$, an index of summation, ranging over $[n]$, is associated with each square of $Y(\lambda)$. Thus, each term is identified by an assignment $T\colon (i,j) \mapsto T_{i,j}$ that can be displayed by writing each index $T_{i,j}$ as entry into its square (i,j); the indices are subject to the conditions $T_{i,j} < T_{i+1,j}$ (strict increase down columns) and $T_{i,j} \leq T_{i,j+1}$ (weak increase along rows) whenever both referenced indices exist. An assignment T satisfying these conditions is called a *semistandard Young tableau* of shape λ and entries in $[n]$; the set of all such tableaux is denoted by $\mathrm{SST}(\lambda, n)$. For instance, $\begin{smallmatrix}0&0&1&1\\ 1&2&&\\ 2&&&\end{smallmatrix}$ depicts a semistandard Young tableau of shape $(4,2,1)$ and entries in $[3]$, i.e., an element of $\mathrm{SST}((4,2,1),3)$. Like in the case of $e_d(n)$ and $h_d(n)$, each term in the summation will be a monomial, containing a factor X_k for each occurrence of k as summation index. Therefore we define for $T \in \mathrm{SST}(\lambda,n)$ its *weight* $\alpha = \mathrm{wt}\, T \in \mathbf{N}^n$ to be such that $\prod_{(i,j)\in Y(\lambda)} X_{T_{i,j}} = X^\alpha$, in other words α_k counts the occurrences of the entry k in T; the tableau just depicted has weight $(2,3,2)$. Then for $\lambda \in \mathcal{P}_{d,n}$, the Schur polynomial $s_\lambda(n)$ is defined by

$$s_\lambda(n) = \sum_{T \in \mathrm{SST}(\lambda,n)} X^{\mathrm{wt}\, T}. \tag{1}$$

For instance, by enumerating $\mathrm{SST}((4,2,1),3)$ one finds that $s_{(4,2,1)}(3)$ is a symmetric polynomial with 15 terms, which equals $m_{(4,2,1)}(3) + m_{(3,3,1)}(3) + 2m_{(3,2,2)}(3)$. The $\mathbf{S}_n$-invariance of $s_\lambda(n)$ is not at all evident from the definition, however. Although this will follow from properties independently derived later, let us prove this key fact right now.

1.2.1. Proposition. *The Schur polynomials are symmetric polynomials, i.e., $s_\lambda(n) \in \Lambda_n^d$ for $\lambda \in \mathcal{P}_{d,n}$.*

Proof. It suffices to prove for any $k < n-1$ that $s_\lambda(n)$ is invariant under the interchange of X_k and X_{k+1}; to this end we construct an involution of the set $\mathrm{SST}(\lambda,n)$, that realises an interchange of the components α_k and α_{k+1} of the weight α. We shall leave all entries $T_{i,j} \notin \{k, k+1\}$ unchanged, as well as the entries k and $k+1$ in any column of T that contains both of them. It is readily checked that the set of squares containing the remaining entries (i.e., entries k or $k+1$ that as such are unique in their column) meets any given row in a contiguous sequence of squares. Let the entries of that sequence be r times k followed by s times $k+1$ (r or s might be 0); we replace them by s times k followed by r times $k+1$. The transformation of T consists of performing this change independently for each row; clearly the operation is an involution, and has the desired effect on $\mathrm{wt}\, T$. □

Remark. This involution, which was introduced in [BeKn], is simple to describe, but not the best one from a mathematical point of view, as it does not give rise to an action of $\mathbf{S}_n$ on $\mathrm{SST}(\lambda,n)$ (denoting interchange of i and $i+1$ by s_i, one may check that application of $s_0 s_1 s_0$ and $s_1 s_0 s_1$ to the tableau T shown above gives different results). There is another involution, the *relèvement plaxique* of [LaSch, §4], that does extend to an $\mathbf{S}_n$-action; it is related to the coplactic operations discussed in §3 below.

Other properties of Schur polynomials are easy to establish. One has $\mathrm{wt}\, T \leq \lambda$ for all $T \in \mathrm{SST}(\lambda, n)$, with equality for exactly one such T, namely the tableau with $T_{i,j} = i$ for all $(i,j) \in Y(\lambda)$; this tableau will be denoted by $\mathbf{1}_\lambda$. This fact follows from the observation that in any $T \in \mathrm{SST}(\lambda, n) \setminus \{\mathbf{1}_\lambda\}$, one can find at least one entry $k+1$ that can be replaced by k (for some $k < n-1$). Being a symmetric polynomial, $s_\lambda(n)$ can be written as a sum of terms $m_\mu(n)$ for $\mu \in \mathcal{P}_{d,n}$ with integer coefficients. These coefficients are all non-negative, they are zero unless $\mu \leq \lambda$, and the coefficient of $m_\lambda(n)$ is equal to 1: the transition from the $s_\lambda(n)$ to the $m_\lambda(n)$ is "unitriangular" with respect to '$\leq$'. Then the fact that $\{\, m_\lambda(n) \mid \lambda \in \mathcal{P}_{d,n} \,\}$ is a $\mathbf{Z}$-basis of Λ_n, implies that $\{\, s_\lambda(n) \mid \lambda \in \mathcal{P}_{d,n} \,\}$ is one as well.

We can now state the problem with which the Littlewood-Richardson rule is concerned, as expressing multiplication in Λ_n on the basis of the Schur polynomials. In somewhat more detail: given $\lambda \in \mathcal{P}_{d,n}$ and $\mu \in \mathcal{P}_{d',n}$, we wish to determine the integer coefficients $c^\nu_{\lambda,\mu}$ for all $\nu \in \mathcal{P}_{d+d',n}$, such that

$$s_\lambda(n) s_\mu(n) = \sum_{\nu \in \mathcal{P}_{d+d',n}} c^\nu_{\lambda,\mu} s_\nu(n). \tag{2}$$

We have suppressed n in the notation $c^\nu_{\lambda,\mu}$, since it will turn out that this coefficient is independent of n (although n must be sufficiently large for $c^\nu_{\lambda,\mu}$ to appear in the formula in the first place). More generally, for identities valid for any number n of indeterminates, we shall sometimes write h_d for $h_d(n)$ and s_λ for $s_\lambda(n)$, etc. (there is an algebraic structure called the ring of symmetric functions that justifies this notation, see [Macd], but we shall not discuss it here). The $c^\nu_{\lambda,\mu}$ are called Littlewood-Richardson coefficients. Representation theoretic considerations show that $c^\nu_{\lambda,\mu} \in \mathbf{N}$; the Littlewood-Richardson rule will in fact describe the $c^\nu_{\lambda,\mu}$ as the cardinalities of certain combinatorially defined sets.

1.3. *Skew shapes, and skew Schur polynomials.*

In order to formulate the Littlewood-Richardson rule, we need to extend the class of tableaux beyond that of the semistandard Young tableaux. Giving $\mathbf{N}^2$ its natural partial ordering (simultaneous comparison of both coordinates), Young diagrams can be characterised as its finite order ideals (if $r \in Y(\lambda)$, then also $q \in Y(\lambda)$ for all $q \leq r$). For the class of skew diagrams, this is relaxed to convexity with respect to '$\leq$': a *skew diagram* D is a finite subset of $\mathbf{N}^2$ for which $p, r \in D$ and $p \leq q \leq r$ imply $q \in D$. The summation analogous to (1) using a skew diagram will still give rise to a symmetric polynomial (one may for instance check that the proof of proposition 1.2.1 remains valid). A skew diagram can always be written as the difference of two Young diagrams: $D = Y(\lambda) \setminus Y(\mu)$ with $Y(\mu) \subseteq Y(\lambda)$; this representation is not unique in general (although in some cases it is, for instance when the sets of rows and columns meeting D are both initial intervals of $\mathbf{N}$). In many cases, for instance when considering the shapes of tableaux, it will be important to fix the partitions λ, μ used to represent a skew diagram, which leads us to define the related but distinct notion of a skew shape.

A *skew shape* is a symbol λ/μ where $\lambda, \mu \in \mathcal{P}$ and $Y(\mu) \subseteq Y(\lambda)$; the set of all skew shapes will be denoted by $\mathcal{S}$. For $\lambda/\mu \in \mathcal{S}$ we define $Y(\lambda/\mu) = Y(\lambda) \setminus Y(\mu)$ and $|\lambda/\mu| = |Y(\lambda/\mu)| = |\lambda| - |\mu|$. A shape of the form $\lambda/(0)$ is called a partition shape; occasionally we consider partitions as skew shapes, in which case λ is identified with $\lambda/(0)$. It is a trivial but useful fact that if $|\lambda/\mu| \leq 3$, no diagonal can meet $Y(\lambda/\mu)$ in more than one square. We shall say that a skew shape χ represents the product $\chi_0 * \chi_1$ of two other skew shapes if $Y(\chi)$ can be written as a disjoint union of skew diagrams $\bar\chi_0$ and $\bar\chi_1$, where $\bar\chi_k$ is obtained by some translation from $Y(\chi_k)$ $(k = 0, 1)$, while for all $(i,j) \in \bar\chi_0$ and $(i',j') \in \bar\chi_1$ one has $i > i'$ and $j < j'$. For instance,

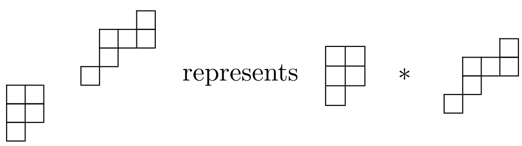

One could define an equivalence relation on $\mathcal{S}$ such that this relation defines a monoid structure on the quotient set, but this would be cumbersome, and it is not really needed for our purposes.

Let $\chi = \lambda/\mu \in \mathcal{S}$; a *skew semistandard tableau* T of shape χ and with entries in $[n]$ is given by specifying χ itself, together with a map $Y(\chi) \to \mathbf{N}$ written $(i,j) \mapsto T_{i,j}$, satisfying $T_{i,j} \in [n]$, $T_{i,j} < T_{i+1,j}$ and $T_{i,j} \leq T_{i,j+1}$ whenever these values are defined. The set of all such T is denoted by $\mathrm{SST}(\chi, n)$, and $\mathrm{SST}(\chi) = \bigcup_{n\in\mathbf{N}} \mathrm{SST}(\chi, n)$; we identify $\mathrm{SST}(\lambda)$ with $\mathrm{SST}(\lambda/(0))$. The weight $\alpha = \operatorname{wt} T \in \mathbf{N}^n$ of $T \in \mathrm{SST}(\chi, n)$ is defined by $\prod_{(i,j)\in Y(\chi)} X_{T_{i,j}} = X^\alpha$, and the skew Schur polynomial $s_\chi(n)$ by

$$s_\chi(n) = \sum_{T\in\mathrm{SST}(\chi,n)} X^{\operatorname{wt} T}. \tag{3}$$

Note that unlike in (1), there is no restriction to $\mathcal{P}_{d,n}$ here, for the partitions λ, μ forming χ; the skew Schur polynomial will be non-zero as long as there are no columns in $Y(\chi)$ of length exceeding n. If a skew shape χ represents the product $\chi_0 * \chi_1$ of two other skew shapes, then there is an obvious weight preserving bijection $\mathrm{SST}(\chi) \to \mathrm{SST}(\chi_0) \times \mathrm{SST}(\chi_1)$; therefore $s_\chi = s_{\chi_0} s_{\chi_1}$. In this case we shall denote s_χ by $s_{\chi_0 * \chi_1}$, or if $\chi_0 = (\lambda/(0))$ and $\chi_1 = (\mu/(0))$, by $s_{\lambda*\mu}$; our problem can be restated as finding the decomposition of the skew Schur polynomial $s_{\lambda*\mu}(n)$ as a sum of ordinary Schur polynomials. In fact the Littlewood-Richardson rule will describe the decomposition of any skew Schur polynomial $s_\chi(n)$.

1.3.1. Definition. *For any skew shape χ and all $\nu \in \mathcal{P}_{|\chi|}$, the numbers c^ν_χ are defined as the coefficients appearing in the decomposition formula*

$$s_\chi = \sum_{\nu\in\mathcal{P}_{|\chi|}} c^\nu_\chi s_\nu. \tag{4}$$

It is clear that if χ represents $(\lambda/(0)) * (\mu/(0))$, then one has $c^\nu_\chi = c^\nu_{\lambda,\mu}$.

1.4. *Littlewood-Richardson tableaux.*

Let $\lambda/\mu \in \mathcal{S}$, and $\nu \in \mathcal{P}$ with $|\lambda/\mu| = |\nu|$; we shall now introduce the objects that are counted by $c^\nu_{\lambda/\mu}$, which will be called Littlewood-Richardson tableaux of shape λ/μ and weight ν. We need a preliminary definition. For $T \in \mathrm{SST}(\lambda/\mu)$ and $k, l \in \mathbf{N}$, define T^l_k to be the number of entries l in row k of T, i.e., the cardinality of the set $\{\, j \in [\lambda_k] \setminus [\mu_k] \mid T_{k,j} = l \,\}$.

1.4.1. Definition. *Two tableaux $T \in \mathrm{SST}(\lambda/\mu)$ and $\bar{T} \in \mathrm{SST}(\nu/\kappa)$ are called companion tableaux if $T^l_k = \bar{T}^k_l$ for every $k, l \in \mathbf{N}$. In this case T is called ν/κ-dominant (and $\bar{T}$ is λ/μ-dominant).*

Here is an example of a pair of companion tableaux, with a table of the pertinent values $T^l_k = \bar{T}^k_l$:

$$T = \begin{array}{ccccc} & & & \boxed{0} & \boxed{1} \\ & \boxed{0} & \boxed{1} & \boxed{1} & \boxed{3} \\ \boxed{0} & \boxed{2} & \boxed{2} & \boxed{3} & \\ \boxed{1} & \boxed{4} & \boxed{4} & \boxed{5} & \\ \boxed{3} & \boxed{5} & & & \end{array}, \quad \bar{T} = \begin{array}{ccccccccc} & & & & & & \boxed{0} & \boxed{1} & \boxed{2} \\ & & & & \boxed{0} & \boxed{1} & \boxed{1} & \boxed{3} & \\ & & & & \boxed{2} & \boxed{2} & & & \\ & & \boxed{1} & \boxed{2} & \boxed{4} & & & & \\ & \boxed{3} & \boxed{3} & & & & & & \\ \boxed{3} & \boxed{4} & & & & & & & \end{array}, \tag{5}$$

$$\left(T^l_k\right)_{\substack{0\leq k<5 \\ 0\leq l<6}} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 2 & 0 & 1 & 0 & 0 \\ 1 & 0 & 2 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \end{pmatrix}. \tag{6}$$

Note that if T is ν/κ-dominant, then $\operatorname{wt} T = \nu - \kappa$, since $(\operatorname{wt} T)_j = \sum_i T^j_i = \sum_i \bar{T}^i_j = \nu_j - \kappa_j$. We shall simply say that T is κ-dominant if it is ν/κ-dominant for $\nu = \kappa + \operatorname{wt} T$; this notion is used in [Litm2]. T may have companion tableaux of different shapes, but at most one of a given shape ν/κ: the multiset of entries of any row is determined by T, and they must be weakly increasing. Therefore different companion tableaux of T differ only by horizontal slides of their rows; in the above example one could for instance shift the first 4 rows of $\bar{T}$ one place to the left. To test for κ-dominance of a tableau T, it suffices to construct the unique candidate for $\bar{T}$ and check that its columns are strictly increasing.

1.4.2. Definition. *A tableau $L \in \mathrm{SST}(\lambda/\mu)$ is a Littlewood-Richardson tableau if it is (0)-dominant. The set of all $\nu/(0)$-dominant tableaux in $\mathrm{SST}(\lambda/\mu)$ is denoted by $\mathrm{LR}(\lambda/\mu, \nu)$.*

Here is an example Littlewood-Richardson tableau L, and its companion Young tableau $\bar{L}$.

$$L = \begin{array}{cccccc} & & & & 0 & 0 \\ & & 0 & 1 & 1 & \\ & 0 & 1 & 2 & 2 & \\ 0 & 1 & 3 & & & \\ 2 & 4 & & & & \end{array}\quad, \qquad \bar{L} = \begin{array}{ccccc} 0 & 0 & 1 & 2 & 3 \\ 1 & 1 & 2 & 3 & \\ 2 & 2 & 4 & & \\ 3 & & & & \\ 4 & & & & \end{array}\quad. \tag{7}$$

The given definition of Littlewood-Richardson tableaux is not the traditional one, but that definition can be easily derived from it, as follows. A pair of companion tableaux $T \in \mathrm{SST}(\lambda/\mu)$ and $\bar{T} \in \mathrm{SST}(\nu/\kappa)$ will determine a bijection $p: Y(\lambda/\mu) \to Y(\nu/\kappa)$, if one fixes for each k, l a bijection between the sets of squares whose entries are counted by T^l_k and $\bar{T}^k_l$. Both these sets are contiguous sequences of squares of some row; we choose the bijection that *reverses* the left-to-right ordering of their squares. For the companion tableaux $T, \bar{T}$ in (5) we get the following bijection, indicated by matching labels:

$$p\colon \begin{array}{ccccc} & & & b & a \\ & f & e & d & c \\ l & k & h & g & \\ q & p & n & m & \\ s & r & & & \end{array} \longrightarrow \begin{array}{ccccccccc} & & & & & & b & f & l \\ & & & & a & d & e & q & \\ & & & & h & k & & & \\ & & c & g & s & & & & \\ & n & p & & & & & & \\ m & r & & & & & & & \end{array} \tag{8}$$

We have the following property: for any square $(i, j) \in Y(\lambda/\mu)$, all squares whose images are on the same row as $p(i, j)$ and to the *left* of it, themselves lie in a column to the *right* of (i, j), and in the same row or above it; in formula: if $p(i, j) = (r, c)$ and $p(i', j') = (r, c')$ with $c' < c$, then $j' > j$ and $i' \le i$. For instance, in the second part of (8) the labels d and a appear in the same row as e and to its left; therefore these labels appear in the first part of (8) strictly to the right of e, and weakly above it. Now given T and κ, one can construct p—and implicitly $\bar{T}$—by traversing $Y(\lambda/\mu)$ in such an order that the mentioned other squares (i', j') are always encountered before (i, j) is (in the example the alphabetic order of the labels has this property): the image $p(i, j)$ can then be taken to be the first currently unused square in row number $T_{i,j}$ of $Y(\nu/\kappa)$. If moreover the traversal is by rows (like the alphabetic order in the example), then strict increase in the column of $\bar{T}$ at the square $p(i, j)$ can be checked as soon as $p(i, j)$ is located: if the square directly above $p(i, j)$ lies in $Y(\nu/\kappa)$, then it must have been included in the image of p before $p(i, j)$ is. What this amounts to, is that at each point during the construction, the union of $Y(\kappa)$ and the image of p so far constructed must be a Young diagram (this can be checked in (8), adding labels in alphabetic order).

1.4.3. Proposition. *A tableau $T \in \mathrm{SST}(\lambda/\mu, n)$ is κ-dominant if and only if the following test succeeds. A variable $\alpha \in \mathbf{N}^n$ is initialised to κ; then the squares $(i, j) \in Y(\lambda/\mu)$ are traversed by weakly increasing i, and for fixed i by strictly decreasing j: at square (i, j) the component $\alpha_{T_{i,j}}$ is increased by 1. The test succeeds if and only if one has $\alpha \in \mathcal{P}$ throughout the entire procedure.* □

For $\kappa = (0)$ this is still not quite the traditional description of Littlewood-Richardson tableaux, which states that the word over the alphabet $[n]$ obtained by listing the entries $T_{i,j}$ in the order described in the proposition (e.g., for the Littlewood-Richardson tableau L shown in (7), the word 0 0 1 1 0 2 2 1 0 3 1 0 4 2) should be a "lattice permutation" (see §3 for a definition; one also finds the terms "lattice word" and "Yamanouchi word"). However, if one expands the definition of that term, one finds that testing whether the indicated word read off from T is a lattice permutation amounts to performing the test of the proposition with $\kappa = (0)$. Incidentally, the (very old) term lattice permutation appears to be related to the fact that the sequence of values assumed by α in the proposition describes a path (from κ to $\kappa + \mathrm{wt}\, T$) in the lattice $\mathbf{N}^n$, that is confined to remain inside the cone $\bigcup_{d \in \mathbf{N}} \mathcal{P}_{d,n} \subseteq \mathbf{N}^n$. It may also be noted that the *original* formulation in [LiRi] is very close to our definition 1.4.2, see §4. We can now state the Rule.

1.4.4. Theorem [Littlewood-Richardson rule]. *For all $\chi \in \mathcal{S}$ and $\nu \in \mathcal{P}$ one has $c^{\nu}_{\chi} = \#\,\mathrm{LR}(\chi, \nu)$.*

We shall present a proof of this theorem in the course of §§2, 3. Given the definitions (4) of c_χ and (1) and (3) of (skew) Schur polynomials, it suffices to construct a bijection

$$\mathcal{R}\colon \mathrm{SST}(\chi, n) \to \bigcup_{\nu \in \mathcal{P}_{|\chi|,n}} \mathrm{LR}(\chi, \nu) \times \mathrm{SST}(\nu, n) \tag{9}$$

for any n, such that whenever $\mathcal{R}(T) = (L, P)$, one has $\mathrm{wt}\, T = \mathrm{wt}\, P$. Once $\mathcal{R}$ is defined, the fact that the same set $\mathrm{LR}(\chi, \nu)$ occurs, independently of n (provided only that ν occurs in the summation) will prove that the coefficient c^{ν}_{χ} is independent of n, as we claimed. We shall refer to $\mathcal{R}$ as Robinson's correspondence, since it was first described in [Rob] (albeit in different terms, not using tableaux). It is natural to describe the correspondence by separately defining its components $\mathcal{R}_0$ and $\mathcal{R}_1$, where $\mathcal{R}(T) = (\mathcal{R}_0(T), \mathcal{R}_1(T))$.

The decomposition of ordinary Schur functions is trivial, so according to the Littlewood-Richardson rule, $\mathrm{LR}(\lambda/(0), \nu)$ should be empty unless $\lambda = \nu$, in which case it should be a singleton. Indeed, for companion tableaux $T \in \mathrm{SST}(\lambda/(0))$ and $\bar{T} \in \mathrm{SST}(\nu/(0))$ one has $\mathrm{wt}\, T = \nu$ and $\mathrm{wt}\, \bar{T} = \lambda$, which together with $\mathrm{wt}\, T \leq \lambda$ and $\mathrm{wt}\, \bar{T} \leq \nu$ imply $\lambda = \nu$ and $T = \bar{T} = \mathbf{1}_\lambda$. Now let $L \in \mathrm{LR}(\chi, \nu)$ be a Littlewood-Richardson tableau whose shape χ represents $\lambda * \mu$. Then the test of proposition 1.4.3 for (0)-dominance succeeds; if one interrupts the test after traversing the squares of the factor μ of $\lambda * \mu$, one sees that the restriction L_μ of L to that factor is (0)-dominant, and hence a Littlewood-Richardson tableau. Being of partition shape, L_μ must then be equal (up to translation) to $\mathbf{1}_\mu$; in particular, the value of α at the point of interruption is μ. The remainder of the test of (0)-dominance of L then shows that its restriction L_λ to the factor λ of $\lambda * \mu$ is ν/μ-dominant. By the symmetry of companion tableaux, this proves:

1.4.5. Proposition. $\#\,\mathrm{LR}(\chi, \nu) = \#\,\mathrm{LR}(\nu/\mu, \lambda)$ *for $\chi \in \mathcal{S}$ representing $\lambda * \mu$.* □

As an example of the correspondence underlying this proposition, we derive from (7) the following pair of corresponding Littlewood-Richardson tableaux, for $\lambda = (5, 4, 3, 1, 1)$, $\mu = (4, 2, 1)$, $\nu = (6, 5, 5, 3, 2)$:

$$\begin{array}{ccccccccc}
 & & & & & 0 & 0 & 0 & 0 \\
 & & & & & 1 & 1 & & \\
 & & & & & 2 & & & \\
0 & 0 & 1 & 2 & 3 & & & & \\
1 & 1 & 2 & 3 & & & & & \\
2 & 2 & 4 & & & & & & \\
3 & & & & & & & & \\
4 & & & & & & & &
\end{array}
\quad \longleftrightarrow \quad
\begin{array}{cccccc}
 & & & & 0 & 0 \\
 & & 0 & 1 & 1 & \\
 & 0 & 1 & 2 & 2 & \\
0 & 1 & 3 & & & \\
2 & 4 & & & &
\end{array}$$

In order to use the Littlewood-Richardson rule to decompose a skew Schur polynomial s_χ on the basis of Schur polynomials, an effective enumeration is required of the union of sets $\mathrm{LR}(\chi, \nu)$, as ν varies over $\mathcal{P}_{|\chi|}$. This can be done using an efficient search procedure, in which the choices for the entries of the squares of $Y(\chi)$ are fixed in the same order as the traversal of proposition 1.4.3 (alternatively, any "valid reading order" as defined below works equally well). For any square (i, j), the possible values $T_{i,j}$ that can be chosen must make the test of that proposition succeed (i.e., they must index a part of the partition α constructed so far that can be increased), and they must satisfy the monotonicity conditions for the rows and columns of T. The reason that we called the search efficient, is that there is always at least one value that satisfies all these conditions (cf. [vLee2, proposition 2.5.3]), so that the search tree will not have unproductive branches. Here we have assumed that no upper limit n is imposed on the number of non-zero parts of ν and hence on the entries in the Littlewood-Richardson tableau T; such a restriction can however be incorporated into the search, by placing an additional condition on the values $T_{i,j}$ tried: they should be sufficiently small to allow column j of T to be completed in a strictly increasing manner. With this extra requirement there will still always remain at least one possible value, unless $Y(\chi)$ has columns of length exceeding n (in which case of course no suitable tableaux T exist at all).

By contrast, no efficient search procedure is known for enumerating just a single set $\mathrm{LR}(\chi,\nu)$, i.e., one for which the size of the search tree that has to be traversed is proportional to $\#\,\mathrm{LR}(\chi,\nu)$. In fact no sufficient condition for $\mathrm{LR}(\chi,\nu)$ to be non-empty, that does not amount to actually finding an element, is even known; therefore it is not possible in general to tell beforehand whether some value tried for an entry will lead to any solutions. So it is worth observing that, while all bijections described in this paper are easily computable, it is in some cases much harder to enumerate the sets themselves linked by these bijections. We add one more remark, prompted by the fact that many texts give the Littlewood-Richardson rule only in the form $c^{\nu}_{\lambda,\mu} = \#\,\mathrm{LR}(\nu/\mu,\lambda)$, while defining Littlewood-Richardson tableaux in terms of a reading of the tableau. This would falsely suggest that the rule is not practical for calculating a product $s_\lambda s_\mu$, since it would seem to require either the construction of complete trial tableaux of varying shapes before testing the Littlewood-Richardson condition (which would fail in most cases), or separate searches for any plausible shape ν/μ and fixed weight λ (which is also unattractive for reasons just mentioned). In reality, by viewing the search strategy outlined above for Littlewood-Richardson tableaux of a shape representing $\lambda * \mu$ from the perspective of the companion tableaux, one easily finds an efficient enumeration procedure for $\bigcup_\nu \mathrm{LR}(\nu/\mu,\lambda)$. Actually, such an enumeration procedure is just what the rule, in its original form given by Littlewood and Richardson, describes; a literal quotation of this description can be found in §4 below.

1.5. *Pictures and reading orders.*

We have constructed a bijection $p\colon Y(\lambda/\mu) \to Y(\nu/\kappa)$ corresponding to a pair of companion tableaux $T \in \mathrm{SST}(\lambda/\mu)$ and $\bar T \in \mathrm{SST}(\nu/\kappa)$, in order to derive proposition 1.4.3. Since T and $\bar T$ can easily be reconstructed from p, one may study the bijections that arise in this way, in place of such pairs of companion tableaux, or of ν/κ-dominant tableaux of shape λ/μ. The conditions that T and $\bar T$ be semistandard tableaux translate into a geometric characterisation of these bijections; this leads to the concept of *pictures* introduced in [Zel1]. Due to the symmetry between companion tableaux, the inverse of any picture is again a picture. One important aspect of pictures is that there are many equivalent ways to define them (like the different characterisations of companion tableaux above). We shall not discuss pictures in depth here, for which we refer to [FoGr] and [vLee2], but it is useful to make a few observations that result from study of pictures.

We have observed above that if one traverses a row from left to right, then the row index of the (inverse) image by p increases weakly, while the column index decreases strictly. A similar condition holds for traversal of a column from top to bottom: the row index of the (inverse) image increases strictly, while the column index decreases weakly (this can be proved by induction on the column considered). Therefore, when verifying κ-dominance with a single traversal of T (as in proposition 1.4.3), it is not necessary to have encountered all squares in rows above it before handling a square (i,j): only those in column j or to its right can influence the test made at (i,j). We might for instance also traverse columns from top to bottom, processing the columns from right to left. We shall consider any ordering of the squares that encounters (i',j') before (i,j) whenever $i' \leq i$ and $j' \geq j$ to be a *valid reading order*. The two cases where one proceeds systematically by rows or by columns merit special names: we shall refer to the former as the *Semitic reading order* (after the Arabic and Hebrew way of writing), and to the latter as the *Kanji reading order* (after the Japanese word for Chinese characters, which are thusly read). In (8), the Semitic reading of the left diagram gives $abcdefghklmnpqrs$, the Kanji reading order gives $acbdgmehnfkprlqs$, while $acbdegfhkmlnprqs$ corresponds to yet another a valid reading order.

In constructing the picture p corresponding to a ν/κ-dominant tableau T, a distinction is forced between the values at all squares of $Y(\lambda/\mu)$, i.e., p is injective even if T is not. Such a distinction can be used in order to apply to T operations that are initially defined only for tableaux with distinct entries (§2 provides an example): it suffices to use an injective map $r : Y(\nu/\kappa) \to \mathbf{N}$ such that $r \circ p$ is a tableau. Taking for r the map corresponding to the Semitic reading of $Y(\nu/\kappa)$, one obtains a tableau $S = r \circ p$ such that $T_{i,j} < T_{i',j'}$ implies $S_{i,j} < S_{i',j'}$, and when $T_{i,j} = T_{i',j'}$ one has $S_{i,j} < S_{i',j'}$ if and only if $j < j'$; this is essentially the operation of standardisation defined in §2. We note however that $r \circ p$ will be a tableau whenever r corresponds to *any* valid reading order; such tableaux may be called specialisations of p. It is shown in [vLee2] that all relevant operations that are defined for T in terms of its standardisation S

could equally well be defined in terms of any specialisation of p. When studying semistandard tableaux however, it is simplest to use the standardisation (as we shall do), mainly because it does not depend on any choice of a shape ν/κ for which T is ν/κ-dominant.

One can develop the theory of the Littlewood-Richardson rule entirely in terms of pictures; doing so clarifies the structure behind many operations and makes certain symmetries explicit. Nevertheless we believe the exposition in terms of tableaux that we shall give is easier to understand, which is in part due to the fact that pictures are, in spite of their name, more difficult to visualise than tableaux are.

§2. Tableau switching and jeu de taquin.

In this section we shall consider a combinatorial procedure that will turn out to be intimately related to the Littlewood-Richardson rule. This procedure is essentially Schützenbergers *jeu de taquin*, but we prefer to introduce it in a slightly different form called "tableau switching" (a term that was introduced in [BeSoSt] for an operation that, although defined in a somewhat different way, constructs the same correspondence between pairs of skew tableaux as we shall do below).

2.1. *Chains in the Young lattice and standardisation.*

Inclusion of Young diagrams defines a partial ordering on the set $\mathcal{P}$ of partitions, which shall be denoted by '$\subseteq$'; the poset $(\mathcal{P}, \subseteq)$ is called the Young lattice. We define a *skew standard tableau* of shape λ/μ to be a saturated increasing chain in the Young lattice from μ to λ, i.e., a sequence of partitions starting with μ and ending with λ such that the Young diagram of each partition is obtained from that of its predecessor by the addition of exactly one square. We shall denote the set of all skew standard tableaux of shape λ/μ by $\mathrm{ST}(\lambda/\mu)$, and say that each of its elements has size $|\lambda/\mu|$. A skew standard tableau of partition shape is called a standard Young tableau, and we write $\mathrm{ST}(\lambda)$ for $\mathrm{ST}(\lambda/(0))$. If its shape λ/μ is given, then specifying some $S \in \mathrm{ST}(\lambda/\mu)$ amounts to putting a total ordering on $Y(\lambda/\mu)$, describing the order in which the squares are added. This can be done by labelling the squares in the desired order with increasing numbers (or elements of some other totally ordered set); these labels will increase along each row and column, whence the name tableau. For instance, the sequence with Young diagrams 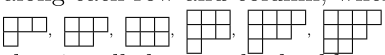 is represented by 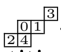. In the literature it is usually this representation that is called a standard tableau, but chains of partitions will be more convenient for us to work with.

In fact any skew semistandard tableau T determines a chain of partitions, with the convention that the squares (i, j) are ordered by increasing value of their entries $T_{i,j}$, or in case these entries are equal, by increasing column number j; this chain (a skew standard tableau) will be called the standardisation of T. For instance, the standardisation of 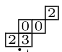 is the skew standard tableau depicted above. Any skew semistandard tableau is determined by its standardisation together with its weight, but a given combination of a skew standard tableau S and a weight α does not necessarily correspond to any skew semistandard tableau T. The condition for the existence of T is that any sequence of squares successively added in S that according to α should have the same entry k in T (so the sequence has length $\alpha_k \geq 2$) should lie in columns whose numbers strictly increase; another formulation is that the indices of the diagonals of the squares should increase (in this case increase is automatically strict, since no two squares on the same diagonal can be successively added). When these equivalent conditions are met, we shall say that S is *compatible* with α.

2.2. *Tableau switching and jeu de taquin.*

Among the skew shapes λ/μ of with $|\lambda/\mu| = 2$, two different kinds can be distinguished: shapes for which the two squares of $Y(\lambda/\mu)$ are incomparable in the natural ordering of $\mathbf{N}^2$, and which shapes therefore admit two different skew standard tableaux, and shapes for which those squares lie in the same row or column (and are adjacent), which shapes admit only one skew standard tableau. The latter kind of shapes will be called dominos. We shall give a construction, based on a certain class of doubly indexed families of partitions, that can be found in [vLee1, 2.1.2]. Let $I = \{\, i \in \mathbf{Z} \mid k \leq i \leq l \,\}$ and $J = \{\, j \in \mathbf{Z} \mid m \leq j \leq n \,\}$ be intervals in $\mathbf{Z}$, and let $(\lambda^{[i,j]})_{i \in I, j \in J}$ be a family of partitions; we shall call this a tableau switching family on $I \times J$ if each "row" $\lambda^{[i,m]}, \ldots, \lambda^{[i,n]}$ and each "column" $\lambda^{[k,j]}, \ldots, \lambda^{[l,j]}$ is a skew standard tableau, and if $\lambda^{[i,j+1]} \neq \lambda^{[i+1,j]}$ whenever $\lambda^{[i+1,j+1]}/\lambda^{[i,j]}$ is not a domino. Here is a small example:

$$\left(\lambda^{[i,j]}\right)_{0\leq i,j<4} = \qquad (10)$$

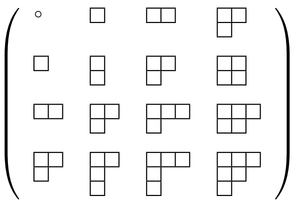

Whenever either the sequence $\lambda^{[i,j]}$, $\lambda^{[i,j+1]}$, $\lambda^{[i+1,j+1]}$ or the sequence $\lambda^{[i,j]}$, $\lambda^{[i+1,j]}$, $\lambda^{[i+1,j+1]}$ is specified to be some skew standard tableau of size 2, there is a unique value for the remaining partition for which one obtains a tableau switching family on $\{i, i+1\} \times \{j, j+1\}$ (because if $\lambda^{[i+1,j+1]}/\lambda^{[i,j]}$ is a domino, then necessarily $\lambda^{[i,j+1]} = \lambda^{[i+1,j]}$). It follows that if values $\lambda^{[i,j]}$ are prescribed for indices $[i,j]$ traversing some "lattice path" going from $[k,m]$ to $[l,n]$ (a zig-zag path in which at each step either i or j increases by 1) by any skew standard tableau, then there is a unique way to extend these values to tableau switching family on $I \times J$. Applying this to the path passing through $[l,m]$ enables the following definition.

2.2.1. Definition. *Let S and T be skew standard tableaux of respective shapes μ/ν and λ/μ. The pair (T', S') of skew standard tableaux obtained from (S,T) by tableau switching, written $(T', S') = X(S,T)$, is defined by the existence of a tableau switching family $(\lambda^{[i,j]})_{k\leq i\leq l;\ m\leq j\leq n}$ such that*

$$\begin{aligned} S &= (\lambda^{[k,m]}, \ldots, \lambda^{[l,m]}), & T &= (\lambda^{[l,m]}, \ldots, \lambda^{[l,n]}); \\ T' &= (\lambda^{[k,m]}, \ldots, \lambda^{[k,n]}), & S' &= (\lambda^{[k,n]}, \ldots, \lambda^{[l,n]}). \end{aligned}$$

As an example, the tableau switching family (10) establishes the fact that $X\Big(\begin{smallmatrix} 0 & 1 \\ 2 & \end{smallmatrix}, \begin{smallmatrix} & & 1 \\ & 2 & \\ 0 & & \end{smallmatrix}\Big) = \Big(\begin{smallmatrix} 0 & 1 \\ 2 & \end{smallmatrix}, \begin{smallmatrix} & & 1 \\ & 0 & \\ 2 & & \end{smallmatrix}\Big)$. A visual way to interpret this definition is the following. Represent S and T together in the skew diagram $Y(\lambda/\nu)$, by filling the squares of their diagrams with labels coming from two disjoint totally ordered sets $A = \{a_k < \cdots < a_{l-1}\}, B = \{b_m < \cdots < b_{n-1}\}$; for instance one can take red numbers a_i for S and blue numbers b_j for T. We shall associate any vertical segment from $[i,j]$ to $[i+1,j]$ in a lattice path with the label a_i, and any horizontal segment from $[i,j]$ to $[i,j+1]$ with the label b_j. Then each lattice path from $[k,m]$ to $[l,n]$ gives rise to a shuffle of the sets A and B, i.e., a total ordering on $A \cup B$ that is compatible with the orderings of A and B individually. Along each lattice path a skew standard tableau of shape λ/ν can be read off from the tableau switching family, which can be represented by filling the squares of $Y(\lambda/\nu)$ with the elements of $A \cup B$, using the total ordering of that set associated to the lattice path.

The initial data of $Y(\lambda/\nu)$ filled with elements of A according to S, and elements of B according to T, defines the part of the tableau switching family along the lattice path that traverses the left and bottom edges. To determine the other members of that family, we shall gradually transform the path, determining a single new member $\lambda^{[i,j]}$ at the time, while updating the filling of $Y(\lambda/\nu)$ with elements of $A \cup B$ so as to correspond to the skew standard tableau read off along the current path, as described above. Eventually the whole family will be determined, and we shall have obtained the lattice path that traverses the top and right edges, from which T' and S' can be read off. The change to the lattice path at each step amounts to transposing one a_i with one b_j in the shuffle, and only a minimal modification, if any, is needed to update the filling of $Y(\lambda/\nu)$: if the squares filled with a_i and b_j are adjacent, then these entries are interchanged, and otherwise the filling is unchanged (this easily follows from definition 2.2.1). Any interchange thus made is a special case of a switch in the sense of [BeSoSt], and this shows that $X(S,T)$ coincides with the result of the tableau switching procedure defined there (in fact that procedure is more liberal, and allows using intermediate fillings that do not correspond to any lattice path).

Here is an example of a transformation of the fillings using the family of (10), and a not very systematic choice of intermediate lattice paths. The paths used are identified by the shuffles written above the tableaux; in many steps this is the only thing that changes. We have used numbers in italics for the a_i, and numbers in bold face for the b_j.

$$
\begin{array}{cccccccccc}
\mathit{012}\mathbf{012} & \mathit{01}\mathbf{0}\mathit{2}\mathbf{12} & \mathit{0}\mathbf{0}\mathit{12}\mathbf{12} & \mathit{0}\mathbf{0}\mathit{1}\mathbf{1}\mathit{2}\mathbf{2} & \mathit{0}\mathbf{01}\mathit{12}\mathbf{2} & \mathit{0}\mathbf{01}\mathit{1}\mathbf{2}\mathit{2} & \mathbf{0}\mathit{0}\mathbf{1}\mathit{1}\mathbf{2}\mathit{2} & \mathbf{0}\mathit{0}\mathbf{12}\mathit{12} & \mathbf{01}\mathit{0}\mathbf{2}\mathit{12} & \mathbf{012}\mathit{012} \\
\begin{array}{|c|c|c|}\hline \mathit{0}&\mathit{1}&\mathbf{1}\\\hline \mathit{2}&\mathbf{2}\\\cline{1-2}\mathbf{0}\\\cline{1-1}\end{array} &
\begin{array}{|c|c|c|}\hline \mathit{0}&\mathit{1}&\mathbf{1}\\\hline \mathbf{0}&\mathbf{2}\\\cline{1-2}\mathit{2}\\\cline{1-1}\end{array} &
\begin{array}{|c|c|c|}\hline \mathit{0}&\mathit{1}&\mathbf{1}\\\hline \mathbf{0}&\mathbf{2}\\\cline{1-2}\mathit{2}\\\cline{1-1}\end{array} &
\begin{array}{|c|c|c|}\hline \mathit{0}&\mathit{1}&\mathbf{1}\\\hline \mathbf{0}&\mathbf{2}\\\cline{1-2}\mathit{2}\\\cline{1-1}\end{array} &
\begin{array}{|c|c|c|}\hline \mathit{0}&\mathbf{1}&\mathit{1}\\\hline \mathbf{0}&\mathbf{2}\\\cline{1-2}\mathit{2}\\\cline{1-1}\end{array} &
\begin{array}{|c|c|c|}\hline \mathit{0}&\mathbf{1}&\mathit{1}\\\hline \mathbf{0}&\mathbf{2}\\\cline{1-2}\mathit{2}\\\cline{1-1}\end{array} &
\begin{array}{|c|c|c|}\hline \mathbf{0}&\mathbf{1}&\mathit{1}\\\hline \mathit{0}&\mathbf{2}\\\cline{1-2}\mathit{2}\\\cline{1-1}\end{array} &
\begin{array}{|c|c|c|}\hline \mathbf{0}&\mathbf{1}&\mathit{1}\\\hline \mathit{0}&\mathbf{2}\\\cline{1-2}\mathit{2}\\\cline{1-1}\end{array} &
\begin{array}{|c|c|c|}\hline \mathbf{0}&\mathbf{1}&\mathit{1}\\\hline \mathit{0}&\mathbf{2}\\\cline{1-2}\mathit{2}\\\cline{1-1}\end{array} &
\begin{array}{|c|c|c|}\hline \mathbf{0}&\mathbf{1}&\mathit{1}\\\hline \mathbf{2}&\mathit{0}\\\cline{1-2}\mathit{2}\\\cline{1-1}\end{array}
\end{array}
$$

In retrospect, some traces of the tableau switching procedure can be already be found in [Haim1] and in the *tables de promotion* of [Schü2]. The symmetry of our definition with respect to i and j has an obvious but important consequence.

2.2.2. Theorem. *Tableau switching is involutive: if $X(S,T) = (T',S')$, then $X(T',S') = (S,T)$.* □

One particular way to transform the lattice path via $[l,m]$ into the one via $[k,n]$, i.e., to go from the shuffle in which all a_i precede all b_j to the one in which all b_j precede all a_i, is to start transposing the final element a_{l-1} of A with all elements of B, then to do the same with a_{l-2}, etc. If, while processing each a_i, we temporarily view the square labelled by it as "empty", then each interchange of entries corresponds to sliding some b_j either up or to the left, into the empty square; the description we so obtain for the transformation of the tableau labelled by B before and after processing a_i is exactly that of in inward *jeu de taquin* slide into the square initially containing a_i (see [Schü3], or for instance [Fult]). It follows that if $X(S,T) = (T',S')$, then T' is obtained from T by a sequence of inward jeu de taquin slides, which we shall write as $T \triangleright T'$. In the example displayed above we obtain the transformation $T \triangleright T'$ (the entries in bold face) as follows:

$$
\begin{array}{|c|c|c|}\hline \ & \ & 1\\\hline \ & 2\\\cline{1-2} 0\\\cline{1-1}\end{array}
\;\triangleright\;
\begin{array}{|c|c|c|}\hline \ & \ & 1\\\hline 0 & 2\\\cline{1-2} \ \\\cline{1-1}\end{array}
\;\triangleright\;
\begin{array}{|c|c|c|}\hline \ & 1 & \ \\\hline 0 & 2\\\cline{1-2} \ \\\cline{1-1}\end{array}
\;\triangleright\;
\begin{array}{|c|c|c|}\hline 0 & 1 & \ \\\hline 2 & \ \\\cline{1-2} \ \\\cline{1-1}\end{array}\,.
$$

By theorem 2.2.2, one also has $S \triangleleft S'$, i.e., S' is obtained from S by a sequence of outward jeu de taquin slides. For the italic entries above:

$$
\begin{array}{|c|c|c|}\hline 0 & 1 & \ \\\hline 2 & \ \\\cline{1-2} \ \\\cline{1-1}\end{array}
\;\triangleleft\;
\begin{array}{|c|c|c|}\hline \ & 1 & \ \\\hline 0 & \ \\\cline{1-2} 2\\\cline{1-1}\end{array}
\;\triangleleft\;
\begin{array}{|c|c|c|}\hline \ & \ & 1\\\hline 0 & \ \\\cline{1-2} 2\\\cline{1-1}\end{array}
\;\triangleleft\;
\begin{array}{|c|c|c|}\hline \ & \ & 1\\\hline \ & 0\\\cline{1-2} 2\\\cline{1-1}\end{array}\,.
$$

The following simple observation will allow us to define tableau switching for skew semistandard tableaux as well as for skew standard tableaux.

2.2.3. Proposition. *If T,T' are skew standard tableaux with $T \triangleright T'$, then T' is compatible with a weight α (as defined in subsection 2.1) if and only if T is.*

Proof. Compatibility of T and T' with α is checked by comparing positions of certain pairs of squares added at successive steps in the chain of partitions; therefore the general case reduces to the one where $X(S,T) = (T',S')$ with T of size 2 and S of size 1. The validity in this case is fairly obvious by inspection, but here is a formal argument. We show that the relative order of the diagonals of the two entries of T is unchanged in T'; then compatibility with α does not change either. For each of the two entries of T, the index of its diagonal changes by at most 1 during the slide, and the two entries certainly do not exchange places. Then since the only other square involved, that of S, does not lie on the same diagonal as either of the entries of T, the relative order of their diagonals remains unchanged by the slide. □

Remark. In fact even more is true: if $T \triangleright T'$ for $T \in \mathrm{ST}(\lambda/\mu)$ and $T' \in \mathrm{ST}(\lambda'/\mu')$, then for any skew shape ν/κ there exists a ν/κ-dominant tableau in $\mathrm{SST}(\lambda/\mu)$ with standardisation T if and only if there is one in $\mathrm{SST}(\lambda'/\mu')$ with standardisation T' (both have weight $\nu - \kappa$). This is a bit more difficult to prove, but the proof is still straightforward: one has to prove that after modifying the relevant companion tableau to reflect the slide $T \triangleright T'$, it still satisfies the tableau condition. This is essentially what is shown in [vLee2, theorem 5.1.1]; see also subsection 3.4 below.

2.2.4. Definition. *Let T_1, T_2 be skew semistandard tableaux such that tableau switching can be applied to the pair (S_1, S_2) of their standardisations; then tableau switching can also be applied to (T_1, T_2), resulting in a pair of skew semistandard tableaux $(T_2', T_1') = X(T_1, T_2)$ determined by the conditions that their standardisations are $(S_2', S_1') = X(S_1, S_2)$, and the weight of T_i' is equal to that of T_i $(i = 1, 2)$. The notion of jeu de taquin slides is also extended to this case; we write $T_2 \triangleright T_2'$ and $T_1 \triangleleft T_1'$.*

We note one further property of tableau switching that is immediate from its definition. From $T \in \mathrm{ST}(\mu/\nu)$ and $U \in \mathrm{ST}(\lambda/\mu)$, a skew standard tableau of shape λ/ν can be formed by joining together the chains of partitions; we shall denote it by $T|U$. As a filling of the diagram $Y(\lambda/\nu)$, it is the union of the fillings for T and U, after making sure (by adding some offset) that all entries used for U exceed those for T. Then by similarly joining tableau switching families we get:

2.2.5. Proposition. *If $X(S,T) = (T', S')$ and $X(S', U) = (U', S'')$, then $X(S, T|U) = (T'|U', S'')$.* $\square$

2.3. *Jeu de taquin equivalence and dual equivalence.*

Jeu de taquin defines an equivalence relation on skew standard tableaux, and on skew semistandard tableaux, generated by the relations $T \triangleright T'$; this relation is called jeu de taquin equivalence. Tableau switching allows us to define another equivalence relation, called dual equivalence (see [Haim2]).

2.3.1. Definition. *Two skew (semi)standard tableaux S_1, S_2 of equal shape are called dual equivalent if for any tableau T of appropriate shape, and with $X(S_i, T) = (T'_i, S'_i)$ for $i = 1, 2$, one has $T'_1 = T'_2$.*

It follows from proposition 2.2.5 that the tableaux S'_1 and S'_2 in this definition are also dual equivalent. The meaning of dual equivalence can be expressed in two ways in terms of jeu de taquin. Firstly, if S_1 and S_2 are dual equivalent, then they have the same shape, and this remains true whenever the same sequence of outward jeu de taquin slides is applied to each of them (i.e., each slide starts with the same empty square in both cases); this is because the tableaux T'_1 and T'_2 record the squares in which the successive slides end. Secondly, when S_1 and S_2 are used to determine sequences of inward slides, then these sequences will always have the same effect when applied to any tableau T.

It is somewhat surprising that there exist pairs of distinct skew standard tableaux that are dual equivalent, since the definition refers to arbitrarily large tableaux T, and therefore for instance to arbitrarily long sequences of jeu de taquin slides. To describe the most elementary cases of such pairs, we need some notation for skew standard tableaux of size 3. As mentioned above, a skew standard tableau $T \in \mathrm{ST}(\lambda/\mu)$ can be described by an ordering of the squares of $Y(\lambda/\mu)$. If $|\lambda/\mu| = 3$ those squares lie on distinct diagonals, which are ordered by index, so we may specify T using a permutation of the letters a, b, c, which we shall call the type of T. The convention used is that the three positions in the word correspond to the diagonals, while a, b, and c, respectively indicate the first second and third square added. As an example, the tableau displayed as $\begin{smallmatrix} & \boxed{1}\boxed{2} \\ \boxed{3} & \end{smallmatrix}$ is of type cab.

2.3.2. Proposition. *Let λ/μ be a shape with $|\lambda/\mu| = 3$. Then $\mathrm{ST}(\lambda/\mu)$ contains a tableau of type bca if and only if it contains a tableau of type acb, and if so, the two tableaux are dual equivalent. Similarly, $\mathrm{ST}(\lambda/\mu)$ contains a tableau of type bac if and only if it contains one of type cab, and if so, the two are dual equivalent.*

Proof. The shape may restrict the possible types of tableaux, due to adjacency of squares: the leftmost or upper one of two adjacent squares must be added before the other one. But within either of the sets of types $\{bca, acb\}$ or $\{bac, cab\}$, the relative order among squares of $Y(\lambda/\mu)$ on successive diagonals is fixed (only the ordering between the outer diagonals differs), so possibility of one type implies that of the other. For the dual equivalence, we shall show that any pair of tableaux matching the hypotheses of the proposition is transformed by a single jeu de taquin slide into another such pair, having in particular equal shapes; then dual equivalence follows by induction on the size of the tableau T in definition 2.3.1 (using proposition 2.2.5). So let $S_1, S_2 \in \mathrm{ST}(\lambda/\mu)$ either have types $\{bca, acb\}$ or types $\{bac, cab\}$, and let T be of size 1 occupying a square t, with $X(S_i, T) = (T'_i, S'_i)$ defined for $i = 1, 2$.

Suppose first that t lies on a diagonal that does not meet $Y(\lambda/\mu)$; in this case we shall show that S'_i has the same type as S_i ($i = 1, 2$), and that $T'_1 = T'_2$ (which implies that S'_1 and S'_2 have the same shape). The former statement is proved in essentially the same way as proposition 2.2.3: no transpositions are possible in the ordering by diagonal index of the squares of S_i, because diagonal indices change by at most one during the slide, and at most one square is available on each diagonal. To show $T'_1 = T'_2$, we view each computation of $X(S_i, T)$ as three successive applications of an inward slide to T; we claim that the sequence of actual moves of the square of T (i.e., ignoring the slides that leave it in place) is the same for $X(S_1, T)$ as for $X(S_2, T)$. The only difference between the two cases is the relative order of the slides into the squares on the outer diagonals of $Y(\lambda/\mu)$, which squares cannot both be adjacent to t; then at least one of these slides leaves the square of T in place (both for $i = 1$ and $i = 2$). Therefore the relative order of these two slides makes no difference, proving our claim.

If t lies a diagonal that does meet $Y(\lambda/\mu)$, then that must be the middle diagonal, and $Y(\lambda/\mu)$ must have the form $\begin{array}{|c|c|}\hline \ & \ \\ \hline \ \\ \cline{1-1}\end{array}$. This is impossible if the types of S_1, S_2 are $\{bca, acb\}$, and the proof for that case is therefore complete. We are left with the case that can be depicted as $S_1 = \begin{array}{|c|c|}\hline 0 & 2 \\ \hline 1 \\ \cline{1-1}\end{array} \lhd \begin{array}{cc}\cline{2-2} & \multicolumn{1}{|c|}{0} \\ \hline \multicolumn{1}{|c|}{1} & \multicolumn{1}{c|}{2} \\ \hline\end{array} = S_1'$ and $S_2 = \begin{array}{|c|c|}\hline 0 & 1 \\ \hline 2 \\ \cline{1-1}\end{array} \lhd \begin{array}{cc}\cline{2-2} & \multicolumn{1}{|c|}{1} \\ \hline \multicolumn{1}{|c|}{0} & \multicolumn{1}{c|}{2} \\ \hline\end{array} = S_2'$; we see that S_1' and S_2' have the same shape (so $T_1' = T_2'$), and types *bca* respectively *acb*, whence they match the case just completed. □

This basic case implies many others by the following immediate consequence proposition 2.2.5:

2.3.3. Proposition. *If S_1 and S_2 are dual equivalent, and also T_1 and T_2, then so are $S_1|T_1$ and $S_2|T_2$.* □

2.4. *Confluence of jeu de taquin.*

A crucial property of jeu de taquin is its confluence, i.e., the property that whenever $T \rhd T_1$ and $T \rhd T_2$, then there exists a tableau U such that $T_1 \rhd U$ and $T_2 \rhd U$. Since any sequence of slides can be extended until a tableau of partition shape is reached, and no further (Young tableaux are the normal forms for jeu de taquin), this is equivalent to saying that for any skew tableau T there is a *unique* Young tableau P such that $T \rhd P$. Although this property will follow as a corollary to our main theorem in §3, we shall give an independent proof here, that does not require any further constructions; what is more, we have in fact already established the essential part of the argument, in proving proposition 2.3.2. Our proof is inspired by the one given in [Eriks] (although it must be admitted that by translating the reasoning using tableau switching, it has become rather similar to the arguments contained in [Haim2]). The confluence we wish to prove can be formulated as follows: given any $T \in \mathrm{ST}(\lambda/\mu)$, all sequences of inward jeu de taquin slides to fill up all squares of $Y(\mu)$ give the same standard Young tableau as result when applied to T. Those sequences of slides are determined by the standard Young tableaux of shape μ, and the statement means that these should all be dual equivalent (by the second interpretation of that term, given after its definition). Therefore we state:

2.4.1. Theorem. *For any $\lambda \in \mathcal{P}$, all tableaux in $\mathrm{ST}(\lambda)$ are dual equivalent.*

Proof. The proof is by induction on $|\lambda|$; the case $|\lambda| = 0$ is trivial (in fact all cases with $|\lambda| \leq 3$ are either trivial or already established). By the induction hypothesis and proposition 2.3.3, one has for any corner s of $Y(\lambda)$ dual equivalence among all members of the subset $C(\lambda, s)$ of $\mathrm{ST}(\lambda)$ of tableaux whose highest entry occupies the square s (as chains of partitions, these tableaux have in common a predecessor of λ in $(\mathcal{P}, \subseteq)$). It will then suffice to establish dual equivalence between a pair of elements chosen from any two such subsets $C(\lambda, p)$, $C(\lambda, r)$; we assume that the diagonal of the corner p has smaller index than that of r. The Young diagram $Y(\lambda) \setminus \{p, r\}$ contains at least one corner q on a diagonal whose index is between those of p and q: one can take for instance its unique corner in the row above p. Let $\mu \in \mathcal{P}$ be such that $Y(\mu) = Y(\lambda) \setminus \{p, q, r\}$, and let $S_1, S_2 \in \mathrm{ST}(\lambda/\mu)$ be the tableaux of respective types *cab* and *bac*; these tableaux are dual equivalent by proposition 2.3.2. For any $R \in \mathrm{ST}(\mu)$ one has $R|S_1 \in C(\lambda, p)$ and $R|S_2 \in C(\lambda, r)$, and $R|S_1$ and $R|S_2$ are dual equivalent by proposition 2.3.3. □

Remark. The statement of the theorem corresponds to a global form of confluence of jeu de taquin, in the sense that any two sequences of jeu de taquin slides applied to T will eventually converge when the normal form (Young tableau) is reached. The proof however shows that confluence can in fact be obtained locally, namely if two different jeu de taquin slides are applied to T, then the resulting tableaux can be made equal by applying at most two more slides to each of them: the dual equivalence of S_1 and S_2 in the proof means that successive slides into p, r, q have the same effect as those into r, p, q.

We defined two tableaux to be dual equivalent if their equality of shape is preserved under sequences of *outward* slides, or equivalently if the two sequences of *inward* slides (applied to some tableau T)

determined by them always have the same effect. One may ask whether this implies the same property with "inward" and "outward" interchanged; indeed that is part of the requirement for dual equivalence in its original definition in [Haim2]. To prove that this property follows from our definition (in fact it is equivalent to it), we use an involutive poset anti-isomorphism of the sub-poset of the Young lattice of partitions contained in some fixed (large) rectangular partition. Fix a partition ρ with n non-zero parts, all equal to m; then this anti-isomorphism, which we shall denote by $\lambda \mapsto \lambda^\diamond$ for $\lambda \subseteq \rho$, is given by $\lambda_i^\diamond = m - \lambda_{n-1-i}$ for $i \in [n]$. For any $T \in \mathrm{ST}(\lambda/\mu)$ with $\lambda \subseteq \rho$, we also define $T^\diamond \in \mathrm{ST}(\mu^\diamond/\lambda^\diamond)$ by applying the anti-isomorphism to all partitions in the chain of T, and reversing their order. By a similar operation applied to tableau switching families, it can be seen that $X(S,T) = (T', S')$ implies $X(T^\diamond, S^\diamond) = (S'^\diamond, T'^\diamond)$.

2.4.2. Proposition. *Two tableaux S_1, S_2 of the same shape are dual equivalent if and only if for any tableau T of appropriate shape, and with $X(T, S_i) = (S'_i, T'_i)$ for $i = 1, 2$, one has $T'_1 = T'_2$.*

Proof. Let us temporarily call the relation in the second clause of the proposition reverse dual equivalence. Suppose first that S_1 and S_2 are reverse dual equivalent and of shape λ/μ. Let $T \in \mathrm{ST}(\mu/(0))$, and put $X(T, S_i) = (S'_i, T')$ for $i = 1, 2$; then S'_1 and S'_2 are dual equivalent by theorem 2.4.1, so S_1 and S_2 are also dual equivalent by the remark after definition 2.3.1. (In fact theorem 2.4.1 is equivalent to the "if" part of the current proposition, since Young tableaux of equal shape are trivially reverse dual equivalent.) Conversely, suppose that S_1 and S_2 are dual equivalent; then for any choice of a rectangular partition $\rho \supseteq \lambda$, the tableaux $S_1^\diamond$ and $S_2^\diamond$ are reverse dual equivalent, whence (by what we just proved) they are dual equivalent, and so $S_1 = S_1^{\diamond\diamond}$ and $S_2 = S_2^{\diamond\diamond}$ are reverse dual equivalent. □

This proposition provides an effective test of dual equivalence. Given tableaux S_1, S_2 of equal shape, one applies successive inward slides into the same squares to both tableaux; if at any moment their shapes become different, then S_1 and S_2 are not dual equivalent, but if the tableaux become Young tableaux without exhibiting shape difference, then one has established the second clause of the proposition, and therefore dual equivalence of S_1 and S_2. For jeu de taquin equivalence one also has a test, due to confluence: one reduces both tableaux to Young tableaux, which will be equal if and only if the original tableaux are jeu de taquin equivalent.

2.4.3. Corollary. *If S_1, S_2 are both dual equivalent and jeu de taquin equivalent, then $S_1 = S_2$.*

Proof. Reducing S_1 and S_2 in parallel preserves shapes and eventually produces equal tableaux; reversing the slides, one must have had $S_1 = S_2$ to begin with. □

2.5. *An alternative Littlewood-Richardson rule.*

Let $\chi \in \mathcal{S}$; the fact that to each skew tableau $T \in \mathrm{ST}(\chi)$ there is associated a unique Young tableau P obtainable from it by jeu de taquin, allows us to subdivide $\mathrm{ST}(\chi)$ according to the shape of this P. Define $\mathrm{ST}(\chi)^{\triangleright P} = \{\, T \in \mathrm{ST}(\chi) \mid T \triangleright P \,\}$ for any standard Young tableau P, and for $\nu \in \mathcal{P}_{|\chi|}$ set $\mathrm{ST}(\chi)^{\triangleright \nu} = \bigcup_{P \in \mathrm{ST}(\nu)} \mathrm{ST}(\chi)^{\triangleright P}$. The sets $\mathrm{ST}(\chi)^{\triangleright P}$ are the fibres of the map $\mathrm{ST}(\chi)^{\triangleright \nu} \to \mathrm{ST}(\nu)$ sending $T \mapsto P$ when $T \triangleright P$. We show that these fibres all have the same number of elements, and that in fact there are canonical bijections between them, so that $\mathrm{ST}(\chi)^{\triangleright \nu}$ is in natural bijection with the Cartesian product of $\mathrm{ST}(\nu)$ and any one such fibre.

2.5.1. Proposition. *Let $\lambda/\mu \in \mathcal{S}$ and $\nu \in \mathcal{P}$, and let $C \in \mathrm{ST}(\nu)$ be fixed. Then there is a bijection $\phi\colon \mathrm{ST}(\lambda/\mu)^{\triangleright C} \times \mathrm{ST}(\nu) \to \mathrm{ST}(\lambda/\mu)^{\triangleright \nu}$, that can be characterised by the conditions that $\phi(L, P) \triangleright P$ and that $\phi(L, P)$ is dual equivalent to L.*

Proof. We can construct $\phi(L, P)$, which is then uniquely determined due to corollary 2.4.3, as follows. Choose any $Q \in \mathrm{ST}(\mu)$, then $X(Q, L) = (C, S)$ for some $S \in \mathrm{ST}(\lambda/\nu)$; then compute $X(P, S)$, which because of $S \triangleright Q$ will be of the form (Q, T) with $T \in \mathrm{ST}(\chi)^{\triangleright P}$. From $X(Q, L) = (C, S)$ and $X(Q, T) = (P, S)$ with $C, P \in \mathrm{ST}(\nu)$, we get dual equivalence of L and T; we set $\phi(L, P) = T$. Since (P, S, L) can be reconstructed from (T, Q, C), the correspondence ϕ is bijective. □

As an example of this construction, let

$$C = \begin{array}{lllll} 0&1&2&7&11\\ 3&5&9&&\\ 4&6&10&&\\ 8&12&&& \end{array}, \qquad L = \begin{array}{lllllll} &&&&2&7&11\\ &&&5&9&&\\ &0&1&10&&&\\ &3&12&&&&\\ 4&6&&&&&\\ 8&&&&&& \end{array}, \qquad P = \begin{array}{lllll} 0&2&4&5&9\\ 1&3&7&&\\ 6&10&12&&\\ 8&11&&& \end{array};$$

choosing $Q = \begin{array}{llll} 0&1&5&8\\ 2&4&7&\\ 3&&&\\ 6&&& \end{array}$ one gets $S = \begin{array}{lllllll} &&&&&1&8\\ &&&2&5&&\\ &&&7&&&\\ &&4&&&&\\ 0&6&&&&&\\ 3&&&&&& \end{array}$, and $X(P,S) = (Q,T)$, where

$$T = \phi(L,P) = \begin{array}{lllllll} &&&&2&5&9\\ &&&4&7&&\\ &0&3&10&&&\\ &1&12&&&&\\ 6&11&&&&&\\ 8&&&&&& \end{array}.$$

The reader may check that the same tableau T is obtained for other choices of Q.

The above proposition can be generalised to skew semistandard tableaux, using proposition 2.2.3. It suffices to replace $\mathrm{ST}(\nu)$ and $\mathrm{ST}(\lambda/\mu)$ respectively by $\mathrm{SST}(\nu,n)$ and $\mathrm{SST}(\chi,n)$, and $\mathrm{ST}(\lambda/\mu)^{\triangleright\nu}$ by $\mathrm{SST}(\chi,n)^{\triangleright\nu}$, which is defined as $\bigcup_{C\in\mathrm{SST}(\nu,n)} \mathrm{SST}(\chi)^{\triangleright C}$ where $\mathrm{SST}(\chi)^{\triangleright C} = \{\, L \in \mathrm{SST}(\chi,n) \mid L \triangleright C \,\}$. We obtain the following corollary, either by applying proposition 2.5.1, or by reusing its proof almost literally.

2.5.2. Corollary. *Let $\chi \in \mathcal{S}$, $\nu \in \mathcal{P}$, $n \in \mathbf{N}$, and fix $C \in \mathrm{SST}(\nu,n)$. Then there is a bijection $\phi: \mathrm{SST}(\chi)^{\triangleright C} \times \mathrm{SST}(\nu,n) \to \mathrm{SST}(\chi,n)^{\triangleright\nu}$, that can be characterised by the conditions that $\phi(L,P) \triangleright P$ and that $\phi(L,P)$ is dual equivalent to L.* □

If one combines the inverses of such bijections ϕ for all $\nu \in \mathcal{P}_{|\chi|,n}$, one obtains a bijection $\mathcal{R}'$ defined on $\bigcup_{\nu\in\mathcal{P}_{|\chi|,n}} \mathrm{SST}(\chi,n)^{\triangleright\nu} = \mathrm{SST}(\chi,n)$, whose codomain strongly resembles that specified for Robinson's bijection $\mathcal{R}$ at the right hand side of (9); the difference is that, in each component of the union over $\nu \in \mathcal{P}_{|\chi|,n}$, the factor $\mathrm{LR}(\chi,\nu)$ is replaced by $\mathrm{SST}(\chi)^{\triangleright C}$, for some $C \in \mathrm{SST}(\nu,n)$ chosen separately for every ν. Writing $\mathcal{R}'(T) = (\mathcal{R}'_0(T), \mathcal{R}'_1(T))$, the map $\mathcal{R}'_1$ preserves weight (since $T \triangleright \mathcal{R}'_1(T)$); therefore the bijection $\mathcal{R}'$ gives us an alternative expression for the value of Littlewood-Richardson coefficients:

2.5.3. Corollary. *If the skew shape χ represents $\lambda * \mu$, then the Littlewood-Richardson coefficient $c^{\nu}_{\lambda,\mu}$ equals $\#\,\mathrm{SST}(\chi)^{\triangleright C}$ for any $C \in \mathrm{SST}(\nu)$. More generally, $\#\,\mathrm{SST}(\chi)^{\triangleright C} = c^{\nu}_{\chi}$ for any skew shape χ.* □

Although this version of the Littlewood-Richardson rule has the advantage of already being proved, it is of little practical use in its present form, since the sets $\mathrm{SST}(\chi)^{\triangleright C}$ cannot be enumerated in any useful manner. However, we shall in the next section prove the identity

$$\mathrm{SST}(\chi)^{\triangleright \mathbf{1}_\nu} = \mathrm{LR}(\chi,\nu); \tag{11}$$

it follows that for the special choice $C = \mathbf{1}_\nu \in \mathrm{SST}(\nu,n)$ for all ν, the bijection $\mathcal{R}'$ exactly matches the specification of Robinson's bijection, providing a proof of the Littlewood-Richardson rule. Indeed, we shall define $\mathcal{R}$ so that it coincides with this specialisation of $\mathcal{R}'$. The remark we made following definition 2.2.4 implies (11), so we could complete our proof of the Littlewood-Richardson rule by proving that remark. Using coplactic operations, as we shall do, is certainly not the simplest way to prove (11), but it provides a better insight into Robinson's bijection, and in particular it gives a simpler description of $\mathcal{R}_0$ than the one that can be extracted from what has been presented so far. It will also lead to a proof of the Littlewood-Richardson rule that does not depend on any of the non-trivial results derived in this section.

By proposition 1.4.5, (11) also implies $\#\,\mathrm{SST}(\chi)^{\triangleright\mathbf{1}_\nu} = \#\,\mathrm{SST}(\nu/\mu)^{\triangleright\mathbf{1}_\lambda}$ when χ represents $\lambda * \mu$, and hence $c^{\nu}_{\lambda,\mu} = c^{\lambda}_{\nu/\mu}$. On the other hand $c^{\lambda}_{\nu/\mu} = c^{\mu}_{\nu/\lambda}$ can be obtained without using (11): the relation $X(\mathbf{1}_\mu, T) = (\mathbf{1}_\lambda, T^*)$ defines a bijection between tableaux $T \in \mathrm{SST}(\nu/\mu)^{\triangleright\mathbf{1}_\lambda}$ and $T^* \in \mathrm{SST}(\nu/\lambda)^{\triangleright\mathbf{1}_\mu}$ (this can be generalised by replacing $\mathbf{1}_\mu$ and $\mathbf{1}_\lambda$ by other fixed tableaux of the same shape). For instance, here is the computation for the tableau L^* corresponding to the Littlewood-Richardson tableau L of (7).

$$L = \begin{array}{|c|c|c|c|c|c|} \cline{5-6} \multicolumn{4}{c|}{} & 0 & 0 \\ \cline{3-6} \multicolumn{2}{c|}{} & 0 & 1 & 1 & \multicolumn{1}{c}{} \\ \cline{2-5} \multicolumn{1}{c|}{} & 0 & 1 & 2 & 2 & \multicolumn{1}{c}{} \\ \cline{1-5} 0 & 1 & 3 & \multicolumn{3}{c}{} \\ \cline{1-3} 1 & 4 & \multicolumn{4}{c}{} \\ \cline{1-2} \end{array} \quad , \quad \begin{array}{|c|c|c|c|c|c|} \hline \mathit{0} & \mathit{0} & \mathit{0} & \mathit{0} & \mathbf{0} & \mathbf{0} \\ \hline \mathit{1} & \mathit{1} & \mathbf{0} & \mathbf{1} & \mathbf{1} & \multicolumn{1}{c}{} \\ \cline{1-5} \mathit{2} & \mathbf{0} & \mathbf{1} & \mathbf{2} & \mathbf{2} & \multicolumn{1}{c}{} \\ \cline{1-5} \mathbf{0} & \mathbf{1} & \mathbf{3} & \multicolumn{3}{c}{} \\ \cline{1-3} \mathbf{2} & \mathbf{4} & \multicolumn{4}{c}{} \\ \cline{1-2} \end{array} \quad \xleftrightarrow{X} \quad \begin{array}{|c|c|c|c|c|c|} \hline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathit{0} \\ \hline \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathit{0} & \multicolumn{1}{c}{} \\ \cline{1-5} \mathbf{2} & \mathbf{2} & \mathbf{2} & \mathit{0} & \mathit{1} & \multicolumn{1}{c}{} \\ \cline{1-5} \mathbf{3} & \mathit{0} & \mathit{2} & \multicolumn{3}{c}{} \\ \cline{1-3} \mathbf{4} & \mathit{1} & \multicolumn{4}{c}{} \\ \cline{1-2} \end{array} \quad , \quad L^* = \begin{array}{|c|c|c|c|c|} \cline{5-5} \multicolumn{4}{c|}{} & 0 \\ \cline{4-5} \multicolumn{3}{c|}{} & 0 & \multicolumn{1}{c}{} \\ \cline{3-4} \multicolumn{2}{c|}{} & 0 & 1 & \multicolumn{1}{c}{} \\ \cline{1-4} 0 & 2 & \multicolumn{3}{c}{} \\ \cline{1-2} 1 & \multicolumn{4}{c}{} \\ \cline{1-1} \end{array} \tag{12}$$

The identity $c^\lambda_{\nu/\mu} = c^\mu_{\nu/\lambda}$ is of course related to the symmetry $c^\nu_{\lambda,\mu} = c^\nu_{\mu,\lambda}$ that is obvious from the definition (2). If we assume (11), we can derive a bijection corresponding to that symmetry. Write $\mathrm{LR}(\lambda * \mu, \nu)$ for $\mathrm{LR}(\chi, \nu)$ when χ represents $\lambda * \mu$, and let $L \in \mathrm{LR}(\lambda * \mu, \nu)$ be given. First determine $\bar L_\lambda \in \mathrm{LR}(\nu/\mu, \lambda)$ corresponding to L by proposition 1.4.5 (the companion tableau of the subtableau L_λ of L), then compute $\bar L^*_\lambda \in \mathrm{LR}(\nu/\lambda, \mu)$ as above by $X(\mathbf{1}_\mu, \bar L_\lambda) = (\mathbf{1}_\lambda, \bar L^*_\lambda)$, and finally apply the bijection corresponding to proposition 1.4.5 in the opposite direction to $\bar L^*_\lambda$ so as to obtain a tableau in $\mathrm{LR}(\mu * \lambda, \nu)$. It does not appear that this somewhat complicated process can be simplified.

§3. Coplactic operations.

In this section we shall introduce another kind of operations on skew semistandard tableaux, which we shall call *coplactic operations*. Unlike jeu de taquin, these transformations do not change the shape of a tableau, but rather its weight, by changing the value of one of the entries. The basic definitions can be formulated most easily in terms of finite words over the alphabet $[n]$. In the application to tableaux, these words will be the ones obtained by listing the entries of skew semistandard tableaux using a valid reading order as described in §1.5; this means in particular that notions of "left" and "right" within words will (unfortunately) get a more or less opposite interpretation within tableaux.

3.1. *Coplactic operations on words.*

We first fix some terminology pertaining to words. A *word* over a set A (called the alphabet) is a finite (possibly empty) sequence of elements of A, arranged from left to right; the elements of the sequence forming a word w are called the *letters* of w. The set of all words of length l over A is denoted by A^l, and $A^* = \bigcup_{l \in \mathbf{N}} A^l$. The concatenation of two words $u, v \in A^*$ will be denoted by uv (the sequence u of letters, followed by the sequence v); this defines an associative product on A^*. Whenever a word w can be written as uv, the word u is called a *prefix* of w, and v a *suffix* of w; a *subword* of w is any word that can be obtained by removing from w a (possibly empty) prefix and a suffix.

For words $w \in [n]^*$ we define the weight $\mathrm{wt}\, w \in \mathbf{N}^n$ in the same way as for semistandard tableaux, i.e., $(\mathrm{wt}\, w)_i$ counts the number of letters i in w. A word $w \in [n]^*$ will be called *dominant* for $i \in [n-1]$ if every prefix u of w satisfies $(\mathrm{wt}\, u)_i \geq (\mathrm{wt}\, u)_{i+1}$, and it will be called *anti-dominant* for i if every suffix v of w satisfies $(\mathrm{wt}\, v)_i \leq (\mathrm{wt}\, v)_{i+1}$. If w is both dominant and anti-dominant for i, it will be called *neutral* for i. If w is either dominant or anti-dominant for i, then it is neutral for i if and only if $(\mathrm{wt}\, w)_i = (\mathrm{wt}\, w)_{i+1}$. In a word w that is neutral for i there is a matching between letters i and letters $i+1$ to their right, like properly matched left and right parentheses (the words in $\{i, i+1\}^*$ that are neutral for i form a so-called Dyck language). Removing from any word a subword that is neutral for i does not affect whether it is dominant, anti-dominant, or neutral for i. A word that contains no non-empty subwords that are neutral for i is of the form $(i+1)^r\, i^s$ with $r, s \in \mathbf{N}$, where exponentiation signifies repetition of the same letter.

A word in $[n]^*$ that is simultaneously dominant for all $i \in [n-1]$ will be simply called *dominant*, or in older terminology a *lattice permutation*. Equivalently, a word is dominant if the weight of every one of its prefixes is a partition (of the length of the prefix into n parts). Observe the similarity with 0-dominance for semistandard tableaux, as characterised in proposition 1.4.3; this is what motivated our choice of terminology. The sequence of weights of the successive prefixes of a dominant word w forms a standard Young tableau, from which w can be readily reconstructed. For later reference we state this as follows:

3.1.1. Proposition. *The set of dominant words $w \in [n]^d$ of weight $\lambda \in \mathcal{P}_{d,n}$ is in bijection with* $\mathrm{ST}(\lambda)$*, associating to w to the sequence of weights of its prefixes.* $\square$

3.1.2. Definition. *A coplactic operation in $[n]^*$ is a transition between words $w = uiv$ and $w' = u(i+1)v$, where $i \in [n-1]$, $u, v \in [n]^*$, and u is anti-dominant for i while v is dominant for i. In this case we shall write $w = e_i(w')$ and $w' = f_i(w)$.*

For instance, in the following word over the alphabet $[6]$, decrementing by 1 any one of the numbers with an underline, or incrementing by 1 any one of the numbers with an overline, constitutes a coplactic operation, and no other coplactic operations are possible; the word is dominant for 1 and neutral for 2.

$$\underline{4}\ 0\ 1\ 5\ 2\ \overline{\underline{1}}\ 3\ \underline{5}\ 0\ 1\ \overline{\underline{4}}\ 2\ \overline{\underline{0}}\ 0\ 1\ 2\ \overline{\underline{3}}\ 3\ 4 \tag{13}$$

In the definition, the letter i in w that is changed to $i+1$ in w' is not contained in any subword $u_0 i v_0$ of w that is anti-dominant for i, for v_0 would then have strictly more letters $i+1$ than letters i, contradicting the dominance for i of v; similarly the indicated letter $i+1$ in w' is not contained in any subword that is dominant for i. In particular the changing letters are not contained in any subword that is neutral for i, and u is the longest prefix of w that is anti-dominant for i, while v is the longest suffix of w' that is dominant for i. Hence the expressions $f_i(w)$ and $e_i(w')$, when defined, are unambiguous.

3.1.3. Proposition. *The expression $e_i(w)$ is defined unless w is dominant for i, and $f_i(w)$ is defined unless w is anti-dominant for i.*

Proof. We shall prove the latter statement, the proof of former being analogous. Let u be the longest prefix of w that is anti-dominant for i; clearly $f_i(w)$ cannot be defined if $u = w$. Otherwise $w = uiv$ for some v, and we show by induction on its length that v is dominant for i, which will prove the proposition. The cases where v is empty or ends with a letter other than $i+1$ are trivial, so assume $v = v'(i+1)$ and suppose v is not dominant for i while (by induction) v' is. Then v' has as many letters i as letters $i+1$, and is therefore neutral for i, so that $w = uiv'(i+1)$ is anti-dominant for i since $ui(i+1)$ is; a contradiction. $\square$

If $w = e_i(w')$, then $\operatorname{wt} w > \operatorname{wt} w'$ (for the ordering of §1); therefore the e_i are called *raising operations*, and the f_i are called *lowering operations*. Starting with any $w \in [n]^*$ one can iterate application of a fixed e_i until, after a finite number of iterations, w is transformed into a word that is dominant for i. More generally any sequence of applications of operations e_i, where i is allowed to vary, must eventually terminate, producing a dominant word. For instance, for the word in (13), if we choose at each step to apply the operation e_i with minimal possible i, the sequence of operations applied is $e_0,\ e_3, e_2, e_1, e_0,\ e_4, e_3, e_2, e_1,\ e_4, e_3, e_2, e_1$, which respectively decrease the letters at positions $5,\ 0, 0, 0, 0,\ 7, 10, 10, 11,\ 3, 3, 3, 4$ from the left, leading finally to the dominant word $0\,0\,1\,2\,1\,0\,3\,4\,0\,1\,2\,1\,0$ $0\,1\,2\,3\,3\,4$. Thus the raising operations e_i define a rewrite system on $[n]^*$, whose normal forms are the dominant words. We shall show below that this rewrite system is confluent; for instance the dominant word just obtained from the word in (13) can also be obtained by always applying the e_i with maximal possible i, which leads to the sequence $e_4,\ e_3,\ e_3, e_4,\ e_2,\ e_2, e_3,\ e_1, e_2,\ e_0, e_1,\ e_0, e_1$, respectively affecting letters at positions $7,\ 10,\ 0, 3,\ 10,\ 0, 3,\ 0, 3,\ 5, 11,\ 0, 4$.

The coplactic operations define a labelled directed graph on the set $[n]^*$, with an edge labelled i going from w to w' whenever $f_i(w) = w'$; we shall call this the *coplactic graph* on $[n]^*$. For each $w \in [n]^*$, we shall call the connected component of the coplactic graph on $[n]^*$ containing w the coplactic graph of w; we consider this to be a rooted graph, with as root the element w itself. For $n = 2$, these coplactic graphs are linear with a dominant word at one end and an anti-dominant word at the other end; distinct raising operations do not commute however, so that for $n > 2$ the coplactic graph associated to w can contain other words with the same weight as w (for instance the coplactic graph of $1\,0\,2$ also contains $2\,0\,1$). Coplactic graphs are isomorphic to the crystal graphs for irreducible integrable $U_q(\mathfrak{gl}_n)$-modules of [KaNa]; this places their properties in a broader perspective. Their structure is intricate, and not easy to describe in an independent way; however, as we shall see, it occurs frequently that distinct words have isomorphic coplactic graphs (for instance the coplactic graphs of $1\,0\,2$ and $0\,2\,1$ are isomorphic). The coplactic graph of the word w in (13) for $n = 6$ contains 53460 words, 120 of which have the same weight as w (for instance $2\,0\,1\,5\,3\,1\,4\,5\,0\,1\,3\,2\,0\,0\,1\,2\,4\,3\,4$); we shall not attempt to depict this graph. However, we encourage the reader to draw some small coplactic graphs, e.g., for the word $0\,1\,0\,1$ and $n = 4$ (cf. [LeTh, Figure 2]).

3.2. *Coplactic operations on tableaux.*

As we mentioned, our interest in coplactic operations is in applying them to tableaux rather than to words. For the purpose of defining coplactic operations, the entries of a tableau will be considered as letters of a word that have been mapped in some order onto the squares of a skew diagram. In view of the similarity between dominance of words and 0-dominance of tableaux, it should come as no surprise that the order in which the entries of a tableau are strung together into a word is a valid reading order as discussed in §1.5. The following properties however, which we shall prove below, are quite remarkable. When coplactic operations are applied to the word read off a tableau, modifying its entries in place, the tableau conditions are preserved; this regardless of which valid reading order is used, and despite the loss of information about rows and columns of the skew diagram caused by the reading process. In fact the changes to the entries caused by the coplactic operations are themselves independent of that order, i.e., the same entry is affected, at whatever place in the word the reading order places it. For instance, here is a tableau with the coplactic operations that can be applied to it, labelled as in (13), and two of its similarly labelled reading words (for the Semitic and Kanji reading order, respectively) that could be used to determine those possible coplactic operations.

$$\begin{array}{|c|c|c|c|c|} \hline & & & 0 & \overline{1} \\ \hline & \overline{0} & \underline{1} & 1 & 3 \\ \hline 0 & 2 & \overline{2} & \underline{3} & \\ \hline 1 & 4 & \overline{4} & \underline{5} & \\ \hline \overline{3} & 5 & & & \\ \hline \end{array} \qquad \overline{1}\,0\,3\,1\,\underline{1}\,\overline{0}\,\underline{3}\,\overline{2}\,2\,0\,\underline{5}\,\overline{4}\,4\,1\,5\,\overline{3} \qquad \overline{1}\,3\,0\,1\,\underline{3}\,\underline{5}\,\underline{1}\,\overline{2}\,\overline{4}\,\overline{0}\,2\,4\,5\,0\,1\,\overline{3}$$

We shall now give a more formal definition. A valid reading order for $\chi \in \mathcal{S}$ is a total ordering '$\leq_r$' on $Y(\chi)$, such that $(i, j) \leq_r (i', j')$ whenever $i \leq i'$ and $j \geq j'$. A corresponding map $w_r\colon \mathrm{SST}(\chi, n) \to [n]^*$ is defined by $w_r(T) = T_{s_0} \cdots T_{s_k}$, where $Y(\chi) = \{s_0, \ldots, s_k\}$ with $s_0 <_r \cdots <_r s_k$; in other words, w_r forms a list of all entries of T, in increasing order for '$\leq_r$' of their squares. Now let c be a coplactic operation that can be applied to $w_r(T)$, say $w_r(T) = uiv$ and $c(w_r(T)) = ui'v$ with $i, i' \in [n]$ and $i \neq i'$. Then if the length of u is l, the square s_l in above enumeration is called the variable square for the application of c to T (it contains the copy of the letter i that is changed by c). If replacing i in square s_l of T by i' results in a tableau $U \in \mathrm{SST}(\chi, n)$, we define $U = c(T, \leq_r)$, so that $w_r(c(T, \leq_r)) = c(w_r(T))$.

3.2.1. Proposition. *Let $\chi \in \mathcal{S}$, let $T \in \mathrm{SST}(\chi, n)$ and let c be a coplactic operation e_i or f_i with $i \in [n-1]$. Then for any valid reading order '$\leq_r$', the tableau $c(T, \leq_r)$ is defined if and only if $c(w_r(T))$ is; moreover, this condition and (when it holds) the value of $c(T, \leq_r)$ do not depend on '$\leq_r$'.*

Proof. We shall use a process of reduction: $T \in \mathrm{SST}(\chi, n)$ is simplified by successively removing certain sets of squares from $Y(\chi)$, restricting T and '$\leq_r$' to the remainder. At each step the change to $w_r(T)$ will be the removal of a subword that is neutral for i, so that neither the condition whether or not $c(w_r(T))$ is defined, nor the variable square, is affected. The reduction steps are of two types. The first type is the removal of squares whose entries are not i or $i+1$, as the corresponding subwords of $w_r(T)$ (of length 1) are neutral for i; this reduces us to the case that T has entries i and $i+1$ only. In that case the second type of reduction applies, which removes a maximal rectangle within $Y(\chi)$ consisting of two rows of squares (since the columns of $Y(\chi)$ now have at most 2 squares, maximality means the rectangle cannot be extended into the columns to its left or right). One easily checks that the letters of $w_r(T)$ corresponding to such a rectangle form a subword that is neutral for i. When no further reduction of this type is possible, T is reduced to a tableau with at most one square in any column, and regardless of the original reading order, $w_r(T)$ lists their entries from the rightmost column to the leftmost one. Assuming now that $c(w_r(T))$ is defined, it follows from the definition of coplactic operations for words that the variable square does not have a left neighbour with entry $i + 1$, nor a right neighbour with entry i; this shows in the reduced case that changing the entry of the variable square does not violate weak increase along rows. Neither is it possible that such a neighbouring entry of the variable square was removed by the second type of reduction (the variable square itself would then have to have been in a column of length 2, which it was not), showing that weak increase along rows is preserved in the unreduced case as well; strict increase down columns is preserved since there are no other entries i or $i+1$ in the column of the variable square. $\square$

3.2.2. Definition. *On* $\mathrm{SST}(\chi, n)$ *the coplactic operations* e_i *and* f_i *(for* $i < n$*) are (partially) defined by* $e_i(T) = e_i(T, \leq_r)$ *and* $f_i(T) = f_i(T, \leq_r)$ *for an arbitrary valid reading order '*$\leq_r$*'; this is taken to mean also that the left hand sides are undefined when the corresponding right hand sides are.*

We call $T \in \mathrm{SST}(\chi)$ dominant if $w_r(T)$ is dominant for any (and hence for every) valid reading order r; it means $e_i(T)$ is undefined for all $i \in \mathbf{N}$. From proposition 1.4.3 we get the following characterisation.

3.2.3. Corollary. *For any* $\chi \in \mathcal{S}$*, the subset of dominant elements of* $\mathrm{SST}(\chi)$ *is equal to* $\mathrm{LR}(\chi)$. $\square$

3.3. *The main theorem: commutation.*

As a result of proposition 3.2.1, one obtains many instances of distinct words with isomorphic coplactic graphs, namely words obtained as $w_r(T)$ for fixed T and different reading orders '$\leq_r$'. In fact we shall presently find even more instances than can be found in this manner. The situation resembles to that of jeu de taquin, where we found remarkably many cases of dual equivalent skew semistandard tableaux. We shall now give a theorem that explains both these phenomena, and at the same time essentially proves the Littlewood-Richardson rule. This theorem, which is the only one in our paper with a somewhat technical proof, simply states that jeu de taquin commutes with coplactic operations. This implies that coplactic operations transform tableaux into dual equivalent ones, and that words obtained from jeu de taquin equivalent tableaux using any reading order always have isomorphic coplactic graphs.

3.3.1. Theorem. *Coplactic operations* e_i *and* f_i *on tableaux commute with jeu de taquin slides in the following sense. If* $S \in \mathrm{SST}(\mu/\nu)$, $T \in \mathrm{SST}(\lambda/\mu)$*, and* $X(S, T) = (T', S')$*, then* $e_i(T)$ *is defined if and only if* $e_i(T')$ *is, and if so, one has* $X(S, e_i(T)) = (e_i(T'), S')$*; the same holds when* e_i *is replaced by* f_i*.*

In the course of the proof we shall need to draw some specific configurations that may occur within the tableaux. These will involve entries i and $i+1$ only; in order to fit them into the squares, we subtract i from each, representing them respectively as 0 and 1 (one might also say the drawings assume $i = 0$).

Proof. Since the standardisation of T can be written as $U|V|W$ where U, V, W are the standardisations of the subtableaux of T of entries less than i, in $\{i, i + 1\}$, and greater than $i + 1$, respectively, we can reduce by proposition 2.2.5 to the case that T has entries i and $i+1$ only; we may also assume $|\mu/\nu| = 1$. Since $T_0 = e_i(T_1)$ means the same as $f_i(T_0) = T_1$, it suffices to consider the operations e_i. We may use any valid reading order to determine coplactic operations on tableaux; it will be convenient to use the Kanji reading order, which we shall denote by $\leq_{\mathrm{K}}$. The only differences between $w_{\mathrm{K}}(T)$ and $w_{\mathrm{K}}(T')$ are due to transitions $\begin{array}{|c|c|}\hline & 0 \\ \hline 1 & 1 \\ \hline\end{array} \to \begin{array}{|c|c|}\hline 0 & \\ \hline 1 & 1 \\ \hline\end{array}$ and $\begin{array}{|c|c|}\hline 0 & 0 \\ \hline & 1 \\ \hline\end{array} \to \begin{array}{|c|c|}\hline 0 & 0 \\ \hline 1 & \\ \hline\end{array}$ during the slide $T \triangleright T'$; the effect of these transitions on $w_{\mathrm{K}}(T)$ amounts to interchanging one letter with a word $i(i + 1)$. Since that word is neutral for i, this does not affect the dominance or anti-dominance for i of any subword containing the change. In particular $w_{\mathrm{K}}(T')$ is dominant for i if and only if $w_{\mathrm{K}}(T)$ is, whence $e_i(T)$ is defined if and only if $e_i(T')$ is; we assume henceforth that this is the case. Let v be the variable square for the application of e_i to T, and v' the variable square for the application of e_i to T'. We can find v' from v by tracing for each intermediate "tableau with empty square" $\tilde{T}$ during the slide $T \triangleright T'$ which of its entries $i+1$ corresponds to the letter affected by e_i in $w_{\mathrm{K}}(\tilde{T})$ (like before, we shall label this entry in drawings by underlining it). This amounts to moving along with that entry if it slides into another square, and switching to the entry to its right if an entry i moves into its present column, giving rise to a transition $\begin{array}{|c|c|}\hline & 0 \\ \hline \underline{1} & 1 \\ \hline\end{array} \to \begin{array}{|c|c|}\hline 0 & \\ \hline 1 & \underline{1} \\ \hline\end{array}$.

We now distinguish two cases, depending on whether or not the path of the inward slide applied to T (i.e., the set of squares whose entries move) is the same as for the slide applied to $e_i(T)$. Suppose first that the paths are the same, then the entry of v makes the same move, if any, during the two slides;

since this entry is $i+1$ in one case and i in the other, it must be unique in its column both before and after the slide. By the above description of v', we see that v' coincides with final position of the entry of v; this implies the theorem for this first case.

Suppose next that the paths of the two slides differ, which must be caused by the change of the entry of v. This means that v occurs in exactly one of these paths; since the rule for an inward jeu de taquin slide is to select the smallest candidate entry to move, v must be in the path of the slide applied to $e_i(T)$, where it has entry i rather than $i+1$. We can see as follows that this move of the entry i of v in $e_i(T)$ cannot be to the left: this would mean that in the slide $T \triangleright T'$, where the entry $i+1$ of v does not move, the square to its left is filled by another entry $i+1$ sliding up; the path of that slide therefore cannot involve the square above v, whence our earlier reasoning gives $v' = v$, which is absurd since $e_i(T')$ must have weakly increasing rows. Therefore the move of the entry i of v in $e_i(T)$ must be upwards, and it is the first move of the slide. It is a transition $\begin{array}{|c|c|}\hline & 0 \\ \hline 0 & 1 \\ \hline\end{array} \to \begin{array}{|c|c|}\hline 0 & 0 \\ \hline & 1 \\ \hline\end{array}$ whereas in T we have $\begin{array}{|c|c|}\hline & 0 \\ \hline \underline{1} & 1 \\ \hline\end{array} \to \begin{array}{|c|c|}\hline 0 & \\ \hline 1 & \underline{1} \\ \hline\end{array}$ (the presence of the entry at the bottom right follows from the fact that v is the variable square: the Kanji reading $w_{\mathrm{K}}(R)$ of the subtableau R of T consisting of the columns to the right of v is anti-dominant for i). The latter transition may be followed by a number of similar transitions further to the right, where each time the anti-dominance of $w_{\mathrm{K}}(R)$ guarantees the presence of the entry at the bottom right; eventually they must be followed by a final transition $\begin{array}{|c|c|}\hline 0 & \\ \hline 1 & \underline{1} \\ \hline\end{array} \to \begin{array}{|c|c|}\hline 0 & \underline{1} \\ \hline 1 & \\ \hline\end{array}$. The relevant part of the slide for T looks like $\begin{array}{|c|c|c|c|}\hline & 0 & 0 & 0 \\ \hline \underline{1} & 1 & 1 & 1 \\ \hline\end{array} \triangleright \begin{array}{|c|c|c|c|}\hline 0 & 0 & 0 & \underline{1} \\ \hline 1 & 1 & 1 & \\ \hline\end{array}$, and the location of the underline after the slide indicates the variable square v' in T'. The corresponding part of the slide for $e_i(T)$ is $\begin{array}{|c|c|c|c|}\hline & 0 & 0 & 0 \\ \hline 0 & 1 & 1 & 1 \\ \hline\end{array} \triangleright \begin{array}{|c|c|c|c|}\hline 0 & 0 & 0 & 0 \\ \hline 1 & 1 & 1 & \\ \hline\end{array}$, whose result matches $e_i(T')$; this proves the theorem for this second and final case. $\square$

Figure 1 illustrates the commutation of jeu de taquin and coplactic operations. In this example the operation e_i with smallest possible i is always chosen; several of them have been combined at each step, to prevent the display from getting excessively large. The reader is urged to study the transitions carefully.

3.3.2. Corollary. *Equation (11) holds, which establishes the validity of the Littlewood-Richardson rule.*

Proof. Using corollary 3.2.3, theorem 3.3.1 implies that jeu de taquin preserves the property of being a Littlewood-Richardson tableau. For $L \in \mathrm{SST}(\chi)$ there exists a tableau of partition shape $P \in \mathrm{SST}(\nu/(0))$ with $L \triangleright P$; then P is a Littlewood-Richardson tableau if and only if $P = \mathbf{1}_\nu$, in which case $\nu = \mathrm{wt}\, L$. $\square$

Note that we did not use the commutation statement of theorem 3.3.1, just the (easier) statement about when $e_i(T)$ is defined. Neither did we use confluence of jeu de taquin here, although it was used to obtain corollary 2.5.3. But in fact theorem 3.3.1 independently proves this confluence, and more.

3.3.3. Corollary. *On the set of all skew semistandard tableaux, the two rewrite systems defined by inward jeu de taquin slides, respectively by the raising operations e_i, are both confluent (their normal forms are the semistandard Young tableaux, and the Littlewood-Richardson tableaux, respectively). Moreover the two normal forms of a skew semistandard tableau T uniquely determine T.*

Proof. Since the two rewrite systems commute in the precise sense of theorem 3.3.1, the set of normal forms for one system is closed under the rewrite rules of the other system. Confluence of that rewrite system on this set of normal forms is clear, since the only tableaux that are normal forms for both systems simultaneously are the tableaux $\mathbf{1}_\lambda$ for $\lambda \in \mathcal{P}$, and for a tableau that is a normal form for one of the systems the pertinent value of λ can be directly read off as the shape (in the case of semistandard Young tableaux) respectively as the weight (in the case of Littlewood-Richardson tableaux).

Now let T be any skew semistandard tableau, let P be a tableau of partition shape obtained from T by a sequence of inward jeu de taquin slides, and let $\tilde{P}$ designate the sequence of shapes of the intermediate tableaux. Similarly let L be a Littlewood-Richardson tableau obtained from T by applying a sequence $\tilde{L}$ of raising operations. As normal forms for one system, P and L each have a unique normal form for the other system. By theorem 3.3.1, the normal form of P for raising operations can be obtained by applying the sequence $\tilde{L}$ of such operations, while the normal form of L for jeu de taquin can be obtained by applying a sequence of slides with $\tilde{P}$ as sequence of intermediate shapes; moreover the two normal forms are the same tableau $\mathbf{1}_\lambda$ (where $P \in \mathrm{SST}(\lambda/(0))$ and $\mathrm{wt}\, L = \lambda$). Now P can be reconstructed from $\mathbf{1}_\lambda$

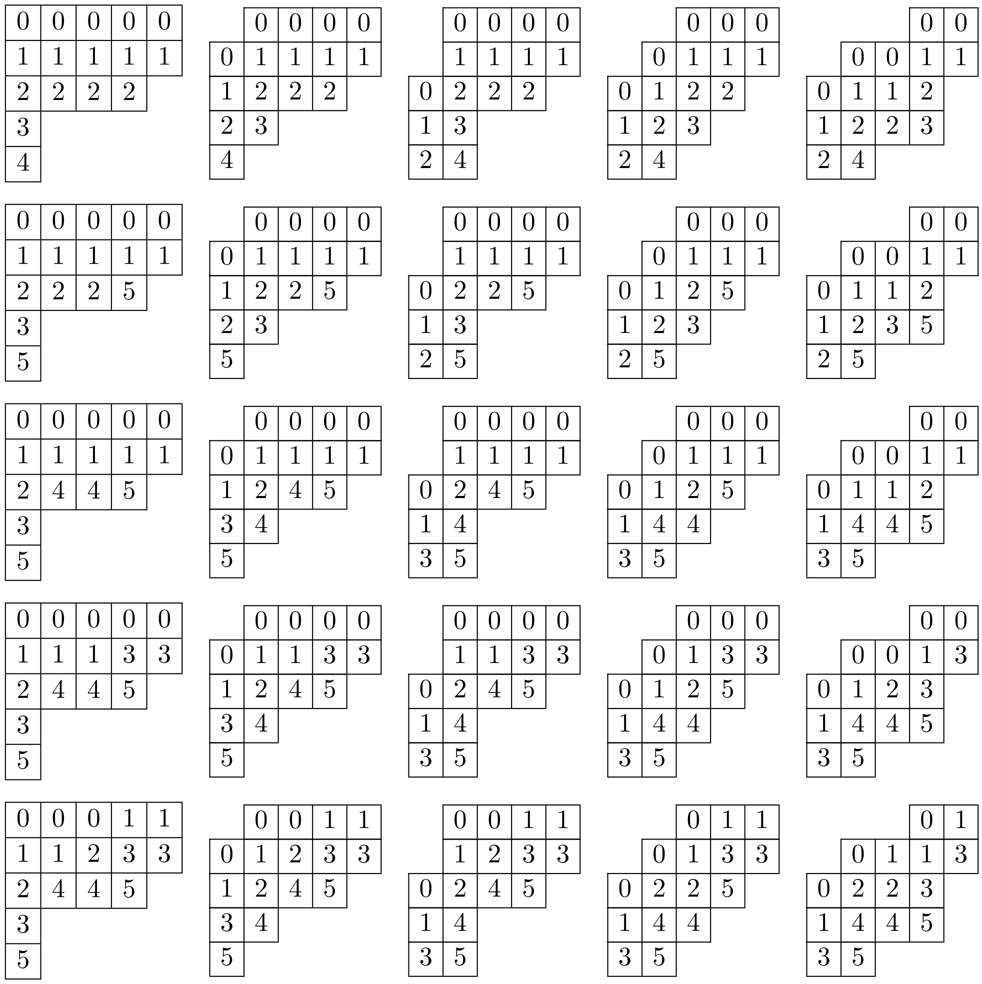

Figure 1. Commutation of jeu de taquin (right to left) and raising operations (upwards). Raising operations are grouped: e_0, e_0, e_1; e_2, e_1, e_2, e_1; e_3, e_2, e_3, e_2; e_4, e_4, e_3, e_2.

by reversing the the raising operations of $\tilde{L}$, and is therefore independent of the sequence $\tilde{P}$ used to obtain it from T; similarly L can be reconstructed using $\tilde{P}$, and therefore independent of $\tilde{L}$. Thus both rewrite systems are confluent. If L and P are given, then T can be reconstructed by a similar process: for instance, a sequence $\tilde{P}$ of shapes can be found by reducing L by inward jeu de taquin slides into $\mathbf{1}_\lambda$, and since the latter is also obtainable from P by raising operations, a sequence of outward jeu de taquin slides can be applied to P involving the shapes of $\tilde{P}$ in reverse order, which produces T as result. □

It may be noted that parts of the argument above could have been formulated in more concrete terms, e.g., by using a skew standard tableau to encode the information of $\tilde{P}$ and using tableau switching; we have not done this in order to stress the conceptual simplicity of the argument, its symmetry with respect to the two rewrite systems, and the fact that no detailed knowledge about these systems is used. On the other hand the computation of T from L and P above can easily be seen to be equivalent to that of $\phi(L,P)$ in corollary 2.5.2, so that the following defines a specialisation of $\mathcal{R}'$, as claimed at the end of §2.

3.3.4. Corollary. *For $\chi \in \mathcal{S}$ and $n \in \mathbf{N}$, Robinson's bijection $\mathcal{R}$ (i.e., one satisfying the specification (9)) can be defined by $\mathcal{R}(T) = (\mathcal{R}_0(T), \mathcal{R}_1(T))$ for $T \in \mathrm{SST}(\chi, n)$, where $\mathcal{R}_0(T) \in \mathrm{LR}(\chi, \nu)$ is the normal form of T for the rewrite system of raising operations, and the $\mathcal{R}_1(T) \in \mathrm{SST}(\nu, n)$ is the normal form of T for the rewrite system of inward jeu de taquin slides.* □

3.4. *Jeu de taquin on companion tableaux.*

Having completed our discussion of the Littlewood-Richardson rule proper, it is worth while to point out a remarkable connection between coplactic operations and jeu de taquin performed on companion tableaux. Note that we have already found, in corollary 3.2.3, that T has a companion tableau that is a normal form for jeu de taquin (i.e., T is a Littlewood-Richardson tableau) if and only if T itself is a normal form for raising operations; the connection we shall establish generalises this. It is closely related to the theory of pictures, where instead of coplactic operations one has a second form of jeu de taquin, with a similar commutation theorem [vLee2, theorem 5.3.1]. The symmetry exhibited by this connection has some important implications, which we shall only indicate briefly here. Firstly, it implies the fundamental symmetry of the Robinson-Schensted correspondence (see for instance [vLee1, 3.2]), which is more manifest for Schensted's formulation of the correspondence [Sche] than for Robinson's bijection; however, pictures clearly exhibit the link between these formulations [vLee2, theorem 5.2.3]. Secondly, the connection clarifies the bijection corresponding to $c^{\nu}_{\lambda,\mu} = c^{\nu}_{\mu,\lambda}$: as remarked at the end of §1, this bijection involves tableau switching performed on a companion tableau of (part of) the Littlewood-Richardson tableau. A detailed discussion of that bijection, and a generalisation, can be found in [vLee4].

Our starting point is proposition 3.1.1. It gives an interpretation of standard Young tableaux in terms of words, but we shall consider more generally skew standard tableaux. To that end we extend the notion of dominance for words, in analogy to that for tableaux, to κ-dominance: a word $w \in [n]^*$ will be called κ-dominant, where $\kappa \in \mathcal{P}_{d,n}$ for some d, if pw is dominant for some dominant word p of weight κ (this clearly does not depend on the choice of p). Now proposition 1.4.3 can be interpreted as saying that a tableau T is κ-dominant if and only if $w_{\mathrm{S}}(T)$ is κ-dominant.

Proposition 3.1.1 now generalises as follows: for any $\nu \supseteq \kappa$, the κ-dominant words of weight $\nu - \kappa$ are in bijection with $\mathrm{ST}(\nu/\kappa)$. To see this one chooses some fixed dominant word p of weight κ, and associates to any κ-dominant word w the skew standard tableau obtained by applying proposition 3.1.1 to pw, and extracting from the resulting chain of partitions $T \in \mathrm{ST}(\nu)$ the subchain from κ to ν (again this subchain does not depend on the choice of p).

3.4.1. Proposition. *Let $w \in [n]^*$ be a κ-dominant word, with corresponding tableau $S \in \mathrm{ST}(\nu/\kappa)$; let $S' \in \mathrm{ST}(\nu'/\kappa')$ be obtained from S by an inward jeu de taquin slide into a square in row k, ending in row l. Then the sequence of words $w = w_k, w_{k+1}, \ldots, w_l$ given by $w_{i+1} = e_i(w_i)$ for $k \leq i < l$ is well defined, and w_l corresponds to S'; moreover w_l is the first word in the sequence that is κ'-dominant.*

Proof. Fix a dominant word p of weight κ. When the prefixes of pw are listed in increasing order, adding a letter i of w will increase part i of the partition in the chain S that is the weight of the prefix, and therefore add a square in row i of its diagram. To make more explicit the correspondence of the letter to that square, we attach to each letter a subscript that we shall call its "ordinate", describing the column of the square: the ordinate of the leftmost occurrence of i in w is κ_i, and for the remaining occurrences of i the ordinates increase from left to right by unit steps. Note that it is also true that among the letters with fixed ordinate j, the value of the letters increases from left to right by unit steps. Let $\widetilde{w}$ be the word obtained by so augmenting the letters of w with ordinates.

The slide $S \triangleright S'$ can be described by tableau switching: $X(E, S) = (S', E')$, where E is the unique element of $\mathrm{ST}(\kappa/\kappa')$ (i.e., the chain $(\kappa' \subset \kappa)$; the only square (k, κ'_k) of E is the initial position of the empty square), and $E' \in \mathrm{ST}(\nu/\nu')$ is similarly unique (its square (l, ν'_l) is the final position of the empty square). Since each letter i_j of $\widetilde{w}$ describes the square (i, j) added at the corresponding place in the chain S, we can mirror this tableau switching computation by operations on $\widetilde{w}$. We add a (subscripted) letter to $\widetilde{w}$, designating the empty square, which is distinguished from the other letters; we shall indicate it by a caret. Initially this is $\widehat{k_{\kappa'_k}}$, added at the left end of $\widetilde{w}$. We successively move this distinguished letter to the right through the word, making adjustments according to the rules for tableau switching: when $\widehat{i_j}$ is moved across a letter i_{j+1} or $(i+1)_j$, it assumes the value and ordinate of that letter, while that letter becomes i_j after the switch (so $\widehat{i_j}\, i_{j+1} \to i_j\, \widehat{i_{j+1}}$ and $\widehat{i_j}\, (i+1)_j \to i_j\, \widehat{(i+1)_j}$); otherwise the letters are simply interchanged, preserving their own value and ordinate ($\widehat{i_j}\, a_b \to a_b\, \widehat{i_j}$).

We shall show that if the successive intermediate words obtained from $\widehat{k_{\kappa'_k}}\, \widetilde{w}$ are "stripped", by which we mean removing the distinguished letter and all ordinates, then we obtain a sequence of words over $[n]$ that reduces to $w_k, \ldots, w_l$ by removing repeated occurrences of words. Switches $\widehat{i_j}\, a_b \to a_b\, \widehat{i_j}$ and

$\widehat{i_j}\, i_{j+1} \to i_j\, \widehat{i_{j+1}}$ clearly leave the stripped word unchanged. What we must show, is that in the remaining case, which with the full context of a prefix $\widetilde{u}$ and suffix $\widetilde{v}$ takes the form $\widetilde{u}\,\widehat{i_j}\,(i+1)_j\,\widetilde{v} \to \widetilde{u}\,i_j\,\widehat{(i+1)}_j\,\widetilde{v}$, the stripped words undergo a coplactic operation e_i; this requires that the word u obtained by stripping $\widetilde{u}$ be anti-dominant for i, and that the word v obtained by stripping $\widetilde{v}$ be dominant for i. With the distinguished letter taken into account, all switches preserve the properties mentioned above of "increase by unit steps". Hence in $\widetilde{u}$, any letter $i_{j'}$ has a letter $(i+1)_{j'}$ somewhere to its right (and one has $j' < j$), whence u is anti-dominant for i; similarly v is dominant for i because any letter $(i+1)_{j'}$ in $\widetilde{v}$ has a letter $i_{j'}$ to its left.

What remains is to prove is that the stripped word is not κ'-dominant until the last switch of the form $\widehat{i_j}\,(i+1)_j \to i_j\,\widehat{(i+1)}_j$. The argument given that such switches affect the stripped word by application of e_i remains valid when a dominant word p' of weight κ', properly augmented, is prepended to $\widehat{k_{\kappa'_k}}\,\widetilde{w}$ and words obtained from it (now among letters with any fixed value, the ordinates start with 0 at the left). So $p' w_{i+1} = e_i(p' w_i)$, showing that w_i is not κ'-dominant. □

We now pass from κ-dominant words to κ-dominant tableaux. If a κ-dominant tableau T has $\bar{T} \in \mathrm{SST}(\nu/\kappa)$ as companion tableau, then it follows from the definitions that the element of $\mathrm{ST}(\nu/\kappa)$ corresponding to the κ-dominant word $w_{\mathrm{S}}T$ of weight $\nu - \kappa$ is equal to the standardisation of $\bar{T}$: the entries in any row i of T, traversed in the Semitic reading order, are the row numbers of the squares with entry i in $\bar{T}$, listed from left to right. Applying this to the proposition, we obtain:

3.4.2. Theorem. *Let $\bar{T} \in \mathrm{SST}(\nu/\kappa)$ be a companion tableau of $T \in \mathrm{SST}(\chi)$, and let $\bar{T}' \in \mathrm{SST}(\nu'/\kappa')$ be obtained from $\bar{T}$ by an inward jeu de taquin slide into a square in row k, ending in row l. Then the sequence $T = T_k, T_{k+1}, \ldots, T_l \in \mathrm{SST}(\chi)$ given by $T_{i+1} = e_i(T_i)$ for $k \le i < l$ is well defined, and $\bar{T}'$ is a companion tableau of T_l; moreover T_l is the first tableau in the sequence that is κ'-dominant.* □

As an example, we apply this construction to the companion tableaux $T, \bar{T}$ of (5), and $k = 0$. The word $w = w_{\mathrm{S}}(T) = 1\,0\,3\,1\,1\,0\,3\,2\,2\,0\,5\,4\,4\,1\,5\,3$, gets augmented as $\widetilde{w} = 1_4\,0_6\,3_2\,1_5\,1_6\,0_7\,3_3\,2_4\,2_5\,0_8\,5_0\,4_1\,4_2\,1_7\,5_1\,3_4$, since $\kappa = (6,4,4,2,1,0)$. To simulate the slide into the square $(0,5)$, we prepend $\widehat{0_5}$ to $\widetilde{w}$, and then move this distinguished letter across the successive letters of $\widetilde{w}$. In the process, it exchanges its contents successively with 0_6, 1_6, and 1_7, with as final result the augmented word $1_4\,0_5\,3_2\,1_5\,0_6\,0_7\,3_3\,2_4\,2_5\,0_8\,5_0\,4_1\,4_2\,1_6\,5_1\,3_4\,\widehat{1_7}$. After stripping the words, we get the relation $e_0(1\,0\,3\,1\,1\,0\,3\,2\,2\,0\,5\,4\,4\,1\,5\,3) = 1\,0\,3\,1\,0\,0\,3\,2\,2\,0\,5\,4\,4\,1\,5\,3$. The pair of companion tableaux is transformed into the following one:

$$e_0(T) = \begin{array}{ccccc} & & & 0 & 1 \\ & 0 & 0 & 1 & 3 \\ 0 & 2 & 2 & 3 & \\ 1 & 4 & 4 & 5 & \\ 3 & 5 & & & \end{array}\,, \quad \bar{T}' = \begin{array}{ccccccccc} & & & & & 0 & 1 & 1 & 2 \\ & & & & 0 & 1 & 3 & & \\ & & & & 2 & 2 & & & \\ & & 1 & 2 & 4 & & & & \\ & 3 & 3 & & & & & & \\ 3 & 4 & & & & & & & \end{array}\;; \tag{14}$$

this shows that $e_0(T)$ is κ'-dominant for $\kappa' = (5,4,4,2,1,0)$. A further slide into the square $(0,4)$ causes $\widehat{0_4}$ to traverse the augmented word, exchanging its contents with successively 1_4, 1_5, and 2_5, and transforming the word into $0_4\,0_5\,3_2\,1_4\,0_6\,0_7\,3_3\,2_4\,1_5\,0_8\,5_0\,4_1\,4_2\,1_6\,5_1\,3_4$. On the stripped word we have application of $e_1 \circ e_0$, and the resulting pair of companion tableaux is

$$(e_1 \circ e_0)(e_0(T)) = \begin{array}{ccccc} & & & 0 & 0 \\ & 0 & 0 & 1 & 3 \\ 0 & 1 & 2 & 3 & \\ 1 & 4 & 4 & 5 & \\ 3 & 5 & & & \end{array}\,, \quad \bar{T}'' = \begin{array}{ccccccccc} & & & & 0 & 0 & 1 & 1 & 2 \\ & & & & 1 & 2 & 3 & & \\ & & & & 2 & & & & \\ & & 1 & 2 & 4 & & & & \\ & 3 & 3 & & & & & & \\ 3 & 4 & & & & & & & \end{array}\;. \tag{15}$$

Writing the letters of the augmented words back into T, one finds the transitions

$$\begin{array}{ccccc} & & & 0_6 & 1_4 \\ & 0_7 & 1_6 & 1_5 & 3_2 \\ 0_8 & 2_5 & 2_4 & 3_3 & \\ 1_7 & 4_2 & 4_1 & 5_0 & \\ 3_4 & 5_1 & & & \end{array} \xrightarrow{e_0} \begin{array}{ccccc} & & & 0_5 & 1_4 \\ & 0_7 & 0_6 & 1_5 & 3_2 \\ 0_8 & 2_5 & 2_4 & 3_3 & \\ 1_6 & 4_2 & 4_1 & 5_0 & \\ 3_4 & 5_1 & & & \end{array} \xrightarrow{e_1 \circ e_0} \begin{array}{ccccc} & & & 0_5 & 0_4 \\ & 0_7 & 0_6 & 1_4 & 3_2 \\ 0_8 & 1_5 & 2_4 & 3_3 & \\ 1_6 & 4_2 & 4_1 & 5_0 & \\ 3_4 & 5_1 & & & \end{array}. \tag{16}$$

This last representation provides another perspective to proposition 3.2.1. First observe that the ordinates of the entries of T with a fixed value i increase *from right to left* by unit steps, starting with the ordinate κ_i. It follows that reading the augmented entries according to any valid reading order '$\leq_r$' will produce a properly augmented word; stated differently, augmenting (for κ) any $w_r(T)$ and then writing the augmented letters back into T along the reading order '$\leq_r$' always produces the same augmented tableau. Also, the values of the entries of the augmented tableau T with a fixed ordinate j increase from top to bottom, by what was said in §1.5. The fact that during the first transition $\widehat{0_5}$ exchanges its contents with 0_6 rather than with 1_5 is due to the fact that 0_6 precedes 1_5 in (the augmentation of) $w_{\mathrm{S}}(T)$, and similarly the exchange with 1_6 was made because 1_6 precedes 0_7 (for the final exchange with 1_7, there is no alternative). It can be verified that these ordering relations are unchanged when an augmentation of any reading $w_r(T)$ is used instead of that of $w_{\mathrm{S}}(T)$. This is due to the placement of these entries in the augmented tableau T, and the same is true for the relevant pairs of entries $(1_4, 0_5)$, $(1_5, 2_4)$ and $(2_5, 1_6)$ in the middle augmented tableau of (16). The fairly easily proved fact that this is always the case (cf. [vLee2, lemma 5.1.2]) provides an alternative proof of proposition 3.2.1.

Using theorem 3.4.2, one can express coplactic operations on T in terms of jeu de taquin slides on a suitable companion tableau of T. To compute $e_k(T)$ (when defined), it suffices to find a companion tableau $\bar{T}$ for which one has $l = k+1$ in the theorem; this can be achieved by choosing ν/κ in such a way that $\kappa_k - \kappa_{k+1}$ is minimal, and that $\nu_{k+2} < \kappa_k$ (the latter condition can be weakened considerably). Conversely, if one wishes to express jeu de taquin slides on a companion tableau of T in terms of coplactic operations on T, it is necessary to replace the condition of κ'-dominance in the theorem by a condition stated in terms of coplactic graphs. This is possible using the following proposition.

3.4.3. Proposition. *Let $T \in \mathrm{SST}(\chi, n)$, and let ν/κ be a skew shape with $\nu - \kappa = \mathrm{wt}\, T$. Then the following are equivalent*

(i) T is κ-dominant;

(ii) no operation e_i can be applied more than $\kappa_i - \kappa_{i+1}$ times successively to T;

(iii) no operation f_i can be applied more than $\nu_i - \nu_{i+1}$ times successively to T.

For instance, for the tableau T of (5) used in the example above we have $\kappa = (6, 4, 4, 2, 1, 0)$ and $\nu = (9, 8, 6, 5, 3, 2)$, so the proposition states that e_1 cannot be applied to T, that e_3, e_4, f_0, f_2 and f_4 cannot be applied more than once to T, and that e_0, e_2, f_1, and f_3 cannot be applied more than twice; this can be verified. One finds moreover that the statement can be sharpened for e_3 (which cannot be applied) and for f_3 (which can be applied only once); this corresponds to the fact that T is also ν'/κ'-dominant, with $\nu' = (8, 7, 5, 4, 3, 2)$ and $\kappa' = (5, 3, 3, 1, 1, 0)$.

Proof. We first reduce to a statement about words, by replacing T by $w = w_{\mathrm{S}}(T)$. Let p be any dominant word of weight κ, denote by r_i and s_i the maximal number of times e_i respectively f_i can be applied to w, and put $c_i = \kappa_i - \kappa_{i+1}$, $d_i = \nu_i - \nu_{i+1}$. We shall prove for any $i \in [n-1]$ equivalence between: (1) pw is dominant for i, (2) $r_i \leq c_i$, (3) $s_i \leq d_i$. Since $\mathrm{wt}\, w = \nu - \kappa$ implies $r_i - s_i = c_i - d_i$, (2) and (3) are equivalent. We prove equivalence of (1) and (2) for any word $w \in [n]^*$. Subwords that are neutral for i neither affect condition (1) nor the value of r_i. We remove such subwords from w, reducing it to $(i+1)^{r_i} i^{s_i}$, and in (1) also from p, reducing it to i^{c_i}. What remains is to show that $i^{c_i}(i+1)^{r_i} i^{s_i}$ is dominant for i if and only if $r_i \leq c_i$, but this is obvious. □

3.4.4. Corollary. *Let $T \in \mathrm{SST}(\chi)$ and $T' \in \mathrm{SST}(\chi')$ be jeu de taquin equivalent, and ν/κ a skew shape. Then T is ν/κ-dominant if and only if T' is so, in which case the companion tableaux of T and T' of shape ν/κ are dual equivalent.*

Proof. By theorem 3.3.1, T and T' have isomorphic coplactic graphs, and proposition 3.4.3 then shows that T is ν/κ-dominant if and only if T' is. Let $\bar{T}$ and $\bar{T}'$ respectively be the companion tableaux of T and T' of shape ν/κ, and consider an inward jeu de taquin slide into the same square applied to each of them. By theorem 3.4.2, the result will in either case be a companion tableau of a tableau obtained by a sequence of raising operations from T respectively from T'. Moreover, since the condition determining the length of those sequences can be expressed in terms of the coplactic graphs of T and T' by proposition 3.4.3, the two sequences will be identical. This implies on one hand that the tableaux $\widetilde{T}, \widetilde{T}'$ obtained by applying the sequences are jeu de taquin equivalent (by theorem 3.3.1), and on the other hand that the jeu de taquin slides applied to $\bar{T}$ and $\bar{T}'$ leave the empty square in the same row, and therefore result in tableaux of equal shape. As these resulting tableaux are companion tableaux of the jeu de taquin equivalent tableaux $\widetilde{T}$ and $\widetilde{T}'$, respectively, we have arrived at a situation similar to our point of depart, but with T, T' replaced by $\widetilde{T}$, $\widetilde{T}'$. By the same reasoning the shapes of the companion tableaux will remain equal when further inward slides are applied, and by proposition 2.4.2 this proves that $\bar{T}$ and $\bar{T}'$ are dual equivalent. $\square$

This proof shows that jeu de taquin slides performed on T commute with jeu de taquin slides performed on a companion tableau of T (since clearly $\widetilde{T}$ and $\widetilde{T}'$ are linked by the same sequence of jeu de taquin slides as T and T'); this is essentially what is stated in the main theorem 5.3.1 of [vLee2], but the proof here is simpler.

One may ask how practical this alternative method of computing coplactic operations is. For performing single coplactic operations, or for following prescribed paths in the coplactic graph (such as the one used in figure 1), the method is cumbersome, as one (repeatedly) has to adapt the shape ν/κ to single out the required coplactic operation. (Note however that the path that always applies e_i with the *largest* possible i can be found without such adaptations.) On the other hand, when the goal is just to compute $\mathcal{R}_0(T)$, the fact that one often gets more than one coplactic operation at a time is an advantage, as is the absence of a need to repeatedly recompute globally defined quantities, such as the location of the square affected by some e_i. In addition, this method is less error prone when used for manual computation. For instance in (5), one can find by computing

$$\begin{array}{ccccccccc} &&&&&&0&1&2\\ &&&&0&1&1&3&\\ &&&&2&2&&&\\ &&1&2&4&&&&\\ &3&3&&&&&&\\ 3&4&&&&&&& \end{array} \;\triangleright\; \begin{array}{ccccc} 0&0&1&1&2\\ 1&1&2&2&3\\ 2&3&3&4&\\ 3&&&&\\ 4&&&& \end{array} \qquad \text{that} \qquad \mathcal{R}_0\left(\begin{array}{ccccc} &&&0&1\\ &0&1&1&3\\ 0&2&2&3&\\ 1&4&4&5&\\ 3&5&&& \end{array}\right) = \begin{array}{ccccc} &&&0&0\\ &0&0&1&1\\ 0&1&1&2&\\ 1&2&2&3&\\ 2&4&&& \end{array}.$$

§4. Some historical comments.

Now that we have seen the Littlewood-Richardson rule from a modern perspective, let us look at some elements of its intriguing history, in particular the two papers [LiRi] and [Rob], written in the 1930's.

4.1. *The paper by Littlewood and Richardson.*

The paper in which the Littlewood-Richardson rule is first stated, is mainly concerned with symmetric group characters and Schur functions (a term introduced in that paper, though mostly contracted to "S-functions"). Out of 16 sections, only §8 deals with the multiplication of S-functions. Remarkably, semistandard tableaux do not occur explicitly, and in particular are not used in the definition of Schur functions; instead these are expressed in terms of power sum symmetric functions using symmetric group characters. In fact no attempt is made at all to express Schur functions in terms of monomials, or even of minimal symmetric polynomials m_λ, although curiously the opposite *is* done, in §5.

Semistandard tableaux do occur implicitly, as follows. For $\mu = (r) \in \mathcal{P}_r$ one has $s_\mu = h_r$, and every skew semistandard tableau of weight μ is automatically a Littlewood-Richardson tableau; call the shape of such a tableau a horizontal strip. As starting point for multiplication of S-functions, it is shown that $s_\lambda h_r$ is the sum of all s_ν with ν/λ a horizontal strip (this agrees with our proposition 1.4.5). By iteration this implies that $s_\lambda h_{\alpha_0} \cdots h_{\alpha_l}$ is the sum of s_ν taken over all ways to successively add horizontal strips of sizes $\alpha_0, \ldots, \alpha_l$ to $Y(\lambda)$ yielding $Y(\nu)$. The number of ways to so obtain a given partition ν is $\#\{\, T \in \mathrm{SST}(\nu/\lambda) \mid \mathrm{wt}\, T = \alpha \,\}$.

For the general case of multiplying S-functions, the following rule is formulated ([LiRi, p. 119]):

> "THEOREM III.—*Corresponding to two* S*-functions* $\{\lambda_1, \ldots, \lambda_p\}$, $\{\mu_1, \ldots, \mu_q\}$ *build tableaux* A *and* B *as in Theorem II. Then in the product of these two functions, the coefficient of any* S*-function* $\{\nu_1, \nu_2, \ldots\}$ *is equal to the number of compound tableaux including all of the symbols of* A *and* B*, and corresponding to* $\{\nu_1, \nu_2, \ldots\}$*, that can be built according to the following rules.*
>
> *Take the tableau* A *intact, and add to it the symbols of the first row of* B*. These may be added to one row of* A*, or the symbols may be divided without disturbing their order, into any number of sets, the first set being added to one row of* A*, the second set to a subsequent row, the third to a row subsequent to this, and so on. After the addition no row must contain more symbols than a preceding row, and no two of the added symbols may be in the same column.*
>
> *Next add the second row of symbols from* B*, according to the same rules, with this added restriction. Each symbol from the second row of* B *must appear in a later row of the compound tableau than the symbol from the first row in the same column.*
>
> *Similarly add each subsequent row of symbols from* B*, each symbol being placed in a later row of the compound tableau than the symbol in the same column from the preceding row of* B*, until all the symbols of* B *have been used."*

The tableaux A and B are Young diagrams of shapes λ and μ, filled with formal symbols. The rows of B are rearranged into horizontal strips filling $Y(\nu/\lambda)$; as indicated above, such a rearrangement corresponds to some $T \in \mathrm{SST}(\nu/\lambda)$ with $\mathrm{wt}\, T = \mu$. The restriction added in the last two paragraphs depends on the particular rule that is used to ensure that each horizontal strip is obtained in only one way from a given row of B; this rule can be rephrased as stating that within each row of B, the number of the destination row in $Y(\nu/\lambda)$ of each symbol increases weakly from left to right. If we replace each symbol in B by this number of its destination row, then the resulting filling of $Y(\mu)$, which has weakly increasing rows by construction, is required in the last two paragraphs to also have strictly increasing columns; in our terminology, they form a semistandard Young tableau $\bar{T} \in \mathrm{SST}(\mu/(0))$. In fact, $\bar{T}$ can be seen to be a companion tableau of T in the sense of the current paper. If we interpret the rearrangement of symbols as a bijection $Y(\mu) \to Y(\nu/\lambda)$, then we almost arrive at the notion of pictures; the only difference is that when some sequence of symbols from a row of B are moved to a common destination row, this happens "without disturbing their order", rather than by reversing their order, as is done in the case of pictures.

Although claimed to be generally valid, Theorem III is only proved in [LiRi] when μ has at most two parts (the first sentence after the statement of the Theorem is "No simple proof has been found that will demonstrate it in the general case"; indeed "simple" can be omitted). We already mentioned the case $\mu = (r)$. The proof for a partition $\mu = (q, r)$ with two parts $q \geq r > 0$ is based on the identity $s_\mu = h_q h_r - h_{q+1} h_{r-1}$ (an instance of a determinantal formula known as the Jacobi-Trudi identity), which means that $c^\nu_{\lambda,\mu}$ can be found by subtracting the coefficient of s_ν in $s_\lambda h_{q+1} h_{r-1}$ from the one in $s_\lambda h_q h_r$; these coefficients can be determined by counting the tableaux in $\mathrm{SST}(\nu/\lambda, 2)$ of weights $(q+1, r-1)$ and (q, r), respectively. To show that the difference matches the number of tableaux given by Theorem III, the rule is reformulated in terms of lattice permutations (cf. proposition 1.4.5), and a bijective correspondence in $\mathrm{SST}(\nu/\lambda, 2)$ is given between tableaux of weight (q, r) that do not satisfy the rule, and all tableaux of weight $(q+1, r-1)$. It is shown that this transformation, which coincides with our e_0, preserves semistandardness.

It is remarkable that the authors state their rule as a Theorem when, by their own admission, they only have a proof for some very simple cases. In part this can be attributed to the general attitude at the time, which appears to have been that combinatorial statements are less in need of a proof than, say, algebraic statements; on the other hand, they do devote three pages to a proof of the special cases. It appears that they viewed their rule mainly as a computational device, useful to find a result, whose correctness may then be verified by other means. They work out a complete example for the computation of $s_{(4,3,1)} s_{(2,2,1)}$ (in their notation $\{4, 3, 1\} \times \{2^2 1\}$), displaying all 34 tableaux contributing to the result. They use

$$\begin{pmatrix} a, & b, & c, & d \\ e, & f, & g & \\ h & & & \end{pmatrix}, \quad \begin{pmatrix} \alpha, & \beta \\ \gamma, & \delta \\ \varepsilon & \end{pmatrix},$$

as tableaux A, B; replacing the symbols of A by 0's because they don't move, they then display tableaux like

$$\begin{array}{lllll} 0 & 0 & 0 & 0 & \alpha \\ 0 & 0 & 0 & \beta & \\ 0 & \gamma & \delta & & \\ \varepsilon & & & & \end{array} ;$$

note the striking resemblance with pictures as displayed in (8) (but the order of γ and δ is unchanged here). It is clear that they use a geometric criterion (γ should remain below α, δ below β, and ε below γ) rather than the lattice permutation condition; their method is therefore an efficient one, as we discussed following proposition 1.4.5. The authors go on to indicate explicitly how the resulting decomposition of the product can be verified by computing the dimension of the corresponding (reducible) representation of the symmetric group $\mathbf{S}_{13}$. Interestingly, while their set of tableaux is correct, they forget the contributions of three of them to the decomposition, whence the dimension of the printed result, 398541, falls short of the correctly predicted dimension 450450; this in spite of their claim that "this equation proves to be correct".

One more curious point concerns the Theorem II referred to in Theorem III. It describes the set of S-functions appearing in a product of two S-functions, without giving their multiplicities. This would seem to contradict our statement that no method is known to decide membership of that set, which does not amount to finding (or failing to find) an appropriate Littlewood-Richardson tableau. This is not so for two reasons. First, the criterion given is even more impractical than using the Littlewood-Richardson rule: it states (in a slightly different formulation) that occurrence of s_ν in $s_\lambda s_\mu$ is equivalent to the existence of a bijection $Y(\lambda * \mu) \to Y(\nu)$ mapping squares of any row to distinct columns, and squares of any column to distinct rows. Enumerating this set of bijections (or proving it empty) is certainly not easier than enumerating its subset of pictures, which is equivalent to enumerating $\mathrm{LR}(\nu/\mu, \lambda)$. Second, the criterion is wrong; for instance the bijection

$$\begin{array}{|c|c|c|c|} \cline{3-4} \multicolumn{1}{c}{} & \multicolumn{1}{c|}{} & a & b \\ \cline{3-4} \multicolumn{1}{c}{} & \multicolumn{1}{c|}{} & c & d \\ \hline e & f & \multicolumn{2}{c}{} \\ \cline{1-2} g & h & \multicolumn{2}{c}{} \\ \cline{1-2} \end{array} \quad \longrightarrow \quad \begin{array}{|c|c|c|c|} \hline a & b & e & h \\ \hline c & d & \multicolumn{2}{c}{} \\ \cline{1-2} f & \multicolumn{3}{c}{} \\ \cline{1-1} g & \multicolumn{3}{c}{} \\ \cline{1-1} \end{array}$$

satisfies the requirements, while $c^{(4,2,1,1)}_{(2,2),(2,2)} = 0$. The proof that is given in fact only estabishes the necessity of the condition, not its sufficiency. Fortunately, Littlewood and Richardson will be remembered more for a true Theorem they did not claim to prove, than for a false Theorem they did claim to prove.

4.2. *The paper by Robinson.*

In [Rob], Robinson builds forth on these ideas, claiming to complete the proof of the rule. The paper is quite difficult to read however, its formulations extremely obscure, and its argumentation mostly implicit; we shall now first try to summarise the argument as it appears to have been intended. In a deviation from the previous paper, the proof does not use the Jacobi-Trudi identity, but rather descending induction on the weight μ with respect to '$\leq$', based on the expression $s_\mu = h_\mu - \sum_{\mu'>\mu} K_{\mu'\mu} s_{\mu'}$, where $h_\mu = h_{\mu_0} h_{\mu_1} \cdots$, and the $K_{\mu'\mu}$ are non-negative integer coefficients (nowadays called Kostka numbers). By taking $\lambda = (0)$ in the formula of [LiRi] for the decomposition of products $s_\lambda h_\mu$, we see that $K_{\mu'\mu}$ is the number of semistandard Young tableaux of shape μ' and weight μ. For $\mu = (q,r)$ for instance, one easily finds $s_\mu = h_q h_r - s_{(q+1,r-1)} - \cdots - s_{(q+r,0)}$.

Robinson defines the correspondence $\mathcal{R}$ of (9), and uses it as follows to show that each $L \in \mathrm{LR}(\nu/\lambda, \mu)$ corresponds to an occurrence of s_ν in $s_\lambda s_\mu$. We already know that each $T \in \mathrm{SST}(\nu/\lambda)$ with $\mathrm{wt}\, T = \mu$ corresponds to an occurrence of s_ν in $s_\lambda h_\mu$. If $\mathcal{R}_0(T) \neq T$, then $\mathcal{R}_0(T)$ corresponds by inductive assumption to an occurrence X of s_ν in $s_\lambda s_{\mu'}$, where $\mu' = \mathrm{wt}\, \mathcal{R}_0(T) > \mu$. Viewing $s_{\mu'}$ as a constituent of h_μ, and thereby $s_\lambda s_{\mu'}$ as a part of $s_\lambda h_\mu$, we let T correspond to the occurrence of s_ν in $s_\lambda h_\mu$ matching X. There may be distinct $T, T' \in \mathrm{SST}(\nu/\lambda)$ of weight μ with $\mathcal{R}_0(T) = \mathcal{R}_0(T')$, but in that case the semistandard Young tableaux $\mathcal{R}_1(T)$ and $\mathcal{R}_1(T')$ of shape μ' and weight μ differ, and can be used to distinguish distinct occurrences of $s_{\mu'}$ in h_μ, and hence to distinguish occurrences of s_ν matching X in distinct contributions $s_\lambda s_{\mu'}$ to $s_\lambda h_\mu$. Indeed the number of such semistandard Young tableaux is precisely the multiplicity $K_{\mu'\mu}$ of $s_{\mu'}$ in h_μ. The tableaux $T \in \mathrm{SST}(\nu/\lambda)$ that remain, namely those with $T = \mathcal{R}_0(T)$ and therefore $T \in \mathrm{LR}(\nu/\lambda, \mu)$, must then correspond to the occurrences of s_ν in the remainder of $s_\lambda h_\mu$ after subtraction of all the $s_\lambda s_{\mu'}$, i.e., in $s_\lambda s_\mu$, as claimed.

In order to give the above argument a certain transparency, we have named the occurring multiplicities $K_{\mu'\mu}$, and related them to semistandard Young tableaux as they occur (implicitly) in [LiRi]. This is not however the way it is done in [Rob]. Rather than using the formula of Littlewood and Richardson for the decomposition of $h_\mu = s_{(0)} h_\mu$ into Schur functions, Robinson uses a formula by Young, given in the cryptic form $h_\alpha = \sum \bigl[\prod S_{rs}^{\lambda_{rs}}\bigr](\alpha)$, and accompanied by the following "explanation", quoted literally from [Young]:

> "S_{rs} where $r < s$ *represents the operation of moving one letter from the s-th row up to the r-th row, and the resulting term is regarded as zero, whenever any row becomes less than a row below it, or when letters from the same row overlap,—as, for instance, happens when $\alpha_1 = \alpha_2$ in the case of $S_{13}S_{23}$."*

This is what is nowadays called Young's rule, although it is usually stated in a somewhat different form. The meaning of this original formulation appears to be as follows. The summation is over certain collections $(\lambda_{rs})_{1\leq r<s\leq n}$ with $\lambda_{rs} \in \mathbf{N}$, each of which gives rise to a monomial $M = \prod S_{rs}^{\lambda_{rs}}$ in commuting indeterminates S_{rs}. An operation on Young diagrams is associated to such M, where each factor S_{rs} moves a square from row s up to row r, all factors acting simultaneously. An application of this operation is assumed implicitly to vanish if any row of the diagram contains too few squares (in the given summation with application to $Y(\alpha)$, only terms with $\sum_r \lambda_{rs} \leq \alpha_s$ need to be considered), and it also vanishes in the explicitly (if not very clearly) described cases. In the summation each remaining operation contributes the Schur function corresponding to the shape of the diagram produced by the application of the operation to $Y(\alpha)$.

In order to understand this formula, it helps to compare it to the decomposition formula in terms of semistandard Young tableaux: any $P \in \mathrm{SST}(\beta)$ with $\mathrm{wt}\, P = \alpha$ should correspond to a monomial M, whose operation transforms $Y(\alpha)$ into $Y(\beta)$. This can be achieved by taking $\lambda_{rs} = P^s_r$ for all r, s (cf. definition 1.4.1): then each entry of P records the row of $Y(\alpha)$ that its square came from (i.e., P

is obtained by applying M as square-moving operation to $\mathbf{1}_\alpha$, with each entry moving along with its square). It remains unclear how the quotation above can be interpreted as making all operations vanish that are not related in this manner to any tableau P. The given restrictions may be read as requiring β to be a partition, and requiring that squares from one row do not end up in the same column (assuming the ambiguity about the destination column of squares is resolved). But even with the most lenient interpretation, it is a mystery how in the result of application to $\mathbf{1}_\alpha$ the monotonicity of columns is enforced: for instance, for $\alpha = (1,1,1,1)$, the operation associated to $S_{12}S_{24}$ must be made to vanish, while the one for $S_{12}S_{23}S_{34}$ should survive (their applications to $\mathbf{1}_\alpha$ result in $\begin{array}{|c|c|}\hline 1 & 2 \\ \hline 4 \\ \cline{1-1} 3 \\ \cline{1-1}\end{array}$ and $\begin{array}{|c|c|}\hline 1 & 2 \\ \hline 3 \\ \cline{1-1} 4 \\ \cline{1-1}\end{array}$, respectively).

In any case, Robinson proceeds to define a process of transforming non-lattice permutations (in the form of words over an alphabet $\{c_1, c_2, \ldots\}$) into lattice permutations, which he calls "association I"; incidentally, he attributes the procedure to D. E. Littlewood. In our terminology it amounts to repeatedly applying the raising operation e_i with smallest possible index i (where the letter c_j is treated like the letter $j-1$ in our description), until no further application is possible. Applying it to the Semitic reading $w_\mathrm{S}(T)$ of a skew tableau T of weight α, and writing the resulting lattice permutation back into the shape of T, this association corresponds to $\mathcal{R}_0$ as described in corollary 3.3.4. Inspired by the fact that the monomials $\prod S_{rs}^{\lambda_{rs}}$ are described as "products of operations", Robinson associates such a monomial M to the sequence of operations e_i used in determining association I, by grouping together maximal sequences of successive operations of the form $e_{s-1}, e_{s-2}, \ldots, e_{r+1}, e_r$, replacing them by S_{rs}, and multiplying all of these. He calls the resulting correspondence between the original word and this monomial "association II". For instance, the sequence used in our figure 1 would be grouped as (e_0), (e_0), (e_1), (e_2, e_1), (e_2, e_1), (e_3, e_2), (e_3, e_2), (e_4), (e_4, e_3, e_2), and would give rise to the monomial $M = S_{01}^2 S_{12} S_{13}^2 S_{24}^2 S_{45} S_{25}$.

Robinson now states that any such M is one of the monomials that arises when Young's rule is applied for h_α (in our terms this means that M corresponds to a semistandard Young tableau P of weight α); this he justifies only by checking one example. We can check it for our example by applying M to $\mathbf{1}_\alpha$ (since $\alpha = (3,4,2,3,2,2)$ is not a partition, we must extend our definitions a bit, allowing intermediate "tableaux" whose shapes are not Young diagrams); alternatively we can find P directly by using $P_r^s = \lambda_{rs}$ for $r < s < 6$, and $P_s^s = \alpha_s - \sum_{r<s} P_r^s$:

$$\left(P_r^s\right)_{0 \le r \le s < 6} = \begin{pmatrix} 3 & 2 & 0 & 0 & 0 & 0 \\ & 2 & 1 & 2 & 0 & 0 \\ & & 1 & 0 & 2 & 1 \\ & & & 1 & 0 & 0 \\ & & & & 0 & 1 \\ & & & & & 0 \end{pmatrix}, \quad \text{whence} \quad P = \begin{array}{|c|c|c|c|c|}\hline 0 & 0 & 0 & 1 & 1 \\ \hline 1 & 1 & 2 & 3 & 3 \\ \hline 2 & 4 & 4 & 5 \\ \cline{1-4} 3 \\ \cline{1-1} 5 \\ \cline{1-1}\end{array}.$$

That this condition is always satisfied (i.e., that P is a semistandard Young tableau) is by no means clear however. If another criterion for selecting e_i were used (say taking i maximal), it would in fact be very hard to associate a proper monomial M to the sequence of operations e_i at all. If we admit that the construction always works, then association II defines a map that, applied to the word $w_\mathrm{S}(T)$, may serve as $\mathcal{R}_1$ (keeping in mind that Robinson works directly with monomials rather than with semistandard Young tableaux). It is easy to see that the shape β of M (i.e., the shape of the corresponding tableau P) matches the weight of the lattice permutation found by association I, but not that the map $\mathcal{R}$ defined by combining $\mathcal{R}_0$ and $\mathcal{R}_1$ is a bijection, as required in Robinson's proof sketched above. For this, it is of vital importance that the sequence of raising operations can be reconstructed from M, which is far from obvious since the S_{rs} commute, but the operations e_i do not. The key property is that in the sequence of factors S_{rs} found, the index s increases weakly, while for fixed s the index r decreases weakly (this is the hard part); once this is established the fact that M corresponds to a semistandard Young tableau can also be proved easily. This is not all however, since the surjectivity of $\mathcal{R}$ must also be established, i.e., that every monomial M of shape β arising in Young's rule applied to h_α is obtained for some word $w_\mathrm{S}(T)$ of weight α that corresponds under association I to a given lattice permutation of weight β.

For reading words $w_\mathrm{S}(P')$ of tableaux P' of *partition* shape β, the mentioned key property is not difficult to establish; moreover the lattice permutation is $w_\mathrm{S}(\mathbf{1}_\beta)$ in this case, while $P = P'$. Since we know (by theorem 3.3.1) that the coplactic graph of any word $w_S(T)$ is isomorphic to that of some such word $w_S(P')$ (with $T \triangleright P'$), we can see that the properties stated above do always hold. However, these facts are highly non-trivial given only the information provided in [Rob]. Nonetheless, Robinson apparently considers them to be obvious: nothing in the paper even suggests that anything needs to be proved. In conclusion, although Robinson gives an interesting construction that actually works, and that could be used in a proof of the Littlewood-Richardson rule, his argument contains such important gaps, that it definitely cannot be considered to provide such a proof.

4.3. *Later developments.*

The more recent history of the Littlewood-Richardson rule is no less interesting than the initial phase, but since it is much more accessible and better known, we shall limit ourselves to a brief overview. The flawed proof given by Robinson is so incomprehensibly formulated, that its omissions apparently go unnoticed for decades; his reasoning is reproduced in [Litw] by way of proof. In the early 1960's, in an unrelated study, Schensted gives a combinatorial construction [Sche] that will later be considered to be essentially equivalent to that of Robinson; this in spite of the fact that it defines a rather different kind of correspondence by an entirely different procedure. Schensted's construction clearly defines a bijection, but without any obvious relation to the Littlewood-Richardson rule (although it does involve tableaux); nonetheless it will later be central to several proofs of that rule. Initially however, no such connection is made, although the combinatorial significance of the construction is soon observed by Schützenberger [Schü1].

This changes in the 1970's, and important new properties of Schensted's construction are found: [Knu], [Gre]. From this development emerge the first proofs of the Littlewood-Richardson rule: Lascoux and Schützenberger introduce jeu de taquin, and using it a proof is given in [Schü3], while Thomas gives a proof that is based entirely on a detailed study of Schensted's construction in [Thom1], [Thom4]. The former proof is by means of a statement similar to our corollary 2.5.3, but it is obtained differently (e.g., for confluence of jeu de taquin, results of [Knu] and [Gre] are used). The latter proof is interesting in that it already derives properties of Schensted's correspondence that are related to pictures. Both approaches differ essentially from [LiRi] and [Rob], in that semistandard tableaux of shape λ and weight α figure in the same manner as in the present paper, identifying monomials X^α in s_λ rather than constituents s_λ of h_α. We also note that, whereas in [Rob] the basic construction is that of the component $\mathcal{R}_0$ of Robinson's bijection, with many questions remaining unanswered about $\mathcal{R}_1$, it is $\mathcal{R}_1$ that is central in [Schü3], and $\mathcal{R}_0$ does not even occur there. Two publications from this period do take up the construction of [Rob]. In [Thom2] the originally deterministic description of $\mathcal{R}_0$ is generalised to a rewrite system (cf. corollary 3.3.4), which is shown to be confluent. In [Macd, I (9.2)] the task of completing Robinson's argumentation (and presenting it in an understandable way) is taken up. A justification is provided for it, by establishing (in the course of several pages of detailed verifications) some crucial combinatorial properties of Robinson's construction; as we noted above, the need to prove such properties is not in any way mentioned in the original paper. Even so, Macdonald does not appear to give a convincing argument proving the surjectivity of $\mathcal{R}$, and some additional verifications seem to be needed.

After the appearance of these three proofs of the Littlewood-Richardson rule, many publications follow; some present interesting new ideas that lead to new proofs, but most of these are based on the same construction as one of these earlier proofs. Of particular interest is [Zel1], whose construction relates to all three approaches. It generalises Schensted's correspondence to pictures, which in essence consists in showing that its bijectivity is preserved when certain restrictions parametrised by shapes (like ν/κ-dominance) are imposed at both sides. This is exactly the information that is required in Schensted-based proofs (cf. [Thom4], [White], [ReWh]), to make the connection with Littlewood-Richardson tableaux. One also obtains as a special case a bijection $\mathcal{R}$ matching the specification (9), which provides a shortcut for the proof of [Macd] (what is not so obvious, is that this is in fact the same correspondence as defined by Robinson). Finally, this establishes a symmetry between $\mathcal{R}_0$ and $\mathcal{R}_1$ (which is what is most notably missing in the approach of [Schü3]), which allows not only $\mathcal{R}_1$ but also $\mathcal{R}_0$ to be defined by jeu de taquin. At the time however, these connections are not made, and the paper apparently does not get much attention (this may be due to its somewhat obscure definition of pictures, which bears no apparent relation to semistandard tableaux). The mentioned observations are made only quite a bit later, in [vLee2].

There are many more recent developments that relate to the Littlewood-Richardson rule; we shall mention a few here, but a detailed discussion is beyond the scope of this paper. Alternative combinatorial expressions for the coefficients $c^{\nu}_{\lambda,\mu}$ have been found (e.g., [BeZe]), as well as generalisations to representations of other groups than $\mathbf{GL}_n(\mathbf{C})$ ([Litm1]), and yet more new proofs. Of the latter we note the "involution style" proofs in [ReSh] and [Gash], which (apart from being particularly simple and concise) are remarkable by their similarity to the original proof for two-part partitions μ in [LiRi]: they use the Jacobi-Trudi identity to express the Schur function s_μ in terms of complete symmetric functions, and proceed to combinatorially cancel terms obtained after expansion of the determinant. The crucial difference (it seems) with that original proof is that the factors in the products h_α of complete symmetric functions are ordered in such a way that the weights α are "wide apart", and only dominant if $\alpha = \mu$: for $\mu = (q, r)$ one uses the expression $s_\mu = h_{(q,r)} - h_{(r-1,q+1)}$ rather than $s_\mu = h_{(q,r)} - h_{(q+1,r-1)}$.

We conclude by returning to the practical use of the rule as a computational tool, which was what it was formulated for in the first place, but which seems to have moved to the background in the course of time. Already in 1968, early in the computer era, the rule is implemented (including dimension checks, as there is no proof at that time!), and used to mechanically produce tables published in [Wyb]. The programme containing that implementation, originally written in FORTRAN, has evolved (with considerable extensions and several conversions to different programming language) into the current programme SCHUR. A more recent paper [ReWh] discusses theoretical aspects of computer implementation of the Littlewood-Richardson rule. Surprisingly, what it calls "a new combinatorial rule for expanding the product of Schur functions", is merely a translation of the problem of multiplying $s_\lambda * s_\mu$ into counting objects that are straightforward encodings of pictures $Y(\lambda * \mu) \to Y(\nu)$, for varying ν; as we have indicated, this differs only marginally from the process described in [LiRi]. Nonetheless, the formulation given is more straightforward to implement efficiently than most other formulations current at that time. While the paper does not specify a detailed algorithm, it has been used in concrete implementations (but we know of none that are available at the time of writing).

Currently, there are several freely available and efficient implementations of the Littlewood-Richardson rule, in various computer algebra systems. We mention an implementation by J. Stembridge, contained in the Maple package SF (`http://www.math.lsa.umich.edu/~jrs/maple.html`) and a similar implementation in ACE (`http://weyl.univ-mlv.fr/~ace/`). The stand-alone program LiE (`http://wwwmathlabo.univ-poitiers.fr/~maavl/LiE/`) contains an implementation, written by the current author; it is available for online use and for consultation of the documented source code [vLee3], via the mentioned WWW-page.

References.

[BeKn] E. A. Bender and D. E. Knuth, "Enumeration of plane partitions", *J. Combin. Theory, Ser. A* **13**, (1972), 40–54.

[BeSoSt] G. Benkart, F. Sottile, J. Stroomer, "Tableau switching: algorithms and applications", *J. Combin. Theory, Ser. A* **76**, (1996), 11–43.

[BeZe] A. D. Berenstein and A. V. Zelevinsky, "Triple multiplicities for $sl(r+1)$ and the spectrum of the exterior algebra of the adjoint representation", *J. of Algebraic Combinatorics* **1**, (1992), 7–22.

[BrTeKr] N. G. de Bruijn, C. van Ebbenhorst Tengbergen, D. Kruyswijk, "On the set of divisors of a number", *Nieuw Archief voor Wiskunde* **23**, (1951), 191–193.

[Eriks] K. Eriksson, *Strongly convergent games and Coxeter groups*, Thesis (KTH Stockholm), (1993).

[Fom] S. Fomin, *Knuth equivalence, Jeu de taquin, and the Littlewood-Richardson rule*, Appendix to: R. P. Stanley, *Enumerative Combinatorics, vol. 2*, Cambridge University Press (1999), pp. 413–439.

[FoGr] S. Fomin and C. Greene, "A Littlewood-Richardson miscellany", *Europ. J. of Combinatorics* **14**, (1993), 191–212.

[Fult] W. Fulton, *Young tableaux*, London Mathematical Society Student Texts 35, Cambridge University Press, 1997.

[Gash] V. Gasharov, "A short proof of the Littlewood-Richardson rule", *Europ. J. of Combinatorics* **19**, (1998), 451–453.

[Gre] C. Greene, "An extension of Schensted's theorem", *Advances in Math.* **14**, (1974), 254–265.

[GrKl] C. Greene and D. J. Kleitman, "Strong versions of Sperner's theorem", *J. Combin. Theory, Ser. A* **20**, (1976), 80–88.

[Haim1] M. D. Haiman, "On mixed insertion, symmetry, and shifted Young tableaux", *J. Combin. Theory, Ser. A* **50**, (1989), 196–225.

[Haim2] M. D. Haiman, "Dual equivalence with applications, including a conjecture of Proctor", *Discrete Math.* **99**, (1992), 79–113.

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

[Knu] D. E. Knuth, "Permutations, matrices and generalized Young tableaux", *Pacific J. of Math.* **34**, (1970), 709–727.

[LaSch] A. Lascoux and M. P. Schützenberger, "Le monoïde plaxique", *Quad. Ricerca Scientifica C.N.R.* **109**, (1981), 129–156.

[LeTh] B. Leclerc and J.-Y. Thibon, "The Robinson-Schensted correspondence, crystal bases, and the quantum straightening at $q = 0$", *Electronic J. of Combinatorics* **3 (no. 2)**, R11, (1996), 24 pp.

[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 Math.* **157**, (1996), 321–362.

[vLee3] M. A. A. van Leeuwen, *The Littlewood-Richardson rule, theory and implementation*, Report MAS-R9709, CWI, Amsterdam, (March 1997).

[vLee4] M. A. A. van Leeuwen, "An analogue of jeu de taquin for Littelmann's crystal paths", *Séminaire Lotharingien de Combinatoire* **B41b**, (1998), 23 pp.

[Litm1] P. Littelmann, "A Littlewood-Richardson rule for symmetrizable Kac-Moody algebras", *Inventiones mathematicae* **116**, (1994), 329–346.

[Litm2] P. Littelmann, "Paths and root operators in representation theory", *Annals of Math.* **142**, (1995), 499–525.

[Litw] D. E. Littlewood, *The theory of group characters, 2nd ed.*, Clarendon press, Oxford, 1950.

[LiRi] D. E. Littlewood and A. R. Richardson, "Group characters and algebra", *Phil. Trans. A* **233**, (1934), 99–141.

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

[ReSh] J. B. Remmel and M. Shimozono, "A simple proof of the Littlewood-Richardson rule and applications", *Discrete Math.* **193**, (1998), 257–266. (In "Selected papers in honor of Adriano Garsia" (Taormina, 1994)).

[ReWh] J. B. Remmel and R. Whitney, "Multiplying Schur functions", *J. of Algorithms* **5**, (1984), 471–487.

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

[Sche] C. Schensted, "Longest increasing and decreasing subsequences", *Canadian J. of Math.* **13**, (1961), 179–191.

[Schü1] M. P. Schützenberger, "Quelques remarques sur une construction de Schensted", *Math. Scandinavica* **12**, (1963), 117–128.

[Schü2] M. P. Schützenberger, "Promotion des morphismes d'ensembles ordonnés", *Discrete Math.* **2**, (1972), 73–94.

[Schü3] M. P. Schützenberger, "La correspondance de Robinson", pp. 59–113 in *Combinatoire et représentation du groupe symétrique* (D. Foata, ed.), Lecture Notes in Math. 579, 1977.

[Thom1] G. P. Thomas, *Baxter algebras and Schur functions*, Ph. D. thesis, University College of Swansea, (1974).

[Thom2] G. P. Thomas, "A generalisation of a construction due to Robinson", *Canadian J. of Math.* **28**, (1976), 665–672.

[Thom3] G. P. Thomas, "On a construction of Schützenberger", *Discrete Math.* **17**, (1977), 107–118.

[Thom4] G. P. Thomas, "On Schensted's construction and the multiplication of Schur functions", *Advances in Math.* **30**, (1978), 8–32.

[White] D. White, "Some connections between the Littlewood-Richardson rule and the construction of Schensted", *J. Combin. Theory, Ser. A* **30**, (1981), 237–247.

[Wyb] B. G. Wybourne, *Symmetry Principles and Atomic Spectroscopy*, J. Wiley and Sons, New York, 1970.

[Young] A. Young, "Quantitative substitutional analysis, Part IV", *Proceedings of the London Mathematical Society (2)* **31**, (1930), 253–272.

[Zel1] A. V. Zelevinsky, "A generalisation of the Littlewood-Richardson rule and the Robinson-Schensted-Knuth correspondence", *J. of Algebra* **69**, (1981), 82–94.

[Zel2] A. V. Zelevinsky, *Representations of finite classical groups, a Hopf algebra approach*, Lecture Notes in Math. 869, Springer-Verlag, Berlin/New York, 1981.