

Spin-preserving Knuth correspondences for ribbon tableaux

Marc A. A. van Leeuwen

Université de Poitiers, Département de Mathématiques,
UFR Sciences SP2MI, Téléport 2, BP 30179, 86962 Futuroscope Chasseneuil Cedex, France

`Marc.van-Leeuwen@math.univ-poitiers.fr`

`http://www-math.univ-poitiers.fr/~maavl/`

Submitted: Dec 1, 2003; Accepted: Jan 3, 2005; Published: Feb 14, 2005

Mathematics Subject Classifications: 05E10

Abstract

The RSK correspondence generalises the Robinson-Schensted correspondence by replacing permutation matrices by matrices with entries in \mathbf{N} , and standard Young tableaux by semistandard ones. For $r \in \mathbf{N}_{>0}$, the Robinson-Schensted correspondence can be trivially extended, using the r -quotient map, to one between r -coloured permutations and pairs of standard r -ribbon tableaux built on a fixed r -core (the Stanton-White correspondence). Viewing r -coloured permutations as matrices with entries in \mathbf{N}^r (the non-zero entries being unit vectors), this correspondence can also be generalised to arbitrary matrices with entries in \mathbf{N}^r and pairs of semistandard r -ribbon tableaux built on a fixed r -core; the generalisation is derived from the RSK correspondence, again using the r -quotient map. Shimozono and White recently defined a more interesting generalisation of the Robinson-Schensted correspondence to r -coloured permutations and standard r -ribbon tableaux; unlike the Stanton-White correspondence, it respects the spin statistic on standard r -ribbon tableaux, relating it directly to the colours of the r -coloured permutation. We define a construction establishing a bijective correspondence between general matrices with entries in \mathbf{N}^r and pairs of semistandard r -ribbon tableaux built on a fixed r -core, which respects the spin statistic on those tableaux in a similar manner, relating it directly to the matrix entries. We also define a similar generalisation of the asymmetric RSK correspondence, in which case the matrix entries are taken from $\{0, 1\}^r$.

More surprising than the existence of such a correspondence is the fact that these Knuth correspondences are not derived from Schensted correspondences by means of standardisation. That method does not work for general r -ribbon tableaux, since for $r \geq 3$, no r -ribbon Schensted insertion can preserve standardisations of horizontal strips. Instead, we use the analysis of Knuth correspondences by Fomin to focus on the correspondence at the level of a single matrix entry and one pair of ribbon strips, which we call a shape datum. We define such a shape datum by a non-trivial generalisation of the idea underlying the Shimozono-White correspondence, which takes the form of an algorithm traversing the edge sequences of the shapes

1 Introduction

involved. As a result of the particular way in which this traversal has to be set up, our construction directly generalises neither the Shimozono-White correspondence nor the RSK correspondence: it specialises to the transpose of the former, and to the variation of the latter called the Burge correspondence.

In terms of generating series, our shape datum proves a commutation relation between operators that add and remove horizontal r -ribbon strips; it is equivalent to a commutation relation for certain operators acting on a q -deformed Fock space, obtained by Kashiwara, Miwa and Stern. It implies the identity

$$\sum_{\lambda \geq_r(0)} G_\lambda^{(r)}(q^{\frac{1}{2}}, X) G_\lambda^{(r)}(q^{\frac{1}{2}}, Y) = \prod_{i,j \in \mathbf{N}} \prod_{k=0}^{r-1} \frac{1}{1 - q^k X_i Y_j};$$

where $G_\lambda^{(r)}(q^{\frac{1}{2}}, X) \in \mathbf{Z}[q^{\frac{1}{2}}][[X]]$ is the generating series by $q^{\text{spin}(P)} X^{\text{wt}(P)}$ of semi-standard r -ribbon tableaux P of shape λ ; the identity is a q -analogue of an r -fold Cauchy identity, since the series factors into a product of r Schur functions at $q^{\frac{1}{2}} = 1$. Our asymmetric correspondence similarly proves

$$\sum_{\lambda \geq_r(0)} G_\lambda^{(r)}(q^{\frac{1}{2}}, X) \check{G}_\lambda^{(r)}(q^{\frac{1}{2}}, Y) = \prod_{i,j \in \mathbf{N}} \prod_{k=0}^{r-1} (1 + q^k X_i Y_j).$$

with $\check{G}_\lambda^{(r)}(q^{\frac{1}{2}}, X)$ the generating series by $q^{\text{spin}^\dagger(P)} X^{\text{wt}(P)}$ of transpose semistandard r -ribbon tableaux P , where $\text{spin}^\dagger(P)$ denotes the spin as defined using the standardisation appropriate for such tableaux.

§1. Introduction.

The Robinson-Schensted correspondence has been generalised by various authors in many different ways. Fomin has even described general schemes that allow defining variants of this correspondence for any combinatorial structures that satisfy certain basic relations. We shall apply such a scheme to find what can be described as a Knuth correspondence extending the Schensted correspondence recently defined by Shimozono and White [ShW2]; it is based however on a quite novel kind of basic construction.

To indicate where our constructions fit in, we shall first need to review various earlier generalisations of the Schensted algorithm, and Fomin's general constructions. That however involves a rather long and abstract discussion, which does not transmit very well the flavour of the operations we shall actually be concerned with. Therefore we prefer, in order to whet the reader's appetite, to first present some enumerative problems that do not require much background, and yet are quite close to the questions related to our main construction; indeed the enumerative claims we formulate below will follow as special case from that construction. These problems allow us to introduce in an informal manner several important ideas behind our constructions. This discussion is for motivation only, so readers who wish to skip this somewhat oversized *hors d'œuvre* can move on to §1.2 without problem; the remainder of the paper will explicitly provide any required notions and results where and when those are needed.

1.1 Some enumerative problems

Although the polynomials $R_{1,1}(w, r)$ are rather laborious to compute by hand, their computation can be quite easily programmed. The basic observation is that after having prefixed w by r bits 1 (more are not necessary), each possible placement of a first r -ribbon is characterised by the simultaneous occurrence of a bit 1 and a bit 0 exactly r places to its right, and that the modification of the path due to placing the ribbon corresponds to changing those two bits and nothing else. The possibilities of adding further ribbons can be computed recursively if one takes care to ensure that they can only be placed further to the right. This can be achieved by removing in the recursive call the initial part of the word that may no longer be altered, i.e., the part up to and including the first bit that changed (from 1 to 0; since the bit disappears anyway there is no need to actually perform this change). To illustrate the simplicity of the algorithm, we present the complete code in the language of the MuPAD computer algebra system. We hope that this is readable even to those not familiar with MuPAD; comments are given after the symbol “//”. The only technical point is the procedure *R11* which prefixes r bits 1 to the word before entering the recursion, and makes sure the result is expanded and presented as a polynomial in $\mathbf{Z}[Y][X]$.

```

R := proc(w, r) // w is a list [w1, ..., wl] with wi ∈ {0, 1}
    local l, result, i, j, ww;
    begin l := nops(w); result := 1 // count the solution with no ribbons
    ; for i from 1 to l - r // when l ≤ r, the loop is skipped
        do if w[i] = 1 and w[i + r] = 0
            then ww := [op(w, i + 1..l)] // copy the sub-list [wi+1, ..., wl]
            ; ww[r] := 1 // change the bit that was copied from wi+r
            ; result := result + R(ww, r) * X * Y ↑ _plus(ww[j] $ j = 1..r - 1)
            // the exponent of Y is ∑j=1r-1 wwj = ∑j=i+1i+r-1 wj
            end_if
        end_for
    ; result // return the polynomial computed
    end_proc;

R11 := proc(w, r) begin poly(poly(R([1 $ r, op(w)], r), [X, Y]), [X]) end_proc;

```

Thus one may compute for $w = [0, 0, 1, 0, 1, 1, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0]$ that $R_{1,1}(w, 4)$ equals

$$\begin{aligned}
& 1 \\
& +X(2 + 2Y + 2Y^2 + Y^3) \\
& +X^2(1 + 4Y + 7Y^2 + 5Y^3 + 5Y^4 + 2Y^5 + Y^6) \\
& +X^3(Y + 6Y^2 + 10Y^3 + 15Y^4 + 12Y^5 + 8Y^6 + 5Y^7 + 2Y^8 + Y^9) \\
& +X^4(2Y^3 + 11Y^4 + 19Y^5 + 23Y^6 + 20Y^7 + 16Y^8 + 8Y^9 + 5Y^{10} + 2Y^{11} + Y^{12}) \\
& +X^5(Y^5 + 10Y^6 + 21Y^7 + 32Y^8 + 29Y^9 + 24Y^{10} + 16Y^{11} + 8Y^{12} + 5Y^{13} + 2Y^{14} + Y^{15}) \\
& +X^6(3Y^8 + 12Y^9 + 28Y^{10} + 34Y^{11} + 33Y^{12} + 24Y^{13} + 16Y^{14} + 8Y^{15} + 5Y^{16} + 2Y^{17} + Y^{18}) \\
& +X^7(Y^{11} + 10Y^{12} + 21Y^{13} + 32Y^{14} + 29Y^{15} + 24Y^{16} + 16Y^{17} + 8Y^{18} + 5Y^{19} + 2Y^{20} + Y^{21}) \\
& +X^8(2Y^{15} + 11Y^{16} + 19Y^{17} + 23Y^{18} + 20Y^{19} + 16Y^{20} + 8Y^{21} + 5Y^{22} + 2Y^{23} + Y^{24}) \\
& +X^9(Y^{19} + 6Y^{20} + 10Y^{21} + 15Y^{22} + 12Y^{23} + 8Y^{24} + 5Y^{25} + 2Y^{26} + Y^{27}) \\
& +X^{10}(Y^{24} + 4Y^{25} + 7Y^{26} + 5Y^{27} + 5Y^{28} + 2Y^{29} + Y^{30}) \\
& +X^{11}(2Y^{30} + 2Y^{31} + 2Y^{32} + Y^{33}) \\
& +X^{12}Y^{36},
\end{aligned}$$

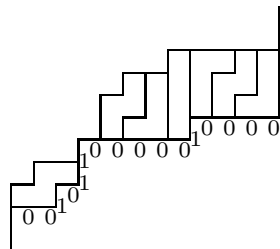
and so does $R_{1,1}(\tilde{w}, 4)$, where $\tilde{w} = [0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1, 1, 0, 1, 0, 0]$.

Our claim above can be interpreted in a geometric fashion. If a placement of r -ribbons below the path corresponding to \tilde{w} is rotated a half turn, one obtains a placement of r -ribbons *above* the path corresponding to w according to similar rules as for the placement below (but note that the order of placement is now from right to left). So the claim can be reformulated as: for any path that is ultimately vertical at both ends, and any specified number of ribbons and total height, there are as many placements possible above the path as there are below the path. Even more than the original formulation this form begs for a bijective proof, a simple rule to move ribbons to the other side of the path, so as to define an invertible map from placements of ribbons on one side to those on the other, preserving the number of ribbons and the total height; one would expect the inverse map to be given by the same rule after rotating the configuration a half turn. Yet we have not been able to find such a rule. Our proof of the claim (given at the end of the paper) will be based on a bijection, but one corresponding an identity obtained by multiplying both polynomials by an identical power series. Although a general method exists to deduce from this a bijection corresponding directly to the claim, the result it is way too complicated to qualify as a bijective proof.

This negative finding notwithstanding, there are quite a few observations we can make about this problem that do involve simple bijective constructions. For instance, when seeing polynomials $R_{1,1}(w, r)$ such as the one displayed above, it is hard to miss a symmetry in the coefficients (although we admit having done so for quite some time): for every i , the polynomial in Y that is the coefficient of X^{12-i} equals the coefficient of X^i multiplied by Y^{36-6i} . In the general case the coefficient of X^{k-i} equals that of X^i multiplied by $Y^{(r-1)(k-2i)}$, where k is the number of bits 0 in w . This suggests that each placement of i r -ribbons should correspond to a placement of $k - i$ such ribbons with the same total *width* (i.e., the number of horizontal adjacencies of squares within

1.1 Some enumerative problems

a same ribbon). Indeed it does, but the latter will be a placement of ribbons *above* the path of w , making the observation about the symmetry of the coefficients equivalent to our claim 1.1.1. To describe the correspondence, first observe that a placement of r -ribbons is completely determined by its top left and bottom right borders, i.e., by the path corresponding to w and that path as modified by the placement of all ribbons. Then if one shifts up the latter path by r units, it will be the top left border of a unique placement of r -ribbons above the path of w . For instance, here is the placement obtained from the one displayed earlier (both have total width 7):



Providing the details to prove that this correspondence is well defined and has the required properties is an instructive exercise that we encourage the reader to solve; those who would like a hint can turn to lemma 5.2.2 below, which provides a closely related result.

We can also give bijections proving several sub-cases of our claim 1.1.1. To begin with, the case $r = 1$ does not pose much difficulty. Now 1-ribbons are just single squares, and the total height, being zero always, plays no role. The placement rule forces squares to be added from left to right below the path of w , advancing at least one column at each step, so that no column can receive more than one square. Conversely any path that remains weakly below that of w and weakly above the same path shifted one unit down (thus leaving room for at most one square in each column) can be obtained for an appropriate placement of squares; the set of squares of such a placement is known as a horizontal strip. Each horizontal component of the path of w (a maximal portion without vertical steps) can be treated in isolation, and can be used to place any number of squares from 0 up to and including its length, in a unique way; this is true both for placement above and below it. So one can bring the horizontal strips above and below the path into bijective correspondence, by requiring that the former has as many squares directly above any horizontal component as the latter has directly below it, in other words that the former has as many squares in any given row as the latter has in the next row. There is an equivalent algorithmic description that treats the squares one at a time: traversing the squares of the horizontal strip above the path of w from left to right, each square is moved one place down, thus crossing the path of w , and then if necessary slid to the left until it finds a place where it can stay, namely where its left edge is part of the path of w as possibly modified by the previous placement of squares below it.

This case can be extended to cover a small part of the claim for general r , namely the part concerning the leading terms (in Y) of the coefficients of the X^i , in other words the placements where all r -ribbons are completely vertical. For such placements only

vertical portions of the path of length at least r effectively produce separate compartments where the ribbons can be placed independently; shorter vertical portions within such a compartment have no effect on the number of vertical ribbons that can be placed in it. So again each compartment can accommodate any number of vertical ribbons from 0 up to and including its width, and does so both above and below. This correspondence too can be described by processing the ribbons above the path from left to right: each one is moved down across the path, and then to the bottom left up to the first place where it will fit. Incidentally a similar procedure also works when there are only horizontal ribbons, but these cases are even more marginal than those involving only vertical ribbons, since generally only relatively few purely horizontal r -ribbons can be placed at all.

All this only scratches the surface of the general problem. It should be noted that one cannot expect a correspondence for the general case where each ribbon above the path gives rise to a ribbon of the same height below the path, for the simple reason that the distributions of heights within the collections of placements that should match are not always the same. This can be seen for the example given (for the given w , there are 5 different placements of five ribbons that, like the one displayed, produce a total height 8 as sum of the multiset $\{0, 1, 2, 2, 3\}$, whereas there are only 4 such placements for \tilde{w}), but a smaller example is more convincing: below the path of 01000 one can place a vertical 3-ribbon and a horizontal one (i.e., 3-ribbons of heights 2 and 0, respectively), but such a pair of 3-ribbons cannot be placed above the same path. This means that it is important that we count only by *total* height, and that any correspondence one would hope to find must have some mechanism for the exchange of height between ribbons (or alternatively it might treat placements of ribbons as a whole without even considering individual ribbons, as our first bijection did).

There are two more bijective correspondences that are worth mentioning in this context, as they provide interesting new points of view, even though they do not tackle the difficulties just indicated. The first proves the specialisation for $Y := 1$ of our claim, i.e., it treats all configurations but ignores the heights of the ribbons. The second handles the equality of the coefficients of X^1 , i.e., it treats configurations consisting of a single ribbon. If one ignores heights, matters become simpler if one forgets the geometric description, and views placement of ribbons simply as operations on bit strings. As we saw, the question of whether an r -ribbon can be placed, and the effect of placing it, can both be expressed in terms of just one pair of bits, at indices i and $i + r$. So placement of different r -ribbons becomes completely independent unless the indices i, i' of the bits involved are congruent modulo r (in the latter case we shall say the ribbons are in the same position class). Thus the possibilities for placing r -ribbons decompose completely following the r different position classes, and the specialisation for $Y := 1$ of $R_{1,1}(w, r)$ decomposes as a product $\prod_{i=0}^{r-1} R_{1,1}(w^{(i)}, 1)$ of polynomials in X , where $w^{(i)}$ is the word extracted from w , of its bits at indices congruent to i modulo r . So for instance for our example, the specialisation $1 + 7X + 25X^2 + 60X^3 + 107X^4 + 149X^5 + 166X^6 + 149X^7 + 107X^8 + 60X^9 + 25X^{10} + 7X^{11} + X^{12}$ factors as $R_{1,1}(0100, 1)R_{1,1}(0100, 1)R_{1,1}(1000, 1)R_{1,1}(0010, 1) = (1 + 2X + 2X^2 + X^3) \times (1 + 2X + 2X^2 + X^3)(1 + X + X^2 + X^3)(1 + 2X + 2X^2 + X^3)$. Thus we are back at

1.1 Some enumerative problems

the case $r = 1$ we know how to handle. We find the following procedure to transform a placement of r -ribbons above the path of w into one below, defined by the final value of a modified copy w' of w . Process the ribbons of any position class from left to right; the relative ordering between members of different classes is irrelevant. For a ribbon with initial bit $w_i = 0$, search in the current value of w' (which still has $w'_i = 0$), testing the bits $w'_{i-r}, w'_{i-2r}, w'_{i-3r}, \dots$ until finding the first bit $w'_{i-kr} = 1$; one has $w'_{i-kr+r} = 0$, and the bit string w' is modified by setting $w'_{i-kr} := 0$ and $w'_{i-kr+r} := 1$. The modifications to w' do not always occur in the right order to describe the ribbons of the placement eventually found, so the independence of operations on different position classes of ribbons is crucial for proving that the same procedure rotated a half turn defines an inverse.

We have seen that preserving heights of individual ribbons is not possible in general, but that ignoring heights altogether makes our problem trivial. The following bijection for the case of single ribbons gives some insight in the role played by height, without the complications of interaction between ribbons; it is based on observations in [ShW2]. When an r -ribbon of height h can be placed below the path of w with initial bit $w_i = 1$, this means that $w_{i+r} = 0$ and $\sum_{j=i+1}^{i+r-1} w_j = h$, which can also be formulated as $\sum_{j=i}^{i+r-1} w_j = h + 1$ and $\sum_{j=i+1}^{i+r} w_j = h$. Thus the places where a ribbon of height h can be placed below the path of w correspond exactly to the places where the sum of r consecutive bits drops from $h + 1$ to h , and similarly a ribbon of height h can be placed above the path precisely in the places where the value of such sums rises from h to $h + 1$. Therefore, starting from a place where such a ribbon can be placed below, one can always find a place further to the top right where a similar ribbon can be placed above (since the path ultimately becomes vertical), and from the first such place, the point of departure can be found back as the first place to its bottom left that will accommodate an r -ribbon of height h below the path. This establishes our bijection for the case of single ribbons. One may visualise all possibilities to place r -ribbons of height h , both below and above, as the points of crossing between the path of w and an appropriately shifted copy of the same path; see the illustrations after corollary 4.5.2 below.

We close our discussion of this problem with an indication of why we think it has no simple bijective solution (although we would love to be proved wrong). When one tries to extend the height preserving procedure for single ribbons to multiple ribbons, the main difficulty is not so much the exchange of heights that may be necessary, as the fact that the left to right order among ribbons cannot be preserved. We believe we could describe a bijection for the case of two ribbons, but it already gets horribly complicated: when the second ribbon placed needs to move beyond the place where the first landed, exchange of height must be taken into account, and it may be necessary to relocate the first ribbon. But the hardest part is to show that one gets a bijection: the original ribbons must be reconstructed from the pair produced without knowing in which order those were placed, so there is no question of a step-by-step inverse; a proof would involve piecing together all the scenarios that can arise. Unless there is some easy way to read off the order in which ribbons have been placed, it is hard to envisage a similar technique handling the case of three or more ribbons.

If we have spent much time on a problem for which we know no solution, it is because it is superficially simpler than a second problem, a variant of the first, but one for which we do have a solution; indeed the solution is closely related to the main result of this paper (and it will not be detailed in this introduction). The variant is simply obtained by extending the path described by the finite bit string w not vertically, but horizontally at both ends; in other words that string is now considered to float in a sea of bits 0. The conditions for placing collections of r -ribbons remain exactly the same, as are the two statistics on placements of ribbons (number of ribbons n and total height t); analogously to the definition of $R_{1,1}(w, r)$, the sum of $X^n Y^t$ over all possible placements below this differently extended path will now be denoted by $R_{0,0}(w, r)$. One still has symmetry between placements of ribbons above and below the path.

1.1.2. Claim. *If \tilde{w} is the word obtained by reversing the bits of w , then $R_{0,0}(w, r) = R_{0,0}(\tilde{w}, r)$.*

If, in keeping with the laws of gravity, we think primarily of placing ribbons above the path, then the path in our first ribbon placement problem resembles a ledge in an otherwise sheer rock-face, while the second problem more resembles a Dutch landscape with a polder to the left, a dike described by the string w , and the sea to the right (the sea being as high as the dike is not quite realistic, fortunately). We shall therefore refer to first ribbon placement problem as the alpine problem, and to this second ribbon placement problem as its polder variant.

The change of landscape modifies the character of our problem in several ways. While ribbons can lean against the rock face, the sea and the space above sea level are inaccessible (the top-rightmost vertical edge of each ribbon must belong either to the dike or to another ribbon). On the other hand, the requirement that the bottom-leftmost horizontal edge of each ribbon lie on the original path does not put a bound on the number of ribbons, since the polder provides an infinite supply of such edges. Indeed, provided w has at least one bit 1, arbitrarily many ribbons can be placed above the path, for instance using only horizontal ribbons in the polder. Hence the identity of our second claim is one of formal power series rather than of polynomials. Considering X to be the principal indeterminate, one has in fact $R_{0,0}(w, r) \in \mathbf{Z}[Y][[X]]$: the coefficient of X^n is a polynomial in Y of degree at most $n(r-1)$.

One cannot compute a complete power series $R_{0,0}(w, r)$, but the recursive procedure R above can be easily adapted to produce an initial part of such power series (in finite time!), by adding as a parameter a limit to the degree in X of the terms it should compute. Thus one verifies that for $w = 1011000001$ (our earlier example without the now superfluous bits 0 at either end), both $R_{0,0}(w, 4)$ and $R_{0,0}(\tilde{w}, 4)$ start as

1.1 Some enumerative problems

$$\begin{aligned}
 & 1 \\
 & +X(2 + Y + Y^2) \\
 & +X^2(2 + 2Y + 4Y^2 + Y^3 + Y^4) \\
 & +X^3(2 + 2Y + 5Y^2 + 4Y^3 + 4Y^4 + 2Y^5 + Y^6) \\
 & +X^4(2 + 2Y + 5Y^2 + 5Y^3 + 7Y^4 + 5Y^5 + 5Y^6 + 2Y^7 + 2Y^8) \\
 & +X^5(2 + 2Y + 5Y^2 + 5Y^3 + 8Y^4 + 8Y^5 + 8Y^6 + 6Y^7 + 6Y^8 + 3Y^9 + 2Y^{10} + Y^{11}) \\
 & \dots
 \end{aligned}$$

Like before, the configurations counted by $R_{0,0}(w, r)$ and by $R_{0,0}(\tilde{w}, r)$ can be interpreted respectively as placements of ribbons below and above the same path, and one would like to prove the claim by means of a bijection between such placements, which preserves the number of ribbons and their total height.

It is interesting to observe how the case $r = 1$, that of the horizontal strips, has changed. In terms of horizontal components of the path, we have effectively gained one such component, with infinite capacity, whether placing squares above or below the path (in more formal terms: assuming that w neither begins nor ends with a bit 0, one has $R_{0,0}(w, 1) = R_{1,1}(w, 1) \left(\sum_{n \in \mathbf{N}} X^n \right)$, which combinatorialists would write as $\frac{R_{1,1}(w, 1)}{1-X}$). But the infinite horizontal component is not the same one in both cases, so if one wants to maintain the bijection based on this decomposition into horizontal components, one has to decree that, while most squares from the strip above the path descend below it and then shift to the left, those that were just above polder level “wrap around at infinity” and come back from the extreme right, just below sea level. There is nothing against that as a bijection for $r = 1$, but as point of departure for the general case it is better to consider a different bijection, one that moves all squares in the same direction; this must be left-to-right when going from a horizontal strip above the path to one below. Doing so, squares arrive under a different horizontal component than the one they belonged to, and since the capacities of those are unrelated, the level at which squares will be placed cannot be as neatly predictable as before.

Yet there is a simple method for placing the squares, which is essentially to take the first available place to the right of their original position, taking into account the other squares. For instance, for a horizontal strip above the path that consists just of n squares on the lowest possible (polder) level, and just to the left of the first vertical edge of the path, the corresponding horizontal strip below the path will occupy the first n columns to the right of that vertical edge, at whatever level is necessary to be just below the path. This is possible for any $n \in \mathbf{N}$ because the squares below any horizontal component of the path can be filled up from left to right, and a horizontal component of infinite capacity is available at the end to absorb whatever number of squares might not fit elsewhere. One wants the same description to define the inverse procedure, which means in this example that the horizontal strip above the path that occupies the n columns directly to the left of the *last* vertical edge of the path should correspond to the strip of n squares to the right of that edge, just below sea level. To obtain that result, one must declare columns that contain a square of the original horizontal strip to be unavailable for placing squares, even if doing so could produce a horizontal strip as

output. Therefore the rule should be: for each square, taken in order from left to right of their original position, move it to the first column to its right that contains neither a square that was already moved nor one that has yet to be moved. One may verify that a square may indeed be placed in that column, just below the path, and that the same rule rotated a half turn will bring back each square to its original position.

Apart from this one-square-at-a-time procedure, there is a description of the same correspondence that treats all squares at the same time. Imagine a bus driving along the path from left to right, taking the squares with it as passengers. Each horizontal edge is a bus stop, where either a square enters the bus (from above), or a square leaves the bus (from below, but never at a stop where a square entered), or finally, in case the bus is empty and there is no square waiting at the bus stop, the bus just drives on. The importance of this alternative description is not so much its cuteness or greater efficiency, as the fact that it treats the squares without regard to their individual identities: while for the purpose of showing equivalence with the earlier description, one may imagine that each square leaving the bus is the one among the current passengers that has been aboard the longest, this “choice” has no effect on the result, and all that matters is keeping track of the number of passengers at each moment.

If we consider the other correspondences that were established for the alpine problem, we see that the one that was introduced in connection with the symmetry in the coefficients of the polynomial $R_{1,1}(w, r)$ has no counterpart in the polder variant, because the power series $R_{0,0}(w, r)$ has no such symmetry; the two other correspondences however (the one counting placements of ribbons ignoring their heights, and the correspondence for placements of single ribbons) can be adapted to the polder variant without much difficulty. The first one of these involves the same reduction to the case $r = 1$ as before, which case is modified as just discussed; the resulting correspondence can be described by transportation of ribbon by means of a bus with r separate compartments, one for each of the position classes. The second one is adapted by inverting the direction of search for a ribbon of appropriate height, due to fact that sums of r consecutive bits now ultimately become 0 in both directions, rather than r .

Unlike for the alpine problem, a bijection handling the general case and preserving height can be given here; this fact is at the heart of our main result. Details about the bijection will be given later, but here is a hint for the impatient. The correspondence is obtained by the passage of an r -deck bus transporting ribbons, but instead of segregation by class (an idea we could not endorse anyway) the level of entry and exit of ribbons is related to height. This level is not identical to the height of the ribbon entering or leaving however (that would force the output to have the same distribution of heights as the input, which, like for the alpine problem, cannot always be achieved). Rather it is the height a ribbon at that position would have in a third placement, one that extends at both sides of the path of w , and occupies the union of the areas occupied by the input and output placement (even if the latter is still under construction). Moreover, the bus operates in a stack-like fashion: whenever a ribbon enters or leaves at some level, all higher levels are empty.

Neither the alpine problem nor its polder variant quite reflect the original problem that motivated our work. The problem deals with Young diagrams, whose boundary

1.2 Background

is given by a path that starts vertically and ends horizontally. By extending w to the left with bits 1 and to the right with bits 0 one obtains such a path, and it can be used to define a generating series $R_{1,0}(w, r)$. For the purpose of counting placements of ribbons above rather than below such a path one similarly defines $R_{0,1}(\tilde{w}, r)$. The symmetry observed for the other problems does not exist for this case however; indeed $R_{1,0}(w, r)$ is a proper power series, while $R_{0,1}(\tilde{w}, r)$ is a polynomial. The simplest case is obtained for the empty word ϵ . Now no ribbons can be placed at all above the path, so that $R_{0,1}(\epsilon, r) = 1$. On the other hand, it can be seen that for any given multiset of heights, there is exactly one way to place ribbons of those heights below the path, which does so by weakly decreasing order of height. Thus one deduces that $R_{1,0}(\epsilon, r) = \prod_{h=0}^{r-1} \frac{1}{1-XY^h}$, the generating series of multisets on $[r] = \{0, \dots, r-1\}$. Instead of the equalities expressed in our first two claims, we observe here that the generating series for placements below the path is always obtained from the generating polynomial for those above by multiplication by this fixed power series $R_{1,0}(\epsilon, r)$.

1.1.3. Claim. *If \tilde{w} is the reverse word of w , then $R_{1,0}(w, r) = R_{0,1}(\tilde{w}, r) \prod_{h=0}^{r-1} \frac{1}{1-XY^h}$.*

In spite of the different nature of the statement, a bijective proof of this final claim will turn out to be deduced immediately from one for claim 1.1.2: each multiset contributing to the factor $\prod_{h=0}^{r-1} \frac{1}{1-XY^h}$ can be interpreted as describing an initial state of the bus when it arrives (instead of all decks being empty initially), the occurrences of $i \in [r]$ occupying deck i . The bus will still leave the scene empty in this case, but that too changes if the path ends vertically as in the alpine problem; by allowing for a non-empty bus both at entry to and at exit from the scene, one obtains a bijective proof not of the identity of claim 1.1.1, but of the identity derived from it by multiplying both sides by the power series $R_{1,0}(\epsilon, r)$.

1.2. Background.

The basic form of the Schensted algorithm constructs a bijection between permutations of n and pairs of standard Young tableaux of equal shape and size n . The two tableaux shall be referred to as the P -symbol and the Q -symbol, and this terminology will be extended to all generalisations of the algorithm considered. Its first generalisation already appears in the original paper [Sche]; it operates on arbitrary sequences of n numbers (with equal entries allowed), and it constructs as P -symbol a semistandard (or column-strict) tableau with the same multiset of entries as the word, while the Q -symbol remains a standard tableau. The symmetry that is lost here is restored in a further generalisation by Knuth [Knu], which operates on matrices with natural numbers as entries, and produces pairs of tableaux which are both semistandard, the multiplicities of their entries being given by the column sums (for the P -symbol) and the row sums (for the Q -symbol) of the matrix. The basic Robinson-Schensted correspondence is recovered when all row and column sums are equal to 1 (the case of permutation matrices); the generalisation given by Schensted corresponds to the case where the row sums are 1 but columns sums are arbitrary.

While this generalises the correspondence considerably, the algorithm itself changes only marginally. The case of a matrix with multiple entries in the same row or column, or

entries exceeding 1, is handled by operating in the same way the basic algorithm would for a permutation matrix derived from it by splitting up rows (each following non-zero entry getting a fresh row below that of the previous one), multiplexing columns similarly, and replacing entries $m > 1$ by $m \times m$ identity sub-matrices. The Schur function s_λ is the generating series of the semistandard tableaux of shape λ , so Knuth's correspondence provides a bijective proof of the Cauchy identity $\prod_{i,j} \frac{1}{1-X_i Y_j} = \sum_\lambda s_\lambda(X) s_\lambda(Y)$. Knuth also defines a second correspondence that provides a bijective proof of a "dual" identity $\prod_{i,j} (1 + X_i Y_j) = \sum_\lambda s_\lambda(X) s_{\lambda^t}(Y)$. The truncation that has occurred here of the factors of the left hand side to their terms of degree ≤ 1 , means that matrix entries are now restricted to the values $\{0, 1\}$; the transposition of λ in the second factor on the right means either that P - and Q -symbols have transpose shapes, or (since we prefer pairs of equal shape) that one is semistandard and the other transpose-semistandard (row-strict). Like the first one, this second correspondence can be constructed by first transforming the given matrix to a permutation matrix and then applying the Schensted algorithm; the only difference is that columns are multiplexed in the opposite sense: each next non-zero entry gets a fresh column to the *left* of the previous ones. In the first correspondence the P - and Q -symbols have symmetric roles, and transposition of the matrix leads to exchanging them. The second correspondence lacks such a symmetry, and we shall refer to it as Knuth's asymmetric correspondence; when not explicitly calling a correspondence asymmetric, it will be assumed to be a symmetric one.

Fomin has shown in a series of papers [Fom1]–[Fom5] that by identifying the various tableaux with paths in a graded partially ordered set or in a directed graph, these correspondences and many other ones can be described in a general framework that links local correspondences in the poset or graph to global correspondences involving pairs of paths. This also means that new correspondences can automatically be defined as soon as a poset or graph with the required local structure is found. As a consequence of the generality of these constructions, the terms "Schensted correspondence" and "Knuth correspondence" now acquire a generic meaning, and the specific correspondences mentioned above will be referred to as the Robinson-Schensted correspondence and the RSK correspondence (use this term only for the symmetric one). Our construction will be an instance of this general framework, so we shall recall the necessary parts of Fomin's work in detail in section 2; here we just sketch the outlines. Although the nature of the poset elements (or vertices of the graph) is not specified, we shall refer to them as "shapes"; they are Young diagrams in the cases of the Robinson-Schensted and RSK correspondences.

The general constructions build a two-dimensional array of shapes, from which the P -symbol and Q -symbol can be read off in the two different directions. To have an analogue of the Robinson-Schensted correspondence one needs a graded poset whose connected components are " r -differential". The most important requirement for this is that for all shapes μ the number of shapes covering μ exceeds the number of shapes covered by μ by a fixed number $r > 0$ (the precise requirement is that the commutator of the "up" and "down" operators for the poset be r times the identity operator). The Young lattice \mathbf{Y} , consisting of Young diagrams ordered by inclusion, is well known to be 1-differential. The r -differential property can be "made bijective" by means of an

1.2 Background

r -correspondence, which defines for every shape μ a bijection between on one side the set of shapes covering μ , and on the other side the union of the set of shapes covered by μ and a set of r extra values. Given such an r -correspondence, Fomin's construction will produce a "Schensted correspondence" between r -coloured permutations of n (where each term has an additional attribute with r possible values, whence their number is $n!r^n$) and pairs of saturated increasing paths of length n in the poset with a fixed minimal element as starting point and a common (but varying) end point. For the Young lattice such paths correspond to standard Young tableaux. For that case there are two natural choices for a 1-correspondence, one of which leads to the usual Schensted correspondence by row insertion, the other to its transposed variant (using column insertion).

For "Knuth correspondences" the general scheme, which is described in [Fom5], is more complicated. The graded set of shapes used has more than a poset structure: it is equipped with a directed graph, where edges may relate shapes any number of levels apart. The entries of the matrices that form the input of the construction come from a graded but otherwise unstructured set S . For the Young lattice there is edge from μ to λ whenever μ/λ is a horizontal strip (so that paths correspond to semistandard tableaux), while $S = \mathbf{N}$. The notion that replaces that of an r -correspondence is what we shall call a shape datum; it gives for every pair (μ, ν) of shapes a bijection from shapes κ with edges toward both μ and ν , to pairs (a, λ) , where λ is a shape with edges from both μ and ν , and $a \in S$. It must satisfy a compatibility with the gradings. This implies that when restricted to edges between shapes at most one level apart, it reduces to an r -correspondence, where r is the size of the rank 1 subset of S .

The shape datum that matches Knuth's original construction can be found by considering how that construction deals with a single matrix entry and one horizontal strip each of the P and Q symbols. Since such a strip is treated just like a skew standard tableaux corresponding to it (by way of "standardisation"), the mentioned shape datum is defined by a localised case of the original Schensted correspondence. Shape data in general however need not be derived from any Schensted correspondence.

The lattice \mathbf{Y}^r is an example of an r -differential poset with $r > 1$, and an r -correspondence for \mathbf{Y}^r can be defined by fixing 1-correspondences on each of its r factors. Similarly a graph on \mathbf{Y}^r , and a shape datum for it with graded set $S = \mathbf{N}^r$, can be derived in a component-wise fashion from the graph on \mathbf{Y} defined by horizontal strips. These structures lead to Schensted- and Knuth correspondences that factor into r independent copies of the original ones, which is not very interesting. Another example is given by the r -rim hook lattices, whose elements are (single) Young diagrams, but whose covering relation, adding an r -ribbon, relates shapes that are r levels apart in \mathbf{Y} ; the construction of [StWh] provides a Schensted correspondence this example. However, these lattices are isomorphic to \mathbf{Y}^r by means of the so-called r -quotient map, and the Stanton-White correspondence thus reduces to the example above.

Yet rim hook lattices can also give rise to Schensted correspondences that do not decompose into independent ones, by choosing an r -correspondence not derived from the r -quotient map. Indeed, in [ShW2] a more interesting r -correspondence for the r -rim hook lattices is defined, which unlike the previous one takes the shapes of the

ribbons into account. To edges in an r -rim hook lattice one may assign a value $h \in \{0, \dots, r-1\}$, namely the height of the associated r -ribbon; with the r extra values occurring in an r -correspondence also labelled with that set of values, Shimozono and White define an r -correspondence preserving these heights. As a consequence, the Schensted correspondence obtained respects the “spin” statistic on standard r -ribbon tableaux that gives half the sum of the heights of the ribbons: the sum of the colours of the input permutation determines the sum of the spins of the output tableaux.

In this paper we define a height respecting shape datum for r -rim hook lattices equipped with the graphs defined by the notion of horizontal r -ribbon strip (which is essentially a placement of ribbons of the previous subsection; this notion underlies that of semistandard r -ribbon tableaux), with graded set $S = \mathbf{N}^r$. The Knuth correspondence derived from it shares the “colour-to-spin” property of the Schensted correspondence of [ShW2], for the natural definitions of the respective statistics on matrices and semistandard r -ribbon tableaux. Given the way the original Knuth correspondence is derived from the Robinson-Schensted correspondence, it may seem a straightforward process to obtain such a shape datum from the r -correspondence of Shimozono and White; in any case, this is what we thought initially. It is not. In fact we were unable to find a shape datum that would reduce, when appropriately restricted, to that r -correspondence; this is essentially for the same reason that we know no bijective proof of our claim 1.1.1. Instead our shape datum reduces to the *transpose* r -correspondence (so to speak its column insertion variant, although that term is not very appropriate in the r -ribbon case). Since horizontal r -ribbon strips (unlike r -ribbons) lack transposition symmetry, this distinction is significant.

Even more surprisingly, our correspondence does not reduce for $r = 1$ to the shape datum corresponding to the RSK correspondence, but to one associated to the so-called Burge correspondence. These shape data are fundamentally different, even if the RSK correspondence and the Burge correspondence are related (the relation also involves the Schützenberger involution on semistandard tableaux). The most crucial observation we had to make in order to find the shape datum used in our construction, was that although the shape datum for the Burge correspondence can be defined by iterating Schensted column-insertion, it has an alternative single-pass description (much like the bus transport in the previous subsection) that can easily be adapted to the context of semistandard r -ribbon tableaux.

1.3. Organisation.

The remainder of this paper is organised as follows. In section 2 we recall in detail Fomin’s general framework to define Schensted and Knuth correspondences. In section 3 we describe the shape data for the RSK correspondence, and for the Burge correspondence, while also indicating how the global correspondences defined by them are related. We close that section by giving the most trivial examples of correspondences with $r > 1$, namely those using \mathbf{Y}^r . In section 4 we recall the definitions involving r -rim hook lattices and semistandard r -ribbon tableaux, and the factoring of many questions concerning r -ribbons due to the r -quotient map; we then discuss the spin statistic and the r -correspondence defined by Shimozono and White, which do not factor in this way.

1.4 Acknowledgement

In section 5 we present our main result, the shape datum that leads to a spin preserving Knuth correspondence from matrices with entries in \mathbf{N}^r to pairs of semistandard r -ribbon tableaux. In section 6 we similarly generalise Knuth's asymmetric correspondence to one from matrices with entries in $\{0, 1\}^r$ to pairs consisting of a semistandard and a transpose semistandard r -ribbon tableau.

Although the presentation of our new constructions is our central goal around which the paper is organised, much of it can also be considered as an expository one of the various constructions on which we build forth. For this reason our pace will often be leisurely and our discussion informal. We have tried to limit the formal definitions and notations used, and most of those that are introduced serve a very localised purpose; giving them will therefore be postponed to the moment they are actually used. We shall also take time to discuss some matters that are not essential to our constructions, notably the reasons why certain approaches we tried were not successful.

1.4. Acknowledgement.

The author would like to thank Thomas Lam for a number of critical comments that helped improve this paper.

§2. Review of Fomin's constructions.

2.1. Schensted correspondences.

Let us recall the framework described in [Fom3]–[Fom4] for defining Schensted correspondences, while simplifying it slightly by omitting some generality that is not needed in the current paper. For a treatment of just the case of the Robinson-Schensted correspondence, we may refer to [vLee1, §3]. We keep in mind that special case throughout the discussion, and keep our notation close to what is customary there.

One starts with graded set \mathcal{P} of shapes, whose rank function $\mathcal{P} \rightarrow \mathbf{N}$ is written $\lambda \rightarrow |\lambda|$ and is such that each set $\mathcal{P}_i = \{ \lambda \in \mathcal{P} \mid |\lambda| = i \}$ is finite. A structure of graded graph on \mathcal{P} is defined by giving a relation contained in $\bigcup_{i \in \mathbf{N}} \mathcal{P}_i \times \mathcal{P}_{i+1}$, in other words a set of edges (λ, μ) with $|\mu| = |\lambda| + 1$. Although two different structures of graded graph on the same set \mathcal{P} are used in the theory (one for each of the “up” and “down” operators introduced below), these always coincide in the cases we shall consider, so we use a single symbol ‘ \prec ’ to denote this relation. Its reflexive and transitive closure will be denoted by ‘ \leq ’, which makes \mathcal{P} into a graded poset, from which one can retrieve the graded graph as its Hasse diagram. The latter need not be connected, and we shall encounter examples where this is not the case; however, since we shall be mainly interested in paths in the graph, we lose no generality by considering one connected component at the time. For $\lambda \leq \mu$ we define a path of shape μ/λ to be a monotonously rising path $\lambda = \lambda^0 \prec \lambda^1 \prec \cdots \prec \lambda^n = \mu$ in the graph, and we shall write $|\mu/\lambda| = n = |\mu| - |\lambda|$. Due to further requirements each connected component of the graph will have a unique minimal element, which by a shift in the grading on that component could be assumed to have rank 0. A path of shape μ/λ where λ is the minimal element of its component will be called simply a path of shape μ .

Let $\mathbf{Z}\mathcal{P}$ denote the free \mathbf{Z} -module on the set \mathcal{P} of shapes; one defines endomorphisms U, D of $\mathbf{Z}\mathcal{P}$ by their action on basis elements: $U: \lambda \mapsto \sum_{\mu \succ \lambda} \mu$ and $D: \lambda \mapsto \sum_{\mu \prec \lambda} \mu$ (these are well defined by the finiteness of each \mathcal{P}_i). The basic assumption we make of our graded graph is the commutation relation

$$D \circ U = U \circ D + r\mathbf{1} \tag{1}$$

for some $r \in \mathbf{N}_{>0}$, where $\mathbf{1}$ is the identity operator. It is clear that in any case both members of this equation preserve the grading. So if one applies the equation to some $\lambda \in \mathcal{P}_i$, it states two things: taking the coefficient of λ itself, one gets

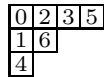
$$\#\{\mu \in \mathcal{P}_{i+1} \mid \mu \succ \lambda\} = \#\{\mu \in \mathcal{P}_{i-1} \mid \mu \prec \lambda\} + r, \tag{2}$$

and taking the coefficient of some $\lambda' \in \mathcal{P}_i$ with $\lambda \neq \lambda'$ one gets

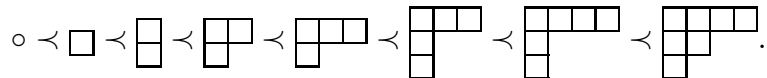
$$\#\{\mu \in \mathcal{P}_{i+1} \mid \mu \succ \lambda; \mu \succ \lambda'\} = \#\{\mu \in \mathcal{P}_{i-1} \mid \mu \prec \lambda; \mu \prec \lambda'\}. \tag{3}$$

The latter condition implies that if there is any path between two shapes, then there is one that is composed of a monotonous descending path followed by a monotonous rising path, so in particular each connected component of \mathcal{P} has a unique minimal element. Therefore equation (1) means that each such component is an r -differential poset as defined in [Stan1, definition 1.1]. One also sees easily that if condition (3) holds for all λ, λ' , then the cardinalities in the equation cannot exceed 1. In case a component of \mathcal{P} is a lattice, that condition amounts to the lattice being modular (i.e., $|\lambda \wedge \mu| + |\lambda \vee \mu| = |\lambda| + |\mu|$).

The classical example of a 1-differential poset is the Young lattice \mathbf{Y} consisting Young diagrams (finite order ideals of \mathbf{N}^2) ordered by inclusion. A partition $\lambda_0 \geq \lambda_1 \geq \dots$ represents a Young diagram $\lambda = \{(i, j) \in \mathbf{N}^2 \mid i < \lambda_j\}$, whose elements are referred to (and displayed) as squares. In \mathbf{Y} , a path $(0) = \lambda^0 \prec \dots \prec \lambda^n = \lambda$ of shape λ corresponds to a standard Young tableau of shape λ , which can be displayed by filling the squares of λ with the numbers in $[n] = \{0, \dots, n-1\}$ in such a way that for each i the set of entries in the diagram λ^i is $[i]$. Thus the standard Young tableau displayed as



corresponds to the path

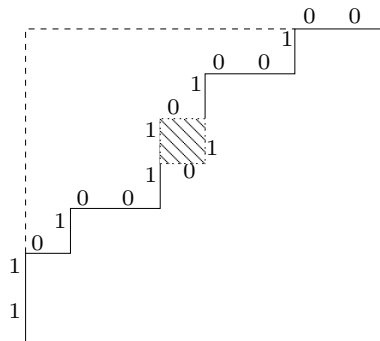


For our purposes we shall consider standard Young tableaux of shape λ just to be paths of shape λ .

Let us check that \mathbf{Y} is a 1-differential poset. Since \mathbf{Y} is a modular (even distributive) lattice (with $\lambda \wedge \mu = \lambda \cap \mu$ and $\lambda \vee \mu = \lambda \cup \mu$), it suffices to check condition (2), which states that for any Young diagram there is one more square that can be added

2.1 Schensted correspondences

to it than there are squares that can be removed from it. Although this is quite easy to see, we give a formal argument, since variations of it will be used later to prove less obvious statements. A Young diagram λ is completely determined by its edge sequence $\delta(\lambda)$, a doubly infinite bit string describing (as in §1.1) the segments of the path that forms the edge of λ , from bottom left to top right (for a formal definition of $\delta(\lambda)$ see [vLee2]; it is essentially the (Comét) code C_λ of λ used in [Stan2, Exercise 7.59], but with all bits complemented). The bits eventually become 1 to the left (for the rows of length 0) and 0 to the right (for the columns of height 0). Occurrences of a substring “10” correspond to squares that can be added to the diagram, causing the substring to be replaced by “01”, occurrences of which therefore correspond to squares that can be removed. We illustrate this for the diagram λ of $(6, 4, 3, 3, 1)$, and the bracketed substring of $\delta(\lambda) = (\dots 1101001[10]100100\dots)$.



In any word over the alphabet $\{0, 1\}$ the occurrences of “10” and “01” are perfectly interleaved from left to right. The limiting behaviour of letters in both directions ensures that the set of these occurrences is finite and non-empty, and that the first and the last ones are occurrences of “10”. Therefore the number of occurrences of “10” exceeds that of occurrences of “01” by one.

As is shown in [Stan1], many enumerative properties can be derived uniformly for all r -differential posets, i.e., they depend only on the identity (1). In [Fom4] bijections corresponding to such identities are constructed, but this requires some additional ingredients (as could be expected, since otherwise there is nothing to discriminate between the elements of \mathcal{P}), namely a family of bijections that correspond to local instances of (1), i.e., linking the sets counted by the members of equations (2) and (3). In fact one needs not bother about (3), since we have observed that those sets are either empty or singletons, which leaves no choice for a bijection. It is not obvious which set should be counted by the final term ‘ r ’ in (2); any r element set disjoint from \mathcal{P} would do. Recall that we denoted by $[r]$ the r element order ideal $\{i \in \mathbf{N} \mid i < r\} = \{0, \dots, r-1\}$ of \mathbf{N} , which we take as our standard r -element set. We could let the mentioned term ‘ r ’ correspond to $[r]$, but it will be clearer, and more flexible for generalising later, to take a set $e_{[r]} = \{e_i \mid i \in [r]\}$ of r symbols that are reserved to serve as exceptional values. We then define an r -correspondence on \mathcal{P} to consist of a family of bijections

$$b_\lambda: \{\mu \in \mathcal{P} \mid \mu \succ \lambda\} \rightarrow \{\mu \in \mathcal{P} \mid \mu \prec \lambda\} \cup e_{[r]} \quad \text{for } \lambda \in \mathcal{P}. \quad (4)$$

(In fact r -correspondences are defined in [Fom4] as bijections between edges rather than between vertices of the graph, but in the situation we are considering, the simpler notion above gives the same information). Assuming that (1) holds, an r -correspondence always exists, although it might be hard to specify; usually however, the proof of (1) suggests a natural choice of an r -correspondence. For \mathbf{Y} , we see that there are in fact two natural choices: one may either associate to any occurrence of “10” in $\delta(\lambda)$ the occurrence of “01” to its left, except that the leftmost occurrence of “10” is sent to the exceptional value e_0 , or one can do the same thing with “left” replaced by “right”. Due to the symmetry of \mathbf{Y} with respect to transposition, there is no fundamental difference between the two choices, but the choice that will lead to the Schensted correspondence in its usual form (i.e., using row insertion) is to move to the right: each square that can be added will then correspond to the square that can be removed from the row directly above it, with the square that can be added to the topmost row corresponding to e_0 .

Once an r -correspondence has been fixed, it may be used to define bijective correspondences that match identities derived from equation (1). Although there are many of these, we focus on the following one:

$$\langle \epsilon \mid D^n(U^n(\epsilon)) \rangle = n!r^n \quad \text{for a minimal element } \epsilon \text{ of } \mathcal{P}, \text{ and } n \in \mathbf{N}, \quad (5)$$

where the scalar product is the canonical one in \mathbf{ZP} , so that the left hand side gives the coefficient of ϵ in $D^n(U^n(\epsilon))$. The identity can be derived by showing that when (1) is used repeatedly to rewrite $D^n \circ U^n$ as a sum of terms of the form $U^i \circ D^i$ with $0 \leq i \leq n$, then $U^0 \circ D^0 = \mathbf{1}$ occurs with coefficient $n!r^n$. Since the operators D and U are adjoint with respect to the scalar product, the left hand side of (5) can also be written as $\langle U^n(\epsilon) \mid U^n(\epsilon) \rangle$. Every $\lambda \in \mathcal{P}_{|\epsilon|+n}$ with $\epsilon \leq \lambda$ occurs in $U^n(\epsilon)$, with as multiplicity the number of paths of shape λ , so (5) states that the sum of the squares of these numbers equals $n!r^n$. For $\mathcal{P} = \mathbf{Y}$, paths of shape λ are standard tableaux of shape λ , and for the value $r = 1$ applicable to this case the number $n!r^n = n!$ of course counts the permutations of n , so here equation (5) gives the enumerative identity that the Robinson-Schensted correspondence proves. Any correspondence similarly proving (5) for some r -differential poset, with the right hand side $n!r^n$ counting the r -coloured permutations of n , can therefore be considered to be a generalisation of the Robinson-Schensted correspondence.

Given an r -correspondence on \mathcal{P} , one can obtain such a correspondence by defining, for any minimal element ϵ of \mathcal{P} and $n \in \mathbf{N}$, an intermediate set of “growth diagrams” that is in bijection both with the set of pairs of paths of some common shape $\lambda \in \mathcal{P}_{|\epsilon|+n}$ with $\epsilon \leq \lambda$ and with the set of r -coloured permutations of n . These bijections are just projections that extract partial information from a growth diagram; their bijectivity therefore means that growth diagrams can always be uniquely reconstructed from such information. We formally define r -coloured permutations of n to be matrices $(A_{i,j})_{i,j \in [n]}$ with entries $A_{i,j} \in \{0\} \cup e_{[r]}$ such that $|A_{i,j}|_{i,j \in [n]}$ is a permutation matrix, where we set $|e_i| = 1$ for all $i \in [r]$. Specifying an r -coloured permutation of n is equivalent to giving the permutation σ of n such that $|A_{i,j}| = 1$ whenever $\sigma(i) = j$, and the n values $A_{i,\sigma(i)} \in e_{[r]}$ ($i \in [n]$), which explains the name.

2.1 Schensted correspondences

2.1.1. Definition. Let \mathcal{P} be a graded poset equipped with an r -correspondence $\{b_\lambda \mid \lambda \in \mathcal{P}\}$ as defined by equation (4), and let a minimal element ϵ of \mathcal{P} and $n \in \mathbf{N}$ be fixed. A growth diagram, or Schensted-growth, relative to these data consists of a pair of maps, the first one mapping $[n+1]^2 \rightarrow \mathcal{P}$ and written $(k, l) \mapsto \lambda^{(k, l)}$, the second one mapping $[n]^2 \rightarrow \{0\} \cup e_{[r]}$ and written $(k, l) \mapsto A_{k, l}$, that satisfy the following conditions for all $k, l \in [n]$:

- (0) $\lambda^{(0, 0)} = \epsilon$;
- (1) $\lambda^{(k, l)} \preceq \lambda^{(k, l+1)}$ and $\lambda^{(k, l)} \preceq \lambda^{(k+1, l)}$;
- (2) $\lambda^{(n, l)} \prec \lambda^{(n, l+1)}$ and $\lambda^{(k, n)} \prec \lambda^{(k+1, n)}$;
- (3) $|\lambda^{(k, l)}| = |\epsilon| + \sum_{(i, j) \in [k] \times [l]} |A_{i, j}|$;
- (4) If $\lambda^{(k, l+1)} = \lambda^{(k+1, l)} \prec \lambda^{(k+1, l+1)}$ then, putting $\mu = \lambda^{(k, l+1)}$, $\kappa = \lambda^{(k+1, l+1)}$ and $v = b_\mu(\kappa)$, one has either $\lambda^{(k, l)} = v$ (if $v \in \mathcal{P}$), or $A_{k, l} = v$ (if $v \in e_{[r]}$).

The alternative term ‘‘Schensted-growth’’ for growth diagram will be used later to distinguish this notion from a similar one used in relation to Knuth correspondences. It follows from (2) and (3) that $(A_{i, j})_{i, j \in [n]}$ is an r -coloured permutation of n . One also has $\lambda^{(i, 0)} = \epsilon = \lambda^{(0, i)}$ for $i \in [n+1]$, so that $P = (\lambda^{(n, 0)} \prec \lambda^{(n, 1)} \prec \dots \prec \lambda^{(n, n)})$ and $Q = (\lambda^{(0, n)} \prec \lambda^{(1, n)} \prec \dots \prec \lambda^{(n, n)})$ are paths of shape $\lambda^{(n, n)}$; this defines the two mentioned projections. For the permutation σ of n corresponding to $(A_{i, j})_{i, j \in [n]}$, one has $\lambda^{(k, l)} = \lambda^{(k+1, l)}$ in condition (1) if and only if $\sigma(k) \geq l$, and $\lambda^{(k, l)} = \lambda^{(k, l+1)}$ if and only if $\sigma^{-1}(l) \geq k$. In particular $|A_{k, l}| = 1$ implies that $\lambda^{(k, l)} = \lambda^{(k+1, l)} = \lambda^{(k, l+1)}$, so in condition (4) one has $A_{k, l} = 0$ in the case $v \in \mathcal{P}$ (since $b_\mu(\kappa) \neq \mu$), while $\lambda^{(k, l)} = \lambda^{(k+1, l)} = \lambda^{(k, l+1)}$ in the case $v \in e_{[r]}$.

It is convenient to display a growth diagram by attaching the shape $\lambda^{(k, l)}$ to the point (k, l) of a grid (with k increasing downwards and l increasing to the right), and writing the matrix entry $A_{k, l}$ into the square of that grid with corners (k, l) , $(k, l+1)$, $(k+1, l)$, and $(k+1, l+1)$. Figure 1 thus illustrates a growth diagram for $\mathcal{P} = \mathbf{Y}$ with the mentioned left-to-right 1-correspondence; non-zero matrix entries, which due to $r = 1$ are necessarily equal to e_0 , are indicated by ‘ \star ’.

The pair (P, Q) gives the values of $\lambda^{(i, j)}$ for $(i, j) \in [n+1]^2 \setminus [n]^2$, so reconstruction of a growth diagram by decreasing values of k and l , from arbitrary (P, Q) , will be possible if for any $k, l \in [n]$ one can uniquely determine the shape $\lambda = \lambda^{(k, l)}$ and the matrix entry $A_{k, l}$ once the shapes $\mu = \lambda^{(k, l+1)}$, $\nu = \lambda^{(k+1, l)}$, and $\kappa = \lambda^{(k+1, l+1)}$ are known. And indeed this is the case: if κ is equal to μ or to ν , then $\lambda^{(k, l)}$ will be equal to the other, and otherwise if μ and ν are distinct, then $\lambda^{(k, l)}$ will be the unique element covered by both (which exists because both are covered by κ); in both these cases $A_{k, l} = 0$. The remaining case $\mu = \nu \neq \kappa$ is the one handled by condition (4); there, as we saw, both λ and $A_{k, l}$ are always determined. From this description one sees that the differential version $|\kappa| - |\mu| - |\nu| + |\lambda| = |A_{k, l}|$ of condition (3) holds in all cases; the verification of the remaining requirements is now easy. A growth diagram can be similarly reconstructed from the matrix $(A_{i, j})_{i, j \in [n]}$, by increasing values of k and l : one determines $\lambda^{(k+1, l+1)}$ once $\lambda^{(k, l)}$, $\lambda^{(k, l+1)}$, $\lambda^{(k+1, l)}$ (and of course $A_{k, l}$) are known; to prove this, similar cases as above can be distinguished, with the case $|A_{k, l}| = 1$ being singled out first.




Figure 1. Schensted-growth for \mathbf{Y} with $\sigma = (0 \ 1 \ 2 \ 3 \ 4 \ 5 \ 6 \ 7)$, $P = \begin{bmatrix} 0 & 2 & 3 \\ 1 & 5 & 7 \\ 4 & 6 \end{bmatrix}$, and $Q = \begin{bmatrix} 0 & 2 & 5 \\ 1 & 4 & 6 \\ 3 & 7 \end{bmatrix}$.

For $\mathcal{P} = \mathbf{Y}$ with the chosen 1-correspondence, the resulting bijection between permutations and pairs of standard Young tableaux is the classical Robinson-Schensted correspondence. The Schensted insertion of the permutation entry $\sigma(i)$ corresponds to the computation, given the line $\lambda^{(i,0)} \preceq \dots \preceq \lambda^{(i,n)}$ of the growth diagram and the entry $A_{i,\sigma(i)} = e_0$, of the next line $\lambda^{(i+1,0)} \preceq \dots \preceq \lambda^{(i+1,n)}$.

The method of reconstructing a growth diagram from (P, Q) or from $(A_{i,j})_{i,j \in [n]}$ can be viewed as a data-flow network consisting of n^2 copies of the r -correspondence respectively of its inverse: for every pair $k, l \in [n]$, a copy of the r -correspondence links the values $A_{k,l}$, $\lambda^{(k,l)}$, $\lambda^{(k+1,l)}$, $\lambda^{(k,l+1)}$, and $\lambda^{(k+1,l+1)}$. In fact we had to complement the r -correspondence with several cases (where no choice was involved) to determine all necessary values, but in the full generality of [Fom4] these cases are already incorporated into the r -correspondence itself. Moreover that r -correspondence operates on the edges rather than on the vertices of the square $\lambda^{(k,l)}$, $\lambda^{(k+1,l)}$, $\lambda^{(k,l+1)}$, and $\lambda^{(k+1,l+1)}$, with each of the two output edges being one of the inputs to a unique other copy of the r -correspondence (unless it is part of the final output), so that the data-flow nature is even more clearly visible there.

One may observe that apart from extraction of the pair (P, Q) of paths in \mathcal{P} , or of A , the r -coloured permutation of n , other projections allow unique reconstruction of a growth diagram as well: it suffices to know $\lambda^{(i,j)}$ for all (i, j) on some lattice path from $(n, 0)$ to $(0, n)$, and all matrix entries $A_{i,j}$ that lie below and to the right of that path. Starting with the shapes along the bottom and right edges of the diagram as given by (P, Q) , one can therefore replace this information by the shapes along paths that are gradually modified to eventually become the path along the left at top edges,

2.2 Knuth correspondences

while retaining all matrix entries that the path has moved across; then at each moment one has complete information to determine the entire growth diagram. This procedure may be interpreted as describing the process of rewriting the left hand side of (5) to its right hand side: it provides a way to trace, and thus obtain a matching between, the contributions to these two expressions and to all intermediate ones.

2.2. Knuth correspondences.

We shall now consider the generalisation of this construction to “Knuth correspondences”, which is described in [Fom5]. The fundamental difference with Schensted correspondences is that matrix entries will be chosen independently, without the restriction that they should give rise to a permutation matrix. Matrix entries will be chosen from some graded set S , whose rank function $S \rightarrow \mathbf{N}$ written $x \rightarrow |x|$ is such that each set $S_i = \{x \in S \mid |x| = i\}$ is finite, and $S_0 = \{0\}$. We shall define growth diagrams similar to those defined in above, for which in particular condition 2.1.1(4) will be required. From the independence of matrix entries it then follows that one has to allow arbitrarily large differences of rank between shapes on adjacent grid points (at least, growing with n), even if all non-zero matrix entries should have rank 1. Therefore one needs to equip \mathcal{P} with a relation that unlike ‘ \preceq ’ can hold between shapes any number of levels apart. We shall denote this relation, which in general will not be transitive, by ‘ \leftarrow ’, with $\mu \rightarrow \lambda$ meaning the same as $\lambda \leftarrow \mu$. At this point we make no assumption about this relation other than that $\lambda \leftarrow \mu$ and $|\lambda| = |\mu|$ imply $\lambda = \mu$. When $\lambda \leftarrow \mu$ we shall as before write $|\mu/\lambda| = |\mu| - |\lambda|$.

The adaptation of the rules governing growth diagrams is quite straightforward. Each matrix entry $a = A_{i,j} \in S$ is associated to a square of the grid, the corners of which have shapes attached to them, which as before we shall designate by $\lambda = \lambda^{(i,j)}$, $\mu = \lambda^{(i,j+1)}$, $\nu = \lambda^{(i+1,j)}$, and $\kappa = \lambda^{(i+1,j+1)}$; these satisfy $\lambda \leftarrow \mu \leftarrow \kappa$ and $\lambda \leftarrow \nu \leftarrow \kappa$. From condition 2.1.1(4) it follows moreover that

$$|\kappa| - |\nu| - |\mu| + |\lambda| = |a|. \quad (6)$$

As in the case of Schensted correspondences, we wish that given μ, ν , and κ , the shape λ and the entry a can be uniquely determined, and that similarly given λ, μ, ν and a , the shape κ can be determined. In other words, for every pair $\mu, \nu \in \mathcal{P}$ there should be a bijection between the sets $\{\kappa \in \mathcal{P} \mid \kappa \rightarrow \mu; \kappa \rightarrow \nu\}$ and $\{(a, \lambda) \in S \times \mathcal{P} \mid \lambda \leftarrow \mu; \lambda \leftarrow \nu\}$ such that equation (6) holds whenever (a, λ) corresponds to κ . Such a family of bijections will be a central notion of this paper, so we state a formal definition.

2.2.1. Definition. *Let \mathcal{P} be a graded set equipped with a relation ‘ \leftarrow ’, and S a graded set. A shape datum for $(\mathcal{P}, \leftarrow, S)$ consists of a family $(b_{\mu,\nu})_{\mu,\nu \in \mathcal{P}}$ of bijections*

$$b_{\mu,\nu}: \{\kappa \in \mathcal{P} \mid \mu \leftarrow \kappa \rightarrow \nu\} \rightarrow \{(a, \lambda) \in S \times \mathcal{P} \mid \mu \rightarrow \lambda \leftarrow \nu\}, \quad (7)$$

such that $|\kappa| - |\nu| - |\mu| + |\lambda| = |a|$ holds whenever $(a, \lambda) = b_{\mu,\nu}(\kappa)$.

In order for a shape datum to exist, $(\mathcal{P}, \leftarrow, S)$ must satisfy an enumerative identity that is most easily expressed using formal power series. We shall use the ring

$\text{End}(\mathbf{ZP})[[X, Y]]$ of formal power series in indeterminates X, Y over the ring of endomorphisms of \mathbf{ZP} . Such a power series can be applied to an element of \mathbf{ZP} to give a formal sum of monomials in X and Y with coefficients in \mathbf{ZP} , i.e., an element of the $\mathbf{Z}[[X, Y]]$ -module $\mathbf{ZP}[[X, Y]]$. This action on elements of \mathbf{ZP} can be extended without problem to an action on power series, in other words we can interpret elements of $\text{End}(\mathbf{ZP})[[X, Y]]$ as endomorphisms of $\mathbf{ZP}[[X, Y]]$. We define two such elements $U_X, D_Y \in \text{End}(\mathbf{ZP})[[X, Y]]$ in terms of their action, by requiring

$$U_X(\lambda) = \sum_{\mu \rightarrow \lambda} X^{|\mu/\lambda|} \mu \quad \text{and} \quad D_Y(\lambda) = \sum_{\mu \leftarrow \lambda} Y^{|\lambda/\mu|} \mu \quad \text{for all } \lambda \in \mathcal{P}. \quad (8)$$

Let $F_S(T) \in \mathbf{Z}[[T]]$ be the rank generating series $\sum_{a \in S} T^{|a|} = \sum_{i \in \mathbf{N}} c_i T^i$ of S , where $c_i = \#S_i$. Then the enumerative requirement for the existence of a shape datum will be

$$D_Y \circ U_X = (U_X \circ D_Y)(F_S(XY)). \quad (9)$$

If one writes $U_X = \sum_{i \in \mathbf{N}} U_i X^i$ and $D_Y = \sum_{i \in \mathbf{N}} D_i Y^i$, then taking the coefficient of $X^i Y^j$ in (9) gives

$$D_j \circ U_i = \sum_{k \leq \min\{i, j\}} c_k (U_{i-k} \circ D_{j-k}). \quad (10)$$

When both members of this equation are applied to μ , the values obtained are of rank $|\mu| + i - j$, and if one takes the coefficient of some $\nu \in \mathcal{P}_{|\mu|+i-j}$, the resulting numbers count the subsets of the domain and codomain of $b_{\mu, \nu}$ for which $|\kappa| = |\mu| + i$ respectively $|\lambda| = |\mu| - j + |a|$; these are subsets that should be in correspondence under $b_{\mu, \nu}$ if it is to satisfy (6). Taking $i = j = 1$ in (10) gives $D_1 \circ U_1 = U_1 \circ D_1 + c_1 \mathbf{1}$, which shows that the current situation extends the one considered earlier: if one defines $\lambda \prec \mu$ when $\lambda \leftarrow \mu$ and $|\lambda| + 1 = |\mu|$, then (2) and (3) will hold, with $r = c_1$. There is however no easy way to extend the structure (\mathcal{P}, \prec) to $(\mathcal{P}, \leftarrow, S)$; indeed it is not obvious why any combinatorial structures satisfying (9) should exist at all. Yet several examples are given in [Fom5], among which the example corresponding to the RSK correspondence (described in the next section), which is the one that will concern us most.

Once one has defined a shape datum, the construction of a global ‘‘Knuth’’ correspondence from it is straightforward, by analogy to the construction of Schensted correspondences. In fact, matters are slightly simpler since there is no need to explicitly distinguish cases according to the differences of rank between shapes associated to neighbouring points of the grid, as all possible cases are already catered for by the shape datum itself. Calling $\epsilon \in \mathcal{P}$ minimal when $\lambda \leftarrow \epsilon$ implies $\lambda = \epsilon$, we have the following notion of growth diagram.

2.2.2. Definition. Let $(b_{\mu, \nu})_{\mu, \nu \in \mathcal{P}}$ be a shape datum for $(\mathcal{P}, \leftarrow, S)$ as defined in 2.2.1, and let a minimal element ϵ of \mathcal{P} , and $n, m \in \mathbf{N}$ be fixed. A growth diagram, or Knuth-growth, relative to these data consists of a pair of maps, the first one mapping $[m + 1] \times [n + 1] \rightarrow \mathcal{P}$ and written $(k, l) \mapsto \lambda^{(k, l)}$, the second one mapping $[m] \times [n] \rightarrow S$ and written $(k, l) \mapsto A_{k, l}$, that satisfy the following conditions, for all $k \in [m]$ and $l \in [n]$:

2.2 Knuth correspondences

- (0) $\lambda^{(0,n)} = \epsilon = \lambda^{(m,0)}$;
 (1) $\lambda^{(k,l+1)} \leftarrow \lambda^{(k+1,l+1)}$, and $\lambda^{(k+1,l)} \leftarrow \lambda^{(k+1,l+1)}$;
 (2) With $\lambda = \lambda^{(k,l)}$, $\mu = \lambda^{(k,l+1)}$, $\nu = \lambda^{(k+1,l)}$, and $\kappa = \lambda^{(k+1,l+1)}$, one has $(A_{k,l}, \lambda) = b_{\mu,\nu}(\kappa)$.

Note that there is no longer a need to require explicitly that $|\lambda^{(k,l)}| = |\epsilon| + \sum_{(i,j) \in [k] \times [l]} |A_{i,j}|$, as this follows easily from the fact that $b_{\mu,\nu}$ respects equation (6). We shall again consider two projections from the set of growth diagrams: one extracting the matrix $A = (A_{k,l})_{k \in [m], l \in [n]}$, the other extracting the pair of paths $P = (\lambda^{(m,0)} \leftarrow \lambda^{(m,1)} \leftarrow \dots \leftarrow \lambda^{(m,n)})$ and $Q = (\lambda^{(0,n)} \leftarrow \lambda^{(1,n)} \leftarrow \dots \leftarrow \lambda^{(m,n)})$. In the latter, P and Q will be called paths of shape $\lambda^{(m,n)}$ (or more explicitly of shape $\lambda^{(m,n)}/\epsilon$) in $(\mathcal{P}, \leftarrow)$. Besides this (final) shape, a property of paths describing the successive ranks of the intermediate shapes is important: we define the *weight* of a path $p = (\lambda^0 \leftarrow \dots \leftarrow \lambda^n)$ to be the vector $\text{wt}(p) = (|\lambda^{i+1}/\lambda^i|)_{i \in [n]}$. One has $\text{wt}(P) = (\sum_{i \in [m]} |A_{i,j}|)_{j \in [n]}$ and $\text{wt}(Q) = (\sum_{j \in [n]} |A_{i,j}|)_{i \in [m]}$, in other words, the weights of the paths P, Q are the vectors of column sums of ranks and of row sums of ranks, respectively, of the matrix A . Both projections can be seen to be bijections, respectively to the set of all matrices with coefficients in S and to the set of all pairs of paths in $(\mathcal{P}, \leftarrow)$ of equal shape, in the same way as for the Schensted-growths. The correspondence between the matrices A and pairs (P, Q) of paths, defined by composing one projection with the inverse of the other, is the Knuth correspondence for $(\mathcal{P}, \leftarrow, S)$ associated to the given shape datum.

Like Schensted correspondences, Knuth correspondences imply enumerative identities, but since the sets linked by any bijection $b_{\mu,\nu}$ are infinite, these follow from the mentioned matching of the weights of (P, Q) and of A . One uses two sets of indeterminates $X_{[n]} = \{X_j \mid j \in [n]\}$ and $Y_{[m]} = \{Y_i \mid i \in [m]\}$ to record the weights of P and Q respectively: the contribution of each pair (P, Q) will be given by the monomial $X^{\text{wt}(P)} Y^{\text{wt}(Q)}$, where as usual $X^{(a_0, \dots, a_{n-1})}$ abbreviates $X_0^{a_0} \dots X_{n-1}^{a_{n-1}}$. By what was observed above, that monomial can also be expressed in terms of the matrix A corresponding to (P, Q) as $\prod_{i \in [m], j \in [n]} (X_j Y_i)^{|A_{i,j}|}$. Summation of that monomial over all $m \times n$ matrices A with entries in S gives $\prod_{i \in [n], j \in [m]} F_S(X_i Y_j)$ (note that we disentangle the indices, the mixing of which was caused by the unfortunate but conventional choice of reading off P from the growth diagram by varying the *second* index). If one defines the generating series $G_\lambda(X_{[n]}) = \sum_P X^{\text{wt}(P)}$, where the sum is over all paths $P = (\epsilon = \lambda^0 \leftarrow \dots \leftarrow \lambda^n = \lambda)$, then one finds the identity

$$\prod_{i \in [n], j \in [m]} F_S(X_i Y_j) = \sum_{\lambda \in \mathcal{P}} G_\lambda(X_{[n]}) G_\lambda(Y_{[m]}) \quad (11)$$

Noting that $G_\lambda(X_{[n]}) = \langle \lambda \mid U_{X_{n-1}} \dots U_{X_0}(\epsilon) \rangle$, this equation can be seen to follow directly from (9):

$$\begin{aligned}
\sum_{\lambda \in \mathcal{P}} G_\lambda(X_{[n]})G_\lambda(Y_{[m]}) &= \langle U_{Y_{m-1}} \cdots U_{Y_0}(\epsilon) \mid U_{X_{n-1}} \cdots U_{X_0}(\epsilon) \rangle \\
&= \langle \epsilon \mid D_{Y_0} \cdots D_{Y_{m-1}} U_{X_{n-1}} \cdots U_{X_0}(\epsilon) \rangle \\
&= \langle \epsilon \mid U_{X_{n-1}} \cdots U_{X_0} D_{Y_0} \cdots D_{Y_{m-1}}(\epsilon) \rangle \prod_{i \in [n], j \in [m]} F_S(X_i Y_j) \\
&= \prod_{i \in [n], j \in [m]} F_S(X_i Y_j),
\end{aligned} \tag{12}$$

where the last step follows from the fact that ϵ is minimal. This derivation demonstrates (again) the straightforwardness of deriving the global Knuth correspondence from the shape datum.

§3. Examples of shape data.

3.1. The shape datum for the RSK correspondence.

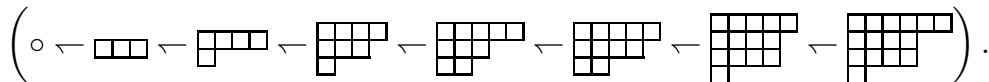
Let us now consider the RSK correspondence, and its shape datum. In this case $S = \mathbf{N}$, $\mathcal{P} = \mathbf{Y}$, and the relation $\lambda \prec \mu$ means that μ/λ is a horizontal strip: $\lambda \subseteq \mu$, and the skew diagram of μ/λ (the set-theoretic difference of diagrams $\mu \setminus \lambda$) has at most one square in any column. When $\lambda \prec \mu$, a specific path in (\mathcal{P}, \prec) (such paths are also called skew standard tableaux) of shape μ/λ is defined by requiring that in the successive diagrams along the path from λ to μ , the squares are added in left to right order, i.e., by (strictly) increasing column number. By interpolating in this way each of its horizontal strips (while eliminating any strips of size 0), a path in (\mathcal{P}, \prec) can be transformed into a (skew) standard tableau of the same shape, called its standardisation. The RSK correspondence, in the “extraction” direction (determining the matrix A from a pair (P, Q) of equal shaped paths in (\mathcal{P}, \prec)), can now be informally described as follows. First the standardisations of P and Q are determined, and then the Schensted-growth corresponding to this pair of standard Young tableaux is constructed. Due to the initial interpolation, this growth diagram is defined on a grid that is too extensive for the matrix A one wishes to find; therefore the grid is now reduced by forgetting its points on horizontal or vertical lines that were introduced by the interpolation, so that the paths remaining along the bottom and right edges are the original P and Q , before standardisation. The shapes associated to the remaining grid points will be the ones of a Knuth-growth; the matrix entry to be associated with a grid square whose corners carry shapes $\begin{pmatrix} \lambda & \mu \\ \nu & \kappa \end{pmatrix}$ is $|\kappa| - |\nu| - |\mu| + |\lambda|$, which is the sum over the corresponding rectangular area of the Schensted-growth of the associated permutation matrix entries. Proving that this description defines (a Knuth-growth for) a shape datum for (\mathbf{Y}, \prec) essentially amounts to showing that the shapes associated to adjacent grid points after reduction differ by horizontal strips; this can be deduced easily from a consideration of local portions of the Schensted-growth consisting of just two horizontally or vertically adjacent squares of the grid for that growth diagram.

Let us demonstrate the construction by a concrete example, which we take from Knuth’s original paper [Knu]. The semistandard tableaux can be readily interpreted as paths, so for instance

3.1 The shape datum for the RSK correspondence

1	1	1	2	4	7
2	3	3	5		
3	4	6	6		
6					

represents



The semistandard tableaux considered are

$$P = \begin{array}{|c|c|c|c|c|c|} \hline 1 & 1 & 1 & 2 & 4 & 7 \\ \hline 2 & 3 & 3 & 5 & & \\ \hline 3 & 4 & 6 & 6 & & \\ \hline 6 & & & & & \\ \hline \end{array} \quad \text{and} \quad Q = \begin{array}{|c|c|c|c|c|c|} \hline 1 & 2 & 2 & 3 & 3 & 6 \\ \hline 3 & 3 & 3 & 4 & & \\ \hline 4 & 5 & 5 & 5 & & \\ \hline 5 & & & & & \\ \hline \end{array}$$

whose standardisations are

$$\begin{array}{|c|c|c|c|c|c|} \hline 0 & 1 & 2 & 4 & 9 & 14 \\ \hline 3 & 6 & 7 & 10 & & \\ \hline 5 & 8 & 12 & 13 & & \\ \hline 11 & & & & & \\ \hline \end{array} \quad \text{and} \quad \begin{array}{|c|c|c|c|c|c|} \hline 0 & 1 & 2 & 6 & 7 & 14 \\ \hline 3 & 4 & 5 & 9 & & \\ \hline 8 & 11 & 12 & 13 & & \\ \hline 10 & & & & & \\ \hline \end{array}$$

It would take too much space to draw the entire Schensted-growth for this example, so figure 2 just displays the Knuth-growth derived from it, and details the Schensted-growth for one of its grid squares. The square that has been refined is indicated by brackets; matrix entries are indicated when non-zero.

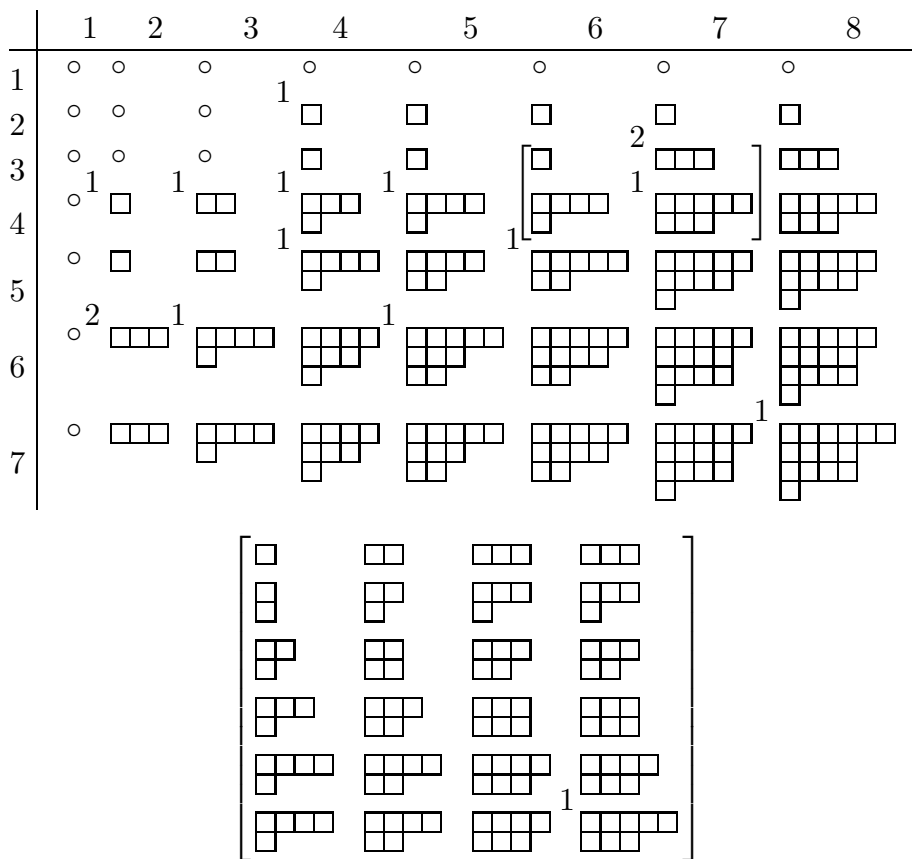


Figure 2. A Knuth-growth for the RSK correspondence, with detail of the associated Schensted-growth.

For each rectangular area of the Schensted-growth for which the paths along the bottom and right edges are standardisations of horizontal strips (like the bracketed area above), any entries 1 in the matrix will be produced along the end of the diagonal that ends in the bottom right corner of the rectangle. Therefore the computation of the shape datum in the “insertion” direction, determining the shape κ at the bottom right corner when the shapes λ, μ, ν at the other corners and the matrix entry a are known, can be done as follows: first complete a Schensted-growth with the standardisations of μ/λ and ν/λ along its top and left edges and with all matrix entries equal to 0, then extend the grid by a more lines at the right and at the bottom, copying the skew standard tableaux found along the right and bottom edges of the rectangle onto each of them, and finally complete the empty $a \times a$ grid square so created (at the bottom right of the grid) as a Schensted-growth containing an $a \times a$ identity matrix. Note that if one combines these individual Schensted-growths to a global one matching the full Knuth correspondence, these extensions of the grid apply to an entire row or column; therefore the entries in a single row of A will correspond to a collection of entries 1 in the final permutation matrix that descend, when traversed from left to right, into consecutive rows; similarly entries in a single column of A give rise to entries 1 that, when traversed from top to bottom, move rightwards into consecutive columns.

Let us describe this the shape datum for the RSK correspondence, in other words the relation between the shapes $\begin{pmatrix} \lambda & \mu \\ \nu & \kappa \end{pmatrix}$ at the corners of a square of the Knuth-growth and the matrix entry a inside, in somewhat different terms. Being defined by a part of a Schensted-growth, it can be computed by a localised version of the Schensted algorithm, in fact by an instance of the algorithm for skew tableaux of [SaSt]. In their terminology, it can be described (in the insertion direction) as follows: successively internally insert the squares of the standardisation of ν/λ into a tableau of shape μ/λ all of whose entries are equal, and then externally insert a more copies of the same entry; the shape of the resulting tableau gives κ/ν . The fact that μ/λ and ν/λ are horizontal strips severely restricts the insertion process, which allows us to describe the shape datum more directly. Only the internal insertion steps can involve “bumping” (corresponding in the Schensted-growth to the case of 2.1.1(4)) and they can do so at most once. This happens whenever the square inserted from ν/λ also occurs in the diagram of μ/λ ; it results in adding a new square on the first available place in the next row. We note that for $\mu = \nu$ and $a = 0$ one finds essentially the bijection given in the introduction to prove the case $r = 1$ of our claim 1.1.1.

A somewhat more formal direct description of this shape datum can be given using a description of horizontal strips as multisets of rows: a horizontal strip of the form λ/μ will be described by specifying for each row i the number $\lambda_i - \mu_i$ of its squares in that row (while this information does not completely describe the strip, it suffices in our setting, where always one of the shapes λ, μ involved is known beforehand). The operation of multiset intersection (in which each element gets as multiplicity the minimum is its multiplicities in the operands) provides a succinct way to describe the common squares of two strips like μ/λ and ν/λ . We also need an operation that shifts one row upwards: if A is a multiset of rows then A^\uparrow denotes another such multiset $\{\{i - 1 \mid i \in A, i > 0\}\}$, in other words the multiplicity of row i in A^\uparrow is the multiplicity of row $i + 1$ in A . Now

3.2 An alternative shape datum

our shape datum b^{Kn} is then specified by $b_{\mu,\nu}^{\text{Kn}}(\kappa) = (a, \lambda)$ where

$$\mu/\lambda = \kappa/\nu - (\kappa/\mu \cap \kappa/\nu) + (\kappa/\mu \cap \kappa/\nu)^\uparrow \quad (13)$$

and

$$\nu/\lambda = \kappa/\mu - (\kappa/\mu \cap \kappa/\nu) + (\kappa/\mu \cap \kappa/\nu)^\uparrow, \quad (14)$$

(the two equations are equivalent), and $a = |\kappa| - |\nu| - |\mu| + |\lambda|$, which is the multiplicity in $\kappa/\mu \cap \kappa/\nu$ of row 0. Equations (3.1) express the relation $\mu/\lambda \cap \nu/\lambda = (\kappa/\mu \cap \kappa/\nu)^\uparrow$, which is the fact we saw above that common squares in row i of μ/λ and ν/λ are in bijection with common squares in row $i+1$ of κ/μ and κ/ν , as well as the fact that the remaining squares of μ/λ and of ν/λ respectively match squares in the same row in κ/ν and in κ/μ (in fact the very same squares; this part is already implied by the inclusions among the shapes $\lambda, \mu, \nu, \kappa$). These equations can be restated as identities in terms of the individual parts of the shapes, which can then be simplified to give the equivalent equations

$$\begin{aligned} a &= \kappa_0 - \max(\mu_0, \nu_0), \\ \lambda_i &= \min(\mu_i, \nu_i) + \max(\mu_{i+1}, \nu_{i+1}) - \kappa_{i+1} \quad \text{for } i \in \mathbf{N}. \end{aligned} \quad (15)$$

Using the fact that $\lambda \leftarrow \mu$ means $\mu_{i+1} \leq \lambda_i \leq \mu_i$ for all $i \in \mathbf{N}$, one easily checks that with these equations, $\mu \leftarrow \kappa \rightarrow \nu$ is equivalent to $a \geq 0$ and $\mu \rightarrow \lambda \leftarrow \nu$, so that b^{Kn} is indeed a shape datum.

3.2. An alternative shape datum.

There is different shape datum for $(\mathbf{Y}, \leftarrow, \mathbf{N})$ that can be obtained by similar methods; it suffices to replace the (row insertion) 1-correspondence used above by its transpose (column insertion) 1-correspondence. One may proceed in exactly the same way to construct a Knuth-growth, interpolating horizontal strips, then building a Schensted-growth, and finally reducing the grid again. This works because in a Schensted-growth for the transpose 1-correspondence, it still holds that any rectangular area for which the path along the right or bottom edge of is the standardisation of a horizontal strip, has such a path along the opposite edge as well (ignoring trivial steps that repeat the same shape). The Knuth-correspondence so obtained is known as the Burge correspondence (although only the case where P - and Q -symbol coincide is actually used in [Bur]). It should not be confused with the Knuth's asymmetric correspondence mentioned in the introduction; that correspondence is not of the type we are currently considering, where the P - and Q -symbol must both be semistandard and of equal shape. We illustrate the construction of the Burge correspondence in figure 3, in a similar way as we did for the RSK correspondence.

To extract a shape datum b^{Bu} from the Burge correspondence, one may again limit the Schensted-growth to a rectangular area like the bracketed one, with standardisations of horizontal strips at the right and bottom. From those standardisations, the rules for a Schensted-growth with the column insertion 1-correspondence determine the other shapes and the matrix entries; the shape λ at the top left corner and the sum a of all matrix entries produced define the shape datum $b_{\mu,\nu}^{\text{Bu}}(\kappa) = (a, \lambda)$. Any matrix entry 1

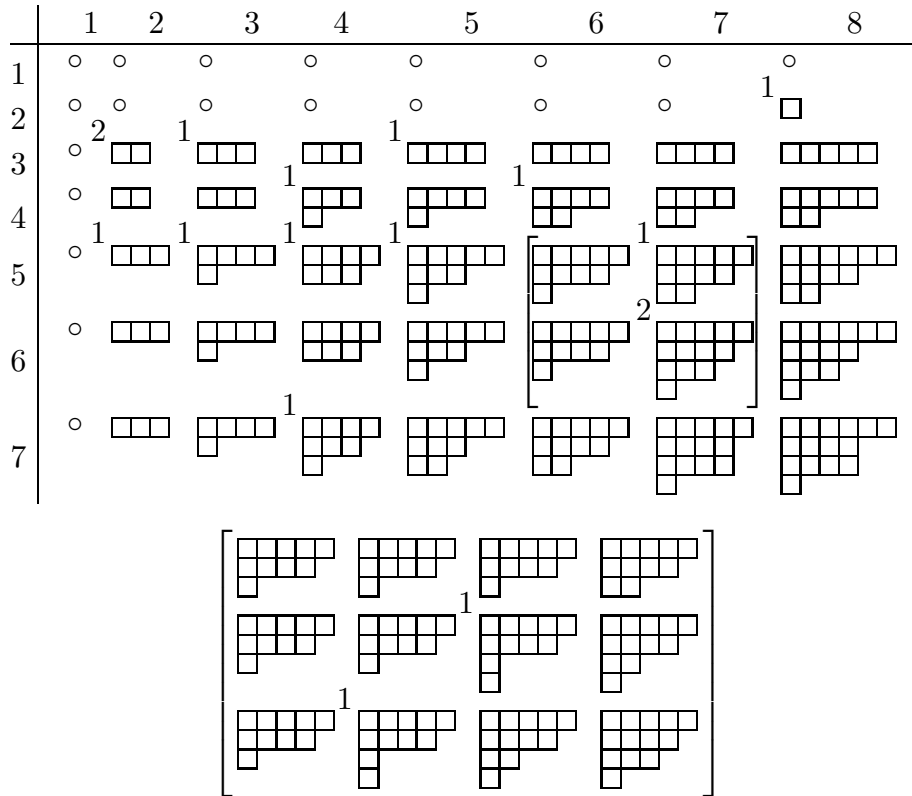


Figure 3. Knuth-growth for the Burge correspondence; detail of the associated Schensted-growth.

occurring in the Schensted-growth now adds a square in column 0, which is the first in any horizontal strip in which it occurs, so it can only be added to the shape λ . Therefore the matrix entries 1 form an anti-diagonal at the top left corner of the rectangle, and the computation in the opposite (insertion) direction should proceed as follows: construct a partial Schensted-growth with matrix entries 1 along the anti-diagonal of an $a \times a$ square in the top left corner and 0 elsewhere, and shapes λ repeated along the top and left of edges that square, followed respectively by standardisations of μ/λ and ν/λ .

Considering the permutation matrix entries 1 contributing to one row or one column of the final matrix, the above description implies a different arrangement than for the RSK correspondence: they are arranged in an anti-diagonal sense (bottom left to top right) in a sequence of consecutive rows respectively columns. In terms of the original “bumping” description, this means that there are, in addition to using column insertion rather than row insertion, two more differences of the Burge insertion algorithm with respect to that of Knuth: the order of insertion is reversed among columns of the two-line array with equal top index (making it weakly *decreasing* by bottom index), and among equal entries in the insertion tableau P , the more recently inserted ones are treated as smaller, so that bumping will replace an entry of the same value if present in the column.

Although it is of no importance to our paper, we should mention that there is a relation between the global RSK and Burge correspondences, which involves the

3.2 An alternative shape datum

Schützenberger involution. It is illustrated by the examples shown in figures 2 and 3: their matrices are vertical mirror images (the same rows appear in reverse order), their P -symbols are the same while their Q -symbols are Schützenberger duals:

$$P = \begin{array}{|c|c|c|c|c|c|} \hline 1 & 1 & 1 & 2 & 4 & 7 \\ \hline 2 & 3 & 3 & 5 & & \\ \hline 3 & 4 & 6 & 6 & & \\ \hline 6 & & & & & \\ \hline \end{array} = P', \quad Q = \begin{array}{|c|c|c|c|c|c|} \hline 1 & 2 & 2 & 3 & 3 & 6 \\ \hline 3 & 3 & 3 & 4 & & \\ \hline 4 & 5 & 5 & 5 & & \\ \hline 5 & & & & & \\ \hline \end{array}, \quad Q^* = \begin{array}{|c|c|c|c|c|c|} \hline 1 & 2 & 2 & 2 & 2 & 4 \\ \hline 3 & 3 & 4 & 4 & & \\ \hline 4 & 4 & 5 & 6 & & \\ \hline 5 & & & & & \\ \hline \end{array} = Q'.$$

This property is readily derived from the well known property of the Robinson-Schensted correspondence that column inserting permutation entries in reverse order gives the same P -symbol as ordinary row insertion, and the Schützenberger dual Q -symbol. Note that reversing the order means vertically reflecting the permutation matrix; to get the same relation between the final matrices for the RSK and Burge correspondences, the rules for multiplexing rows and columns must also be reversed.

In spite of this relation between the global correspondences, and the formal similarity between their definitions, the shape data for the RSK and Burge correspondences have quite different characteristics. Contrary to what we saw for the former shape datum, the partial Schensted-growth that defines the latter can have multiple “bumping” configurations 2.1.1(4) in the same row or column of the grid. In fact, when a horizontal strip contains squares in some range of consecutive columns, any insertion or extraction step that bumps one of these squares will go on bumping the other squares until an unoccupied column is reached (or possibly the left edge of the diagram in case of extraction). By studying the effect of successive insertions one can deduce a description of the new shape datum in terms of occupancy of columns; we shall give this description, but we omit the details of the reasoning since the description will anyway serve merely as a motivation for or the construction that is our final goal.

If horizontal strips μ/λ and ν/λ are given as well as a matrix entry a , the corresponding shape κ for the Burge shape datum can be found as follows. We shall treat a as a modifiable variable, and traverse the columns from left to right; after visiting column j we shall have determined the length κ_j^t of that column in the diagram κ . When visiting column j we first set $c = (\mu_j^t - \lambda_j^t) + (\nu_j^t - \lambda_j^t) \in \{0, 1, 2\}$, which counts the occurrences in the two given horizontal strips of a square in column j . If $c = 1$ we have $\kappa_j^t = \lambda_j^t + 1 = \max\{\mu_j^t, \nu_j^t\}$, and we continue to the next column without further action. If $c = 2$, then $\kappa_j^t = \lambda_j^t + 1 = \mu_j^t = \nu_j^t$, and we increase the value of a by 1 before continuing to the next column. In the final case $c = 0$ we have $\lambda_j^t = \mu_j^t = \nu_j^t$; in this case if $a > 0$ then $\kappa_j^t = \lambda_j^t + 1$ and we decrease the value of a by 1, but if $a = 0$ we put $\kappa_j^t = \lambda_j^t$ and we continue to the next column without changing a . This procedure can be continued indefinitely, but there is no need to go on once a zero-length column of κ is found. Note that $\kappa_j^t - \lambda_j^t \in \{0, 1\}$ for all j , which means that $\lambda \leftarrow \kappa$.

As an example, in the bracketed square of the Knuth-growth of figure 3 one finds for c successively the values $0, 1, 0, 0, \dots$ as $j = 0, 1, 2, 3, \dots$ while $a = 2$ initially; it follows from the above description that $\kappa_j^t = \lambda_j^t + 1$ for $j = 0, 1, 2$ while $\kappa_j^t = \lambda_j^t$ for all other values of j , which can be checked in the diagram. For a case where a also increases,

consider the square with $\lambda = \square\square$ in the row labelled 4 and the column labelled 2, which also has $\mu = \nu = \square\square\square$, and $a = 1$. Then for $j = 0, 1, 2, 3$ one finds $c = 0, 0, 2, 0$, so that a decreases to 0 for $j = 0$, stays so for $j = 1$, raises again to 1 for $j = 2$ and finally drops back to 0 for $j = 3$; this causes $\kappa_j^t = \lambda_j^t + 1$ for $j = 0, 2, 3$ and $\kappa_j^t = \lambda_j^t$ for other values of j , whence $\kappa = \square\square\square\square$.

It may be checked independently from the relation with Schensted insertion that this description defines a shape datum. The variable a ultimately becomes $a = 0$, since $c = 0$ ultimately, after which a must decrease until it is 0. Then one easily sees that (6) holds (for the initial value of a of course). One never has $\kappa_j^t > \kappa_{j-1}^t$ for $j > 0$ (a column of κ made longer than its predecessor), since this requires $\kappa_j^t > \lambda_j^t = \mu_j^t = \nu_j^t$, and therefore $a > 0$ when column j is visited, but if visiting column $j - 1$ leaves $a > 0$ one always has $\kappa_{j-1}^t = \lambda_{j-1}^t + 1$. Finally, an inverse operation (which is the actual direction of the shape datum b^{Bu}) is easily defined in the same terms. The variable a will assume the same sequence of values as before, but of course in the opposite order; in particular it starts at 0 and its final value gives the number a to be determined. For the number c , which is now determined by the horizontal strips κ/μ and κ/ν , the values are not the same as for the forward direction, but they can be read off in all specified cases of the forward operation. They can be checked to indeed always produce the inverse change of the variable a and reconstruct the column length λ_j^t , thus guaranteeing a step-by-step inverse procedure.

Our specific interest in the shape datum for the Burge correspondence is due to two properties: the first is that it admits a description that treats the squares collectively rather than one at a time, the second is that this description can be easily stated in terms of edge sequences. To demonstrate the second point, whose importance will become clear when ribbons are considered, we shall use (in addition to the variable a) a modifiable doubly infinite bit sequence w . For the insertion direction that we shall describe, it starts as $w = \delta(\lambda)$ and ends as $w = \delta(\kappa)$. We shall traverse w from left to right, considering a pair of adjacent bits at the time; such a pair w_{k-1}, w_k corresponds to edges crossing the diagonal $d_k = \{(i, j) \in \mathbf{N}^2 \mid j - i = k\}$, and as done above for columns, we count the occurrences in μ/λ and ν/λ of a square in the diagonal d_k , calling the resulting number $c_k \in \{0, 1, 2\}$.

For each index k we proceed as follows. If $(w_{k-1}, w_k) \neq (1, 0)$, we move to $k + 1$ without further action. If $(w_{k-1}, w_k) = (1, 0)$ and $c_k = 1$, then we set $(w_{k-1}, w_k) := (0, 1)$ before proceeding to $k + 1$. If $(w_{k-1}, w_k) = (1, 0)$ and $c_k = 2$, then we also set $(w_{k-1}, w_k) := (0, 1)$, but in addition $a := a + 1$. If $(w_{k-1}, w_k) = (1, 0)$ and $c_k = 0$ (the case where a square of d_k can be added in κ that is not already in μ or ν), the action depends on the current value of a : if $a = 0$ nothing happens, but if $a > 0$ we set $(w_{k-1}, w_k) := (0, 1)$ and $a := a - 1$. For this procedure to work properly it must traverse a sufficiently large range of diagonals; we may start with the smallest k for which $(w_{k-1}, w_k) = (1, 0)$ initially (i.e., with $w = \delta(\lambda)$; in fact this gives $k = -\lambda_0^t$), and we may stop once $a = 0$ is obtained after treating the last diagonal with $c_k > 0$. Here is the calculation for the two examples cited above.

3.3 An r -correspondence with $r > 1$, and a corresponding shape datum

k	w before	μ/λ	ν/λ	c_k	action	w after	κ/μ	κ/ν
	$\delta(\lambda) =$				$a = 2$			
-3	$\cdot\cdot 1[10]10001010\cdot\cdot$			0	$a := a - 1 = 1$, swap	$\cdot\cdot 1[01]10001010\cdot\cdot$	\square	\square
-2	$\cdot\cdot 10[11]0001010\cdot\cdot$			0	none	$\cdot\cdot 10[11]0001010\cdot\cdot$		
-1	$\cdot\cdot 101[10]001010\cdot\cdot$	\square		1	swap	$\cdot\cdot 101[01]001010\cdot\cdot$	\square	
0	$\cdot\cdot 1010[10]01010\cdot\cdot$			0	$a := a - 1 = 0$, swap terminate	$\cdot\cdot 1010[01]01010\cdot\cdot$ $= \delta(\kappa)$	\square	\square
	$\delta(\lambda) =$				$a = 1$			
-1	$\cdot\cdot 1[10]01000\cdot\cdot$			0	$a := a - 1 = 0$, swap	$\cdot\cdot 1[01]01000\cdot\cdot$	\square	\square
0	$\cdot\cdot 10[10]1000\cdot\cdot$			0	none ($a = 0$)	$\cdot\cdot 10[10]1000\cdot\cdot$		
1	$\cdot\cdot 101[01]000\cdot\cdot$			0	none	$\cdot\cdot 101[01]000\cdot\cdot$		
2	$\cdot\cdot 1010[10]00\cdot\cdot$	\square	\square	2	$a := a + 1 = 1$, swap	$\cdot\cdot 1010[01]00\cdot\cdot$		
3	$\cdot\cdot 10100[10]0\cdot\cdot$			0	$a := a - 1 = 0$, swap terminate	$\cdot\cdot 10100[01]0\cdot\cdot$ $= \delta(\kappa)$	\square	\square

To show that the inverse of this procedure defines a shape datum requires a bit of effort, but is not too difficult. To see that $\mu \cup \nu \subseteq \kappa$ requires proving (by induction on k) that $c_k > 0$ implies that $(w_{k-1}, w_k) = (1, 0)$ holds when this pair of bits is considered. Termination with $a = 0$, and equation (6) for the initial value of a , are then proved like for the column description. Then one proves (again by induction on k) that cases with $w_{k-1} = 0$ can be encountered only when $a = 0$ holds. This fact allows an inverse procedure to be formulated; its description is in fact very similar to the one given above.

For the RSK correspondence a similar description of the shape datum is possible, but would be considerably messier than this one; this is essentially due to the fact that squares are bumped in the direction opposite to the traversal of the horizontal strips (namely towards the bottom left).

3.3. An r -correspondence with $r > 1$, and a corresponding shape datum.

So far we have only shown examples of 1-correspondences, and related shape data. Our goal however are Knuth correspondences whose shape data restrict to r -correspondences with $r > 1$. Before discussing the ones that really interest us, let us treat a construction that builds such correspondences in a fairly trivial way. For these correspondences the natural replacement for the graded set $S = \mathbf{N}$ with generating series $F_{\mathbf{N}}(T) = \sum_{i \in \mathbf{N}} X^i = \frac{1}{1-T}$ will be $S = \mathbf{N}^r$, graded by $|a| = \sum_{i \in [r]} a_i$, with generating series

$$F_{\mathbf{N}^r}(T) = F_{\mathbf{N}}(T)^r = \frac{1}{(1-T)^r} = \sum_{i \in \mathbf{N}} \binom{r}{i} T^i \quad (16)$$

where $\binom{r}{i} = \#\{a \in \mathbf{N}^r \mid |a| = i\} = \binom{r+i-1}{i}$.

The simplest example of an r -differential poset is \mathbf{Y}^r , the set of r -tuples of Young diagrams, partially ordered by inclusion of each of the r components separately. Like \mathbf{Y} this is a distributive lattice, and to check equation (2), one observes that any element covering a given r -tuple of shapes is obtained by adding a square to one of the component

3.3 An r -correspondence with $r > 1$, and a corresponding shape datum

shapes, while keeping the rest fixed; since for each component there is one more shape covering it than there are covered by it, the number of elements covering a given r -tuple exceeds the number of elements covered by it by r . In fact it is easy to see that the Cartesian product of an r -differential poset and an s -differential poset always gives an $r + s$ -differential poset.

It is not difficult either to define an r -correspondence for this situation. One starts by choosing a 1-correspondence for (\mathbf{Y}, \prec) that is to be used in the individual components; it will be denoted by $(b_\lambda)_{\lambda \in \mathbf{Y}}$ (it would be perfectly legal to make different choices for each component, but such frivolity would only complicate notation). If $\vec{\lambda} = (\lambda^i)_{i \in [r]}$ is an r -tuple of shapes covered by another such r -tuple $\vec{\mu} = (\mu^i)_{i \in [r]}$, there is a unique index i for which $\lambda^i \neq \mu^i$, and one has $\lambda^i \prec \mu^i$. Then if $b_{\lambda^i}(\mu^i) = e_0$ one defines $b_{\vec{\lambda}}(\vec{\mu}) = e_i \in e_{[r]}$, while in other cases $b_{\vec{\lambda}}(\vec{\mu})$ is defined by replacing μ^i in $\vec{\mu}$ by $b_{\lambda^i}(\mu^i) \in \mathbf{Y}$.

If $\begin{pmatrix} \vec{\lambda} & \vec{\mu} \\ \vec{\nu} & \vec{\kappa} \end{pmatrix}$ is a grid square of a Schensted-growth for such a r -correspondence, with matrix entry 0, then either $\vec{\lambda}$ and $\vec{\mu}$ differ in the same component as $\vec{\nu}$ and $\vec{\kappa}$, or there is equality in both cases; a similar relation holds for $\vec{\nu}/\vec{\lambda}$ and $\vec{\kappa}/\vec{\mu}$. A grid square with matrix entry e_i introduces a difference in component i across both its row and column of the grid, and by the above property these differences will propagate along row and column into steps of the Q - and P -symbols that will still involve a change in component i . Consequently, the Schensted correspondence may be computed in each of the components separately, as follows: for each $i \in [r]$ the positions of the entries e_i determine on one hand the sets of steps in the P - and Q -symbols that will involve a change in component i , and on the other hand a permutation matrix (by removing all other rows and columns, and replacing each e_i by 1). For each i the Schensted correspondence for \mathbf{Y} then defines a pair of paths (P, Q) in \mathbf{Y} , which r pairs can be spliced together into pairs of paths in \mathbf{Y}^r by taking at each step the next change in the appropriate component i .

The same idea also works to create new Knuth correspondences. On \mathbf{Y}^r one defines $\vec{\lambda} \leftarrow \vec{\mu}$ to mean $\lambda^i \leftarrow \mu^i$ for all $i \in [r]$, and for S one takes the graded set \mathbf{N}^r . Then a shape datum can be defined by components: $b_{\vec{\mu}, \vec{\nu}}(\vec{\kappa}) = (\vec{a}, \vec{\lambda})$ where $(a_i, \lambda^i) = b_{\mu^i, \nu^i}(\kappa^i)$ for all $i \in [r]$; it is trivial to verify the conditions of definition 2.2.1. The Knuth correspondence for this shape datum operates independently in each component even more evidently than the Schensted correspondence above: matrices with entries in $S = \mathbf{N}^r$ can be viewed as r -tuples of matrices with entries in \mathbf{N} , and paths in $(\mathbf{Y}^r, \leftarrow)$ as r -tuples of paths in (\mathbf{Y}, \leftarrow) (i.e., of semistandard tableaux); the Knuth-growth is defined with no interaction whatsoever between different components, so one may compute Knuth-growths separately for each component.

§4. Ribbons, and the Shimozono-White r -correspondence.

4.1. Ribbons, edge sequences, r -cores and r -quotients.

Partitions, partially ordered by repeated removal of r -ribbons (also called r -rim hooks), provide a less artificial example of a poset that satisfies the commutation relation (1) for some $r > 1$. The relation $\lambda \prec_r \mu$ between Young diagrams is said to hold if $\lambda \subset \mu$, and if the diagram of μ/λ consists of one square on each of r consecutive diagonals (in other words: that skew diagram is connected, has r squares, and contains no 2×2 blocks). In this case λ is said to be obtained by removing an r -ribbon from μ , and μ by adding an r -ribbon to λ . The reflexive transitive closure of ' \prec_r ' defines a partial ordering ' \leq_r ' on \mathbf{Y} .

In order to make (\mathbf{Y}, \prec_r) into a graded graph as discussed in §2.1, the grading on \mathbf{Y} must be adapted so that λ and μ have consecutive ranks when $\lambda \prec_r \mu$. This is easily done by defining a grading $|\lambda|_r = \lfloor \frac{|\lambda|}{r} \rfloor$, the quotient of the Euclidean division of $|\lambda|$ by r . The rank of individual shapes is of no importance, since only differences in rank between comparable elements in (\mathbf{Y}, \leq_r) are used; therefore we put $|\mu/\lambda|_r = |\mu|_r - |\lambda|_r = \frac{|\lambda| - |\mu|}{r}$ whenever $\lambda \leq_r \mu$. Paths of shape μ/λ in (\mathbf{Y}, \prec_r) are called standard r -ribbon tableaux of shape μ/λ ; such tableaux have $|\mu/\lambda|_r$ individual r -ribbons.

The operation of adding or removing an r -ribbon is best understood in terms of the edge sequences $\delta(\lambda)$ and $\delta(\mu)$. We shall discuss these matters here summarily, referring to [vLee2] for an extensive discussion and examples. When $\lambda \prec_r \mu$, the edge sequences $\delta(\lambda)$ and $\delta(\mu)$ differ only in two places, which are at distance r , and at those places $\delta(\lambda)$ has $\dots 1 \dots 0 \dots$ while $\delta(\mu)$ has $\dots 0 \dots 1 \dots$ (the dots represent unchanged bits). Hence addition or removal of a single r -ribbon affects bits of the edge sequence whose positions are in the same congruence class modulo r ; we shall say two r -ribbons are in the same position class if the positions of the bits affected by their addition or removal are congruent modulo r . As operations on edge sequences, the addition or removal of r -ribbons in distinct position classes always commute (although the ribbons themselves, viewed as skew diagrams, may change). The bits with positions in a fixed congruence class form, up to a shift, the edge sequence of a unique Young diagram, which changes by a single square for any modification by an r -ribbon in the corresponding position class. Thus the induced sub-poset of (\mathbf{Y}, \leq_r) on the subset of shapes reachable from a given one by addition or removal of r -ribbons in a single position class is isomorphic to $\mathbf{Y} = (\mathbf{Y}, \subseteq)$. It follows that any connected component of (\mathbf{Y}, \leq_r) is isomorphic to \mathbf{Y}^r . In particular each such component has a unique minimal element for ' \leq_r ', which is called an r -core; γ is an r -core if and only if each of the r sequences extracted from $\delta(\gamma)$ by selecting a congruence class modulo r of bit positions is of the form $\dots 111000 \dots$ (these extracted sequences may be shifted with respect to each other). Any $\lambda \in \mathbf{Y}$ is uniquely determined by the minimal element in its connected component of (\mathbf{Y}, \leq_r) , called the r -core of λ , together with its image in \mathbf{Y}^r under the isomorphism, an r -tuple of partitions called the r -quotient of λ .

4.2. Rim hook lattices and semistandard r -ribbon tableaux.

Being isomorphic to \mathbf{Y}^r , any connected component of (\mathbf{Y}, \leq_r) is clearly a distributive

lattice; it is called an r -rim hook lattice. As all r -rim hook lattices are isomorphic as posets, one often considers only the one containing the empty partition as r -core, but there is no good reason to do so: if one were only interested in the abstract poset structure, one could study \mathbf{Y}^r , and it would be useless to introduce ribbons in the first place. Meanwhile the isomorphism, which is called the r -quotient map, makes it immediately clear that r -rim hook lattices are r -differential posets, and it translates any choice of an r -correspondence for \mathbf{Y}^r into such a choice for any r -rim hook lattice. In the case of the r -correspondence for \mathbf{Y}^r of §3.3, using in each component \mathbf{Y} the usual row-insertion (respectively column-insertion) 1-correspondence, the resulting r -correspondence for (\mathbf{Y}, \leq_r) can be described directly as follows. Let $\mu \prec_r \kappa$, then $b_\mu(\kappa) = \lambda \prec_r \mu$ is such that μ/λ is the first r -ribbon, if it exists, that is removable from μ , in the same position class as κ/μ and to its top right (respectively to its bottom left). If no such ribbon exists, $b_\mu(\kappa) = e_i$ where $i \in [r]$ represents the position class of κ/μ : if the top right square of the ribbon, which we shall call its head, lies on diagonal d_k then $i = k \bmod r$. The validity of this description rests on the fact that the r -ribbons in a single position class that can be added to respectively removed from μ , are perfectly interleaved in bottom left to top right order, with at both extremes ribbons that can be added; this can be seen directly by the same argument as given in §2.1 for the 1-correspondence for \mathbf{Y} , but applied to the sequence extracted from $\delta(\mu)$ of bits at positions congruent to i modulo r . Note that for these r -correspondences in (\mathbf{Y}, \leq_r) the names row- and column-insertion would not be very appropriate, since the ribbons κ/μ and μ/λ might be separated by any number of rows and columns in either case.

The r -quotient maps can of course also be used to transport the relation ' \leftarrow ' and the shape data for it from \mathbf{Y}^r to (\mathbf{Y}, \leq_r) . We shall define $\lambda \leftarrow_r \mu$ when $\lambda \leq_r \mu$ and moreover $\lambda^i \leftarrow \mu^i$ for each pair of corresponding components λ^i, μ^i of their respective r -quotients. To better understand this relation, note that $\lambda \leftarrow \mu$ means that $\delta(\lambda)$ can be transformed into $\delta(\mu)$ by a sequence of replacements of a substring '10' by '01', proceeding from left to right with overlap allowed (the bit 1 of the replacement may participate in the next replacement). This fact, which we already saw implicitly in the description of the Burge shape datum by edge sequences, follows from the simple observation that the sequence of squares added in a skew standard tableau has strictly increasing column numbers if and only if it has increasing diagonal numbers (here strictness is for free). In this description, the leftmost bit of a replacement string cannot be modified afterwards, so while considering from left to right occurrences of '10' for replacement, the decision whether to or not replace is prescribed by the target $\delta(\mu)$ (if a possibility exists at all); therefore this description gives a direct, backtrack-free, procedure to decide whether or not $\lambda \leftarrow \mu$.

To decide whether $\lambda \leftarrow_r \mu$ holds, we must make a similar test for each pair of sequences of bits extracted from λ and μ at positions in the same conjugacy class modulo r . Although the r traversals required for these tests are independent, we might as well combine them into a single left to right pass over the edge sequences. Thus we find that $\lambda \leftarrow_r \mu$ if and only if $\delta(\lambda)$ can be transformed into $\delta(\mu)$ by a left to right sequence of replacements of a substring '1x0' by '0x1', where $x \in \{0, 1\}^{r-1}$ is any string of $r - 1$ bits unaffected by the replacement, and with overlap between successive

4.3 Ribbon Schensted and Knuth correspondences that factor

replacements allowed. Each intermediate bit string occurring during this transformation is edge sequence of a Young diagram, and their sequence defines a standard r -ribbon tableau of shape μ/λ ; the skew shape μ/λ is called a *horizontal r -ribbon strip*, and the standard r -ribbon tableau its standardisation. In geometric terms, the standardisation of a horizontal r -ribbon strip μ/λ gives the unique decomposition of its diagram into a sequence of r -ribbons such that the head of each ribbon has its top edge on the inner border of μ/λ (i.e., on the boundary of λ), or equivalently such that the tail (bottom left square) of each ribbon has its bottom edge on the outer border of μ/λ (i.e., on the boundary of μ). A monotonically rising path in $(\mathbf{Y}, \leftarrow_r)$ from λ to μ is called a *semistandard r -ribbon tableau* of shape μ/λ . Here are graphic representations of the standardisation of a horizontal 5-ribbon strip, and of a semistandard 6-ribbon tableau.



Although we used the r -quotient map to define ' \leftarrow_r ', this is not evident from the final description; usually horizontal r -ribbon strips and semistandard r -ribbon tableaux are defined without using r -quotients.

4.3. Ribbon Schensted and Knuth correspondences that factor.

Fomin's constructions reduce the question of defining Schensted and Knuth correspondences for ribbon tableaux to the question of defining r -correspondences for (\mathbf{Y}, \leq_r) , respectively of defining shape data for $(\mathbf{Y}, \leftarrow_r, \mathbf{N}^r)$. The r -quotient map provides an easy way to do this, as we already mentioned for r -correspondences. We note that the Schensted correspondence obtained from the r -correspondence based on the row insertion in each of the r components was originally defined by a direct construction in [StWh], and was later found to factor via the r -quotient map. To define a shape datum b^r for $(\mathbf{Y}, \leftarrow_r, \mathbf{N}^r)$ one may proceed similarly, leading to the following description. Whenever $(a, \lambda) = b_{\mu, \nu}^r(\kappa)$ is to be defined, each of μ, ν, κ has the same r -core γ ; denoting by μ^i, ν^i, κ^i the components of their respective r -quotients, λ will be the partition with r -core γ , and with the components of its r -quotient defined, together with the components of a , by $(a_i, \lambda^i) = b_{\mu^i, \nu^i}(\kappa^i)$ for all $i \in [r]$, where one fixes either $b = b^{\text{Kn}}$ or $b = b^{\text{Bu}}$.

In this construction of shape data, the standardisations of the horizontal r -ribbon strips that arise are not used at all. We know these standardisations exist, since within a r -rim hook lattice $\lambda \leftarrow_r \mu$ just means that $\lambda^i \leftarrow \mu^i$ for all i , but to find the standardisation of a given horizontal r -ribbon strip amounts to carefully ordering the individual squares of the r horizontal strips contributing to it, according to the placements of the associated edges within the edge sequence merged together from r individual ones. Since the Knuth correspondences used in the components were derived via standardisation from Schensted correspondences, one might assume that these shape data, and

the derived Knuth correspondences, can also be obtained by combining the Stanton and White correspondence with standardisation of horizontal r -ribbon strips; indeed one can for instance find an informal statement to this effect in [ShW2, §5]. Things do not work out that nicely however; in particular, one does not obtain the proper standardisations directly from the Stanton and White algorithm, even for the half-semistandard case (with the Q -symbol standard). To illustrate this point, we shall briefly digress to study this Knuth correspondence for ribbon tableaux (whose existence, based on arguments like those above, has been mentioned occasionally, but which does not appear to have been explicitly described anywhere) in some more detail.

The easiest way to define a correspondence based on standardisation is to start with semistandard r -ribbon tableaux P, Q , and (just like we did for the RSK correspondence) interpolate into standard tableaux, compute the Schensted-growth for the Stanton and White r -correspondence in the extraction direction, and afterwards reduce the grid again, summing coloured permutation matrix entries over rectangular areas that collapse to a square, to find the matrix entries $a \in \mathbf{N}^r$ of the final result (addition is done for each colour separately). This produces the same Knuth-growth and hence the same matrix as the Knuth correspondence for the shape datum b^r defined above (with $b = b^{\mathbf{Kn}}$), since both standardisation and application of the r -correspondence commute with the decomposition of shapes via the r -quotient map. More precisely, decomposing the shapes in the standardisation of a horizontal r -ribbon strip produces a sequence of r -tuples of shapes, in which only one component changes at each step; to extract from this an r -tuple of skew standard tableaux, one must eliminate in each component the steps where the shape does not change, but then the resulting tableaux are indeed the standardisations of the r horizontal strips obtained from the decomposition of the original horizontal r -ribbon strip via the r -quotient map.

However, while the correct shapes are assigned to the grid points that remain after reduction, the interpolations between them produced by the Schensted growth are not the standardisations of the respective horizontal r -ribbon strips, and as a consequence the coloured permutation entries are not placed in any predictable manner within their rectangle of the Schensted-growth; this means that it is impossible to reconstruct the same Schensted-growth from the matrix found by an inverse (insertion) procedure. The reason for this is that the relative order of ribbons from different position classes contributing to a horizontal r -ribbon strip can freely change during the extraction process, since the rules of the growth treat the r classes of positions of ribbons completely independently. The only thing that is guaranteed is that ribbons within one position class keep their relative order, so that under the r -quotient map one gets standardisations of horizontal strips. Here are some very simple examples of what can happen, with $r = 2$ and without non-zero matrix entries. The first one uses the Stanton and White r -correspondence, the second one uses the transpose r -correspondence. Note that the top rows are not the standardisations of their respective horizontal r -ribbon strips, while the bottom rows are such standardisations.

4.4 Form and height of ribbons



Note that in both cases one ribbon keeps its head on the diagonal, while the other is bumped, moving two diagonals and thereby inverting the relative order of the two ribbons. Such behaviour occurs during insertion as well as extraction, so regardless of the rule one would choose to insert the different contributions of a single matrix entry into the Schensted-growth, it is impossible to achieve that the sequences interpolating horizontal r -ribbon strips remain their standardisations throughout the insertion process. All in all, there is no reason that would justify using a Schensted-growth instead of computing the Knuth correspondence directly via the r -quotient map.

There is one more point that is worth noting about the shape datum b^r of this Knuth correspondence obtained from $b = b^{\text{Bu}}$, the shape datum of the Burge correspondence. In this case the description of b^{Bu} in terms of edge sequences given in §3.2 can be easily adapted to the ribbon setting, merging r independent traversals of the edge sequence into one in the same manner as for the definition of horizontal r -ribbon strips. Although it is still a parallelised correspondence in disguise, we sketch it here because of its similarity in form to the shape datum that we shall present later.

As before we consider insertion, so let shapes $\mu \rightarrow_r \lambda \leftarrow_r \nu$ and a matrix entry $a \in \mathbf{N}^r$ be given; also, k shall be the index of a diagonal, for which this time we count the number $c_k \in \{0, 1, 2\}$ of r -ribbons in the standardisations of μ/λ and of ν/λ whose head lies on the diagonal d_k . A doubly infinite bit string w is again initialised to $\delta(\lambda)$, but instead of inspecting pairs (w_{k-1}, w_k) of adjacent bits, one inspects the pairs (w_{k-r}, w_k) that would be affected by an r -ribbon with its head on d_k . The action taken for k may swap these two bits and may change the component a_i of a where $i = k \bmod r$, according to the same case distinction as before; after this processing the value of k is increased by 1. The range of values that must be traversed by k is determined in a similar manner as before, where we must now detail that all components of a must have become 0 before termination can be decided.

A verification that (the inverse of) this procedure defines a shape datum for $(\mathbf{Y}, \leftarrow_r, \mathbf{N}^r)$ can be given along the same lines as indicated at the end of §3.2. Since w_k , the rightmost of the bits considered and possibly modified, is put aside for some time before being reconsidered, the condition that enables the formulation of an inverse procedure now read as follows: after treating the pair (w_{k-r}, w_k) , one has for all i with $k - r < i \leq k$ that $w_i = 0$ implies $a_{i \bmod r} = 0$. Surprisingly this algorithm, which is rather trivial because it factors into r copies simultaneously computing the shape datum for the Burge correspondence, can be transformed into the spin preserving shape correspondence that is the main subject of this paper by just one change: the selection $i = k \bmod r$ at each step of the component of a to possibly modify is replaced by $i = \sum_{j=1}^{r-1} w_{k-j}$. Proving the existence of an inverse will be a bit harder though.

4.4. Form and height of ribbons.

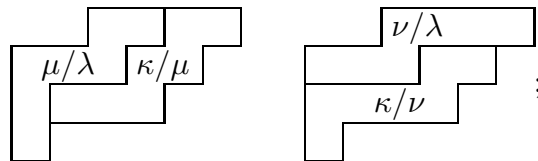
From the above consideration we learn that from the enumerative point of view

$(\mathbf{Y}, \leftarrow_r, \mathbf{N}^r)$ allows the existence of Knuth correspondences, but the ones found so far are not very interesting because they avoid really dealing with ribbons by directly applying the r -quotient map. The key to defining interesting alternative correspondences is to focus on a property of r -ribbons that is not related to the r -quotient maps. One such property is their form, where two ribbons are considered to have the same form if their diagrams are equal up to a translation. There are 2^{r-1} different forms of r -ribbons. If $\delta(\lambda)$ can be transformed into $\delta(\mu)$ by the replacement of a substring $1x0$ by $0x1$ with $x \in \{0, 1\}^{r-1}$, then x describes the form of the r -ribbon μ/λ by telling for each following square whether it is above or to the right of its predecessor; we therefore define $\text{form}(\mu/\lambda) = x$. Now an important statistic on ribbons is their height, where $\text{ht}(\mu/\lambda) \in [r]$ is defined as the sum of the bits in $\text{form}(\mu/\lambda)$. This is the number of vertical steps encountered when going from the tail of μ/λ to its head, or (since even ribbons without vertical steps occupy one row) one less than the number of rows that meet the diagram of μ/λ .

The r -quotient map is not very well suited for studying the height of ribbons, since the bits that contribute to the height of a ribbon are found only in the components of the r -quotient distinct from the one to which the ribbon contributes a square. In fact, although r -core and r -quotient together preserve complete information about a shape, the r -quotients of λ and μ alone do not suffice to determine $\text{ht}(\mu/\lambda)$, since it cannot be determined *which* bits of the other components should be added (the edge sequences of the components of the r -quotient are shifted by amounts determined by the r -core). Consequently, there is very little one can say about the heights of the ribbons related by the r -correspondences for (\mathbf{Y}, \prec_r) and shape data for $(\mathbf{Y}, \leftarrow_r, \mathbf{N}^r)$ discussed above. We shall presently study correspondences that do respect the heights of ribbons; it is these that will justify our interest in r -rim hook lattices. Let us start however by studying how the height of ribbons behaves in the cases encountered in Schensted correspondences that are independent of any particular choice of r -correspondence.

Consider shapes $\begin{pmatrix} \lambda & \mu \\ \nu & \kappa \end{pmatrix}$ with $\lambda \prec_r \mu, \nu \prec_r \kappa$ and $\mu \neq \nu$, in other words κ is the unique element covering μ and ν in (\mathbf{Y}, \prec_r) , and λ is the unique element covered by them. Then the r -ribbons μ/λ and κ/ν will have their head on the same diagonal (their heads may even coincide), as will ν/λ and κ/μ . If these diagonals are respectively d_s and d_t , then $s - t$ cannot lie in $\{-r, 0, r\}$ (one can never successively add two ribbons with their heads on the same diagonal, and $|s - t| = r$ would imply that there is only one shape strictly between λ and κ in (\mathbf{Y}, \leq_r) , contradicting $\mu \neq \nu$). If $|s - t| > r$, then $\text{ht}(\mu/\lambda) = \text{ht}(\kappa/\nu)$ and $\text{ht}(\nu/\lambda) = \text{ht}(\kappa/\mu)$, since the ribbons are separated by at least one diagonal, and the indicated pairs of ribbons have identical diagrams. On the other hand, if $|s - t| < r$, then both equalities of height fail, since the fact of having added a ribbon with its head on diagonal s to λ before forming a ribbon with its head on diagonal t will have changed one bit in the form of the latter ribbon. In more detail: if $-r < s - t < 0$ (κ/μ lies further to the top right than μ/λ) then $\text{ht}(\kappa/\mu) = \text{ht}(\nu/\lambda) + 1$ and $\text{ht}(\kappa/\nu) = \text{ht}(\mu/\lambda) - 1$:

4.5 The Shimozono-White r -correspondence



similarly, if $0 < s - t < r$ then $\text{ht}(\kappa/\mu) = \text{ht}(\nu/\lambda) - 1$ and $\text{ht}(\kappa/\nu) = \text{ht}(\mu/\lambda) + 1$. Note that in all cases $\text{ht}(\kappa/\mu) + \text{ht}(\kappa/\nu) = \text{ht}(\nu/\lambda) + \text{ht}(\mu/\lambda)$. In the Schensted correspondence the ribbons at one side of this equation are computed from those at the other side, so this amounts to a conservation of total height.

4.5. The Shimozono-White r -correspondence.

Now we consider r -correspondences that respect the height of ribbons. For any $\mu \in \mathbf{Y}$, we know that the set of shapes covering μ in (\mathbf{Y}, \prec_r) has r more elements than the set of shapes covered by μ . A key observation, made in [ShW2], is that these sets, labelled by the height of the r -ribbon involved in the covering relation, have a very regular structure. There exists for instance, for every possible height $h \in [r]$, at least one $\kappa \succ_r \mu$ with $\text{ht}(\kappa/\mu) = h$, a fact that does not seem immediately obvious (except for $h = 0$ and $h = r - 1$). The heads of r -ribbons of the form μ/λ or κ/μ all lie on distinct diagonals, so there is a natural total ordering on the set of those ribbons. For any diagonal d_k , the value $\sum_{j=1}^{r-1} \delta(\mu)_{k-j}$ gives the height of the ribbon with its head on that diagonal *if it exists*, and that sum changes by steps at most 1 as k varies. The following somewhat surprising proposition states that, even though many diagonals have no associated ribbons, the heights of the ribbons that are present still change by steps at most 1 as the diagonals are traversed, with even some additional constraints.

4.5.1. Proposition. *Let $\mu \in \mathbf{Y}$, and let ξ_0, ξ_1 be r -ribbons, both either of the form μ/λ or κ/μ , such that the diagonals d_s, d_t containing their heads satisfy $s < t$, and no diagonal d_i with $s < i < t$ contains the head of any r -ribbon of the form μ/λ or κ/μ . Then one of the following cases applies:*

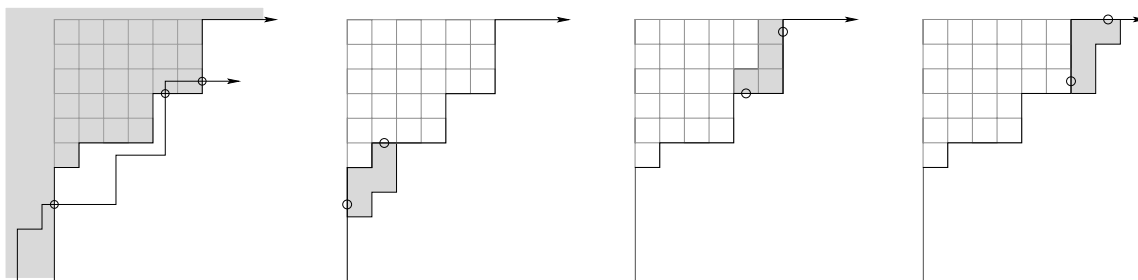
- (1) *both ξ_0 and ξ_1 are of the form κ/μ , and $\text{ht}(\xi_1) = \text{ht}(\xi_0) - 1$;*
- (2) *there is one ξ_i of each of the forms μ/λ and κ/μ , and $\text{ht}(\xi_1) = \text{ht}(\xi_0)$;*
- (3) *both ξ_0 and ξ_1 are of the form μ/λ , and $\text{ht}(\xi_1) = \text{ht}(\xi_0) + 1$.*

Moreover the r -ribbons ξ_{\ll}, ξ_{\gg} of one of the given forms with their heads on diagonals d_k with k minimal respectively maximal, are both of the form κ/μ , while $\text{ht}(\xi_{\ll}) = r - 1$ and $\text{ht}(\xi_{\gg}) = 0$.

Proof. Instead of considering sums of $r - 1$ consecutive bits of $\delta(\mu)$, consider sums of r consecutive bits. As we shift the range of summation upwards so as to include $\delta(\mu)_k$ while dropping $\delta(\mu)_{k-r}$, the sum changes if and only if there is an r -ribbon of one of the given forms with its head on diagonal d_k , with an increase of the sum for ribbons of the form μ/λ , and a decrease of the sum for ribbons of the form κ/μ ; in either case the smaller of the two sums gives the height of the ribbon. Considering two consecutive changes of the sum gives rise to the stated three cases. The final statement corresponds to the fact that the sum, which takes values in $\{0, \dots, r\}$, tends to r as $k \rightarrow -\infty$, and to 0 as $k \rightarrow +\infty$. \square

4.5.2. Corollary. *Let $\mu \in \mathbf{Y}$ and $h \in [r]$ be fixed. The r -ribbons of the form $\xi = \kappa/\mu$ with $\text{ht}(\xi) = h$ and those of the form $\xi = \mu/\lambda$ with $\text{ht}(\xi) = h$ are perfectly interleaved when ordered from bottom left to top right (by the diagonal of their head), with ribbons of the first kind at both ends. In particular, the number of ribbons of the first kind exceeds the number of those of the second kind by 1, exactly. \square*

The easiest way to understand this corollary is by the same considerations as in the proof of the proposition: each time the sum over r consecutive bits of $\delta(\mu)$ descends from the range $\{h + 1, \dots, r\}$ to its complement $\{0, \dots, h\}$, there is an r -ribbon of the form $\xi = \kappa/\mu$ with $\text{ht}(\xi) = h$; each time it rises back to $\{h + 1, \dots, r\}$ there is an r -ribbon of the form $\xi = \mu/\lambda$ with $\text{ht}(\xi) = h$. There is also a nice visual presentation of this proof. Draw the diagram of μ and its boundary, and superimpose a copy of the boundary shifted down by $h + \frac{1}{2}$ units and leftwards by $r - h - \frac{1}{2}$ units. The boundary and its copy can only cross in the middle of edge segments, and whenever they do, these two segments correspond to an r -ribbon of height h that can either be added to μ (when the copy passes the boundary from the “inside”, the side of the diagram, to the outside) or removed from μ (when the copy passes back to the inside). Clearly these types of crossings alternate as one traverses the boundary, and due to the direction of the shift, one eventually passes from the inside to the outside. We illustrate this process for $\mu = (6, 6, 6, 4, 4, 1)$, $r = 4$, and $h = 2$; for those values two ribbons can be added to μ , and one removed.



It is remarkable that whether we group the r -ribbons that can be added to or removed from μ by their position class or by their height, they will be interleaved within each group in exactly the same way. Like for position classes, this property allows us to define an r -correspondence for (\mathbf{Y}, \prec_r) in two ways (moving to the top right or to the bottom left). With the choice of the r -correspondence moving to the top right (recall that this direction refers to the extraction process), such an r -correspondence was first described in [ShW2, propositions 19,20], and we shall call it the Shimozono-White r -correspondence b^{SW} . For the corresponding Schensted insertion algorithm, bumping of r -ribbons is towards the bottom left, like bumping in the Robinson-Schensted algorithm. Explicitly, if $\lambda \prec_r \mu$ and $\text{ht}(\mu/\lambda) = h$, then $\lambda = b_\mu^{\text{SW}}(\kappa)$ where κ/μ is an r -ribbon of height h that can be added to μ , and the first among such ribbons to the bottom left of μ/λ . Also $e_h = b_\mu^{\text{SW}}(\kappa)$ where κ/μ is the top-rightmost r -ribbon of height h that can be added to μ . Corollary 4.5.2 guarantees that this properly describes a unique r -correspondence. The opposite choice, obtained by replacing “bottom left”

4.5 The Shimozono-White r -correspondence

by “top right” in the description and vice versa, is also possible; we shall call the r -correspondence so defined the transposed Shimozono-White r -correspondence, and denote it by b^{WS} (which is more convenient than $(b^{\text{SW}})^t$, although some might have preferred b^{NE}).

These r -correspondences, and the Schensted-growths defined by them, behave well with respect to height. Let $b = b^{\text{SW}}$ or $b = b^{\text{WS}}$, and consider a square $\begin{pmatrix} \lambda & \mu \\ \nu & \kappa \end{pmatrix}$ of a Schensted-growth for b , with matrix entry $a \in \{0\} \cup e_{[r]}$. We assume that $\kappa \notin \{\mu, \nu\}$ (for otherwise one has $\{\lambda, \kappa\} = \{\mu, \nu\}$ and $a = 0$, which causes our conclusions below to hold trivially). We have seen that if $\mu \neq \nu$, then one always has $\text{ht}(\kappa/\mu) + \text{ht}(\kappa/\nu) = \text{ht}(\nu/\lambda) + \text{ht}(\mu/\lambda)$ and $a = 0$. In the case that $\mu = \nu$ and $b_\mu(\kappa) = \lambda \in \mathbf{Y}$, one has the stronger set of conditions $\text{ht}(\kappa/\mu) = \text{ht}(\kappa/\nu) = \text{ht}(\nu/\lambda) = \text{ht}(\mu/\lambda)$ and $a = 0$. Finally if $\mu = \nu$ and $b_\mu(\kappa) = a \in e_{[r]}$, then $\lambda = \mu = \nu$ and $a = e_h$ where $h = \text{ht}(\kappa/\mu) = \text{ht}(\kappa/\nu)$. Therefore, with the conventions that $n(e_h) = h$ (the “colour” of the entry of a coloured permutation, which we are really interpreting as an integer here) and $\text{ht}(\lambda/\lambda) = n(0) = 0$, we may conclude that in all cases one has the identity $\text{ht}(\kappa/\mu) + \text{ht}(\kappa/\nu) = \text{ht}(\nu/\lambda) + \text{ht}(\mu/\lambda) + 2n(a)$. This can be seen as a “law of conservation”: if for lattice paths from $(0, n)$ to $(n, 0)$ in the grid of the Schensted-growth one takes the sum of the heights of all ribbons formed by adjacent shapes on the path, and adds to it twice the sum of the colours of all matrix entries of the Schensted-growth to the lower right of the path, then this sum will be invariant under moving the path across a square. Consequently, that sum is the same for all such paths. For the path going via the origin one gets twice the sum of the colours of all matrix entries, and for the path via the opposite corner one gets the sum of the heights of the ribbons in the P - and Q -symbols; these numbers are therefore equal. Since half the sum of the heights of the ribbons in a standard r -ribbon tableau is called its “spin” (see [CaLe] and [LLT]; the motivation for the terminology is not quite clear), the computation of the Schensted correspondence using b^{SW} can be called a “colour-to-spin Schensted algorithm”. We note in passing that since P, Q may be any pair of standard r -ribbon tableaux of the same shape, and the sum of the heights of their ribbons is always even, the spins of r -ribbon tableaux of a given shape must either all lie in \mathbf{N} or all lie in $\mathbf{N} + \frac{1}{2}$; this is a well known fact that is related to the notion of the r -sign of a shape.

To formulate the enumerative consequences of the existence of such height respecting correspondences, we give a “ q -analogue” of equation (1) that holds for the current situation. With q an indeterminate, we consider endomorphisms D_q^r, U_q^r of the free $\mathbf{Z}[q]$ -module $\mathbf{Z}[q]\mathbf{Y}$ on the set \mathbf{Y} , defined by their action on basis elements $U_q^r: \lambda \mapsto \sum_{\mu \succ_r \lambda} q^{\text{ht}(\mu/\lambda)} \mu$ and $D_q^r: \lambda \mapsto \sum_{\mu \prec_r \lambda} q^{\text{ht}(\lambda/\mu)} \mu$. Then corollary 4.5.2, together with the conservation of total height observed for ribbons related to shapes with a common successor and a common predecessor in (\mathbf{Y}, \prec_r) , leads to the identity

$$D_q^r \circ U_q^r = U_q^r \circ D_q^r + [r]_{q^2} \mathbf{1}, \quad \text{where } [r]_{q^2} = \sum_{i \in [r]} q^{2i} = \frac{1-q^{2r}}{1-q^2} \in \mathbf{Z}[q^2]. \quad (17)$$

From it one derives for any r -core γ (by writing $\langle \gamma \mid (D_q^r)^n ((U_q^r)^n(\gamma)) \rangle$ and expanding $(D_q^r)^n \circ (U_q^r)^n$ into terms $(U_q^r)^i \circ (D_q^r)^i$) the identity:

$$\langle (U_q^r)^n(\gamma) \mid (U_q^r)^n(\gamma) \rangle = n!([r]_{q^2})^n. \quad (18)$$

We denote by $\mathbf{RH}(\gamma)$ the connected component of (\mathbf{Y}, \leq_r) containing the r -core γ (an r -rim hook lattice), and by $F_{q^{1/2}}^{\lambda/\gamma} \in \mathbf{Z}[q^{\frac{1}{2}}]$ (using a square root $q^{\frac{1}{2}}$ of q) the spin generating series $\sum_T q^{\text{spin}(T)}$ of the set of standard r -ribbon tableaux T of shape λ/γ . We can then write (18) in a form that directly matches the colour-to-spin Schensted correspondence:

$$\sum_{\lambda \in \mathbf{RH}(\gamma)} (F_{q^{1/2}}^{\lambda/\gamma})^2 = n!([r]_q)^n = n! \left(\frac{1 - q^r}{1 - q} \right)^n, \tag{19}$$

whose right hand side can be interpreted as the sum over all r -coloured permutations σ of $q^{n(\sigma)}$, where $n(\sigma)$ is the sum of the colours of the entries of σ .

We have finally arrived at the point where we can state our main goal: to find a Knuth correspondence for $(\mathbf{Y}, \leftarrow_r, \mathbf{N}^r)$ that satisfies a similar conservation law involving heights. If one identifies $e_{[r]}$ with the “standard basis” of \mathbf{N}^r , this Knuth correspondence should in particular reduce to a “colour-to-spin” Schensted correspondence when all matrix entries a satisfy $|a| \leq 1$ (i.e., $a \in \{0\} \cup e_{[r]}$). To make our goal precise, one needs to define the spin of a horizontal r -ribbon strip, and the contribution to the spin of a matrix entry (for semistandard r -ribbon tableaux and complete matrices these quantities will then be defined by summing over their constituent parts). By definition $\text{spin}(\mu/\lambda)$ is equal to the spin of the standardisation of μ/λ (cf. [LLT]); note that it is here that the notion of standardisation of a horizontal r -ribbon strip becomes essential, since other interpolations into a standard r -ribbon tableau will in general exist, but these will have different spins (the standardisation in fact achieves the maximal possible spin, much to the satisfaction of the spin doctors). The contribution $n(a)$ of a matrix entry $a \in \mathbf{N}^r$ to the spin extends by linearity the relation $n(e_h) = h$, in other words it is $\sum_{i \in [r]} ia_i$ (one might call this the “total colour” of the entry a , but we are no longer dealing with coloured permutations, and we wish to drop the colourful terminology). So we want $(a, \lambda) = b_{\mu, \nu}(\kappa)$ to be defined in such a way that one always has

$$\text{spin}(\kappa/\mu) + \text{spin}(\kappa/\nu) = \text{spin}(\mu/\lambda) + \text{spin}(\nu/\lambda) + n(a), \quad \text{where } n(a) = \sum_{i \in [r]} ia_i. \tag{20}$$

To formulate an enumerative identity that is required for the existence of such a shape datum, let us introduce $q^{\frac{1}{2}}$ -analogues of the generating series U_X, D_Y :

$$U_{q^{1/2}, X}^r(\lambda) = \sum_{\mu \rightarrow_r \lambda} q^{\text{spin}(\mu/\lambda)} X^{|\mu/\lambda|_r} \mu, \quad D_{q^{1/2}, Y}^r(\lambda) = \sum_{\mu \leftarrow_r \lambda} q^{\text{spin}(\lambda/\mu)} Y^{|\lambda/\mu|_r} \mu, \tag{21}$$

(one has $U_{q^{1/2}, X}^r, D_{q^{1/2}, Y}^r \in \text{End}(\mathbf{Z}[q^{\frac{1}{2}}]\mathbf{Y})[[X, Y]]$), and a q -analogue of $F_{\mathbf{N}^r}$:

$$F_{\mathbf{N}^r}(q, T) = \sum_{a \in \mathbf{N}^r} q^{n(a)} T^{|a|} = \prod_{i \in [r]} \left(\sum_{a_i \in \mathbf{N}} q^{ia_i} T^{a_i} \right) = \prod_{i \in [r]} \frac{1}{1 - q^i T} \in \mathbf{Z}[q][[T]]. \tag{22}$$

Then the required enumerative identity (the one that will be bijectively proved by our main construction) is analogous to (9) with the expression for $F_{\mathbf{N}^r}(q, XY)$ substituted:

$$D_{q^{1/2}, Y}^r \circ U_{q^{1/2}, X}^r = (U_{q^{1/2}, X}^r \circ D_{q^{1/2}, Y}^r) \prod_{i \in [r]} \frac{1}{1 - q^i XY}. \tag{23}$$

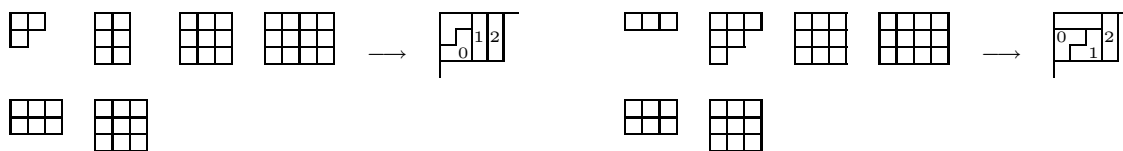
4.5 The Shimozono-White r -correspondence

Taking the coefficient of XY gives an identity equivalent to (17). What we have seen so far only proves that special case, but (23) is known to hold: it can be derived from the commutation relation [LLT, (21)] proved in [KMS], for operators B_i that satisfy $\exp(\sum_{i>0} \frac{B_{-i}}{i} X^i) = U_{-q^{-1}, X}^r$ and $\exp(\sum_{i>0} \frac{B_i}{i} Y^i) = D_{-q^{-1}, Y}^r$.

By giving a shape datum $b_{\mu, \nu}: \kappa \leftrightarrow (a, \lambda)$ for $(\mathbf{Y}, \leftarrow_r, \mathbf{N}^r)$ satisfying (20), we shall prove (23) bijectively. Before we describe the shape datum (by an algorithm already hinted at), let us first mention an obvious idea that does not work. Like for the RSK correspondence, one might compute (a, λ) by taking standardisations of κ/μ and κ/ν , and completing the Schensted-growth for b^{SW} or b^{WS} . The fundamental problem with this approach is that, like for the r -correspondences derived from r -quotient maps, the property of being a standardised horizontal strip is not preserved across Schensted-growths, which spoils the construction (when a ribbon is bumped using b^{SW} or b^{WS} , its head may move to a quite distant diagonal, destroying the initial ordering among the ribbons). In this case “shuffling the ribbons” to get back a standardisation is not an option (even if μ/λ and κ/ν should happen to be horizontal r -ribbon strips), since that changes the spin and invalidates (20). One might hope to find an alternative r -correspondence that preserves standardisations, by ensuring that any ribbon of the standardisation that risks being overtaken by a bumped ribbon would itself get bumped instead; this is what always happens in the case $r = 1$. But for $r \geq 3$, no Schensted-growth can achieve this, as we shall now prove.

4.5.3. Proposition. *For any r -correspondence b for (\mathbf{Y}, \leq_r) with $r \geq 3$, there exists a Schensted-growth for b in which the standard r -ribbon tableau along the top edge of the grid is the standardisation of a horizontal r -ribbon strip, but the one along the bottom edge of the grid is not.*

Proof. We show this by exhibiting a pair of partial Schensted-growths, the completion of at least one of which will provide a witness for the proposition. We give the partial growths explicitly for $r = 3$, and then indicate how the example can be extended for $r > 3$. Here are the partial Schensted-growths:



All matrix entries are 0, and the leftmost grid squares have been completed independently of b . If the shape λ that is to complete the middle grid square adds a 3-ribbon below the 3×3 shape $\mu = (3, 3, 3)$, that is if $\lambda \in \{(3, 3, 3, 1, 1, 1), (3, 3, 3, 2, 1), (3, 3, 3, 3)\}$, then already the standard r -ribbon tableau along the bottom edge of the grid cannot be the standardisation of a horizontal r -ribbon strip. If λ adds a 3-ribbon with its head on a higher diagonal than the one added along the top edge of the grid, that is if $\lambda \in \{(5, 4, 3), (6, 3, 3)\}$, then the rightmost grid square can be completed independently of b and one has no standard r -ribbon tableau along the bottom either. But the only remaining possibility $\lambda = (4, 4, 4)$ cannot be realised in both Schensted-growths,

by the bijectivity of b . To construct similar examples for $r > 3$, take $\mu = (3^{(r)})$ (a $r \times 3$ rectangular shape), the last shape on the first row ($4^{(r)}$), the first shape on the second row $(3, 3, 2^{(r-3)})$ and the second shapes of the first row respectively $(2^{(r)})$ and $(3, 2^{(r-2)}, 1)$. \square

In the domino case ($r = 2$), standardisations of horizontal strips are not preserved by the Schensted-growths for $b = b^{\text{SW}}$ or $b = b^{\text{WS}}$, but they *are* preserved by the Barbasch-Vogan-Garfinkle correspondence (see [vLee1, definition 4.2.1]). Its 2-correspondence b^{BVG} uses the form of the 2-ribbon involved: one has $b_{\mu}^{\text{BVG}}(\kappa) = b_{\mu}^{\text{SW}}(\kappa)$ if $\text{form}(\kappa/\mu) = 0$ (a horizontal domino), and $b_{\mu}^{\text{BVG}}(\kappa) = b_{\mu}^{\text{WS}}(\kappa)$ if $\text{form}(\kappa/\mu) = 1$ (a vertical domino). Thus b^{BVG} will find the very first domino in the chosen direction that can be added or removed (the sum of 2 bits it monitors for changes starts at an extremal value, 0 or 2), which leaves no opportunity for overtaking other dominoes. The domino found also has the same height as the one bumped, which ensures bijectivity, and preservation of spin. A spin preserving Knuth correspondence for $(\mathbf{Y}, \leftarrow_2, \mathbf{N}^2)$ defined via standardisation from b^{BVG} is described (in different terms) in [ShW1].

In [ShW2] a “half semistandard” extension of the Schensted correspondence for b^{SW} is given, in which a semistandard r -ribbon tableau is allowed for the P -symbol, but the Q -symbol is still standard, and the matrix entries are still in $0 \cup e_{[r]}$. It is not defined directly by a Schensted-growth, but whenever a ribbon is “bumped” off its diagonal, a replacement ribbon is inserted *elsewhere* among the ribbons, constructed in a bijective and spin-preserving way. This ingenious construction has motivated our work, but it proved ill suited for adaptation to the “full semistandard” case: when multiple ribbons need to bump at the same time, the constructions of replacement ribbons interfere badly with each other; even if a suitable horizontal strip could be constructed, a step-by-step inverse seems impossible.

§5. A spin preserving shape datum.

In this section we shall describe a shape datum for $(\mathbf{Y}, \leftarrow_r, \mathbf{N}^r)$ that satisfies (20), and which therefore bijectively proves (23). An essential point is that the Schensted correspondence that it generalises is not the Shimozono-White correspondence but its transpose, and similarly that for $r = 1$ it reduces to the Burge correspondence rather than to the RSK correspondence. We shall first give an informal description, in a form similar to the algorithm given at the end of §4.3, without yet proving that this method actually works as it should. After that we shall give a more formal description, in a form that is more suited to the necessary proofs, which will then in fact be provided. Notably we shall prove that the correspondence has an inverse (given by a very similar procedure) and that equation (20) is indeed always satisfied.

5.1. Informal description, examples.

Let shapes μ, ν be given, in the same connected component of (\mathbf{Y}, \leq_r) . Our task is to establish a bijection between on one hand the shapes κ with $\mu \leftarrow_r \kappa \rightarrow_r \nu$ and on the other hand the pairs of a shape λ with $\mu \rightarrow_r \lambda \leftarrow_r \nu$ and an r -tuple $a \in \mathbf{N}^r$, such that equations (6) and (20) are satisfied in all cases. Although edge sequences are central to

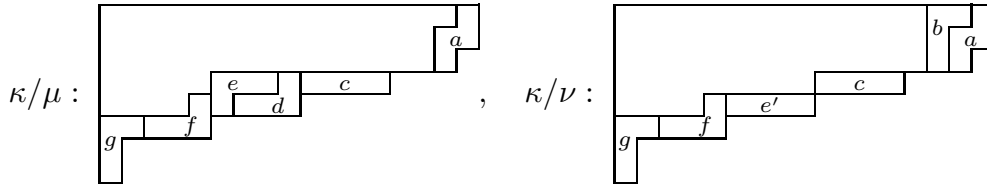
5.1 Informal description, examples

our considerations, we shall avoid referring to them in our informal description; however, we shall often compare r -ribbons by the diagonal containing the head of the ribbon, so to alleviate the terminology a bit, we shall say that a diagonal contains a ribbon, and that the ribbon is on the diagonal, when in fact the diagonal contains the head of the ribbon. We start with the extraction direction, i.e., the determination of (a, λ) given κ ; we shall assume that the standardisations of the horizontal r -ribbon strips κ/μ and κ/ν have been computed. Our description will be an algorithm that treats λ and a as variables, which are initialised as $\lambda := \kappa$ and $a := 0 \in \mathbf{N}^r$, and which will contain the desired values at termination.

One starts at the rightmost diagonal d_k containing a ribbon of at least one of the standardisations of the strips κ/μ and κ/ν ; should both strips be empty, then one terminates immediately. Each time a diagonal d_k has been processed as indicated below, one checks if the diagram of λ has an empty intersection with any of the diagonals d_j for $k - r \leq j < k$. If it has, one terminates (there are no r -ribbons left to remove from λ); if not, one continues to consider the diagonal d_{k-1} . No action is required for d_k unless there exists an r -ribbon on d_k that can be removed from λ . If so, let λ/λ' be that r -ribbon, and put $h = \text{ht}(\lambda/\lambda')$. If both the standardisations of the strips κ/μ and κ/ν have r -ribbons on d_k , modify a by setting $a_h := a_h + 1$. If at least one of those standardisations has an r -ribbon on d_k , put $\lambda := \lambda'$, and processing of d_k is completed. In the remaining case (neither of the standardisations has a ribbon on d_k), processing is also completed (without action) if $a_h = 0$; otherwise one sets $a_h := a_h - 1$ and $\lambda := \lambda'$, completing the processing for d_k . After the algorithm terminates, κ/λ will be a horizontal r -ribbon strip whose standardisation is given by the intermediate values of λ , in other words, it has ribbons on those diagonals for which the variable λ was modified. Then μ/λ will also be a horizontal r -ribbon strip, and its standardisation has ribbons on those diagonals for which λ was modified and that do not contain a ribbon of the standardisation of κ/μ ; the situation is similar for ν/λ and κ/ν .

The reverse procedure is quite similar, but this time the variables λ and a are initialised from the given shape and matrix entry, and one starts with $k = -\lambda_0^t$, so that the diagonal d_k contains the leftmost r -ribbon that can be added to λ , whose head is the square $(0, \lambda_0^t)$, the first one below the leftmost column of λ . One proceeds from there for successive values of k , until $a = 0$ holds and there are no more ribbons of the standardisations of the strips μ/λ and ν/λ any diagonal d_i with $i \geq k$. Action is needed for d_k only if an r -ribbon exists on d_k that can be added to λ . If so, let λ'/λ be that r -ribbon, and put $h = \text{ht}(\lambda'/\lambda)$. If both the standardisations of the strips μ/λ and ν/λ have r -ribbons on d_k , set $a_h := a_h + 1$, and then if at least one of them has an r -ribbon on d_k set $\lambda := \lambda'$, completing the processing of d_k . In the remaining case processing is also completed (without action) if $a_h = 0$; otherwise the processing of d_k consists of setting $a_h := a_h - 1$ and $\lambda := \lambda'$. This second procedure exactly retraces the steps of the first one, but this is not as obvious as it might seem at first glance. Notably, encountering a diagonal where an r -ribbon of height h can be added to λ triggers no action during the first procedure, but during the reverse procedure this circumstance *will* cause action, unless $a_h = 0$. We shall see below that in these cases $a_h = 0$ always holds, which resolves the mystery.

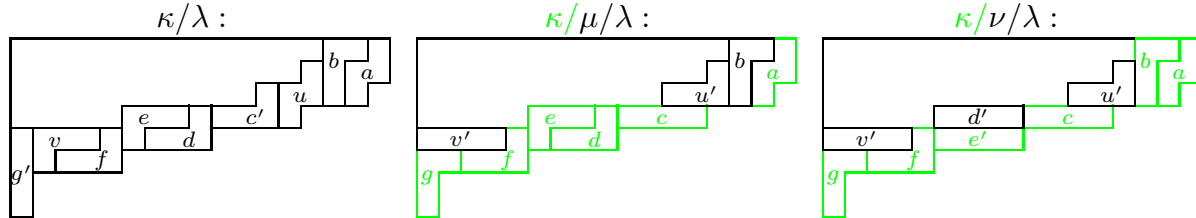
Now let us illustrate these computations in an example. In order to exercise the different cases that can arise, we need rather large diagrams, in particular horizontally. Our example is for $r = 4$; we take $\mu = (16, 15, 15, 5, 4)$, $\nu = (14, 14, 14, 9, 4)$, and $\kappa = (17, 17, 16, 13, 9, 5, 1, 1)$. Then the standardisations of the horizontal 4-ribbon strips involved are as displayed below. We have labelled individual ribbons to facilitate our discussion of the procedure, giving the same label to ribbons whose diagram is identical (even if they are not identical as ribbons, i.e., the step in (\mathbf{Y}, \prec_r) in the two standardisations may differ; this should not cause any confusion).



The procedure starts with $\lambda := \kappa$ and $a := (0, 0, 0, 0)$, on the diagonal containing the ribbon a . Since this ribbon has height 2 and occurs both in κ/μ and in κ/ν , we set $a_2 := 1$, and then remove the ribbon from λ . The next diagonal (to its left) contains b , which only occurs in κ/ν ; therefore $a_2 = 1$ remains unchanged, and b is removed from λ , which becomes $(14, 14, 14, 13, 9, 5, 1, 1)$. No 4-ribbon can be removed from λ on the next two diagonals, but this is possible on the diagonal after that. This ribbon, which we shall call u , has $\text{form}(u) = (1, 0, 1)$ and $\text{ht}(u) = 2$; its diagonal contains no ribbons either in κ/μ or in κ/ν and so since $a_2 = 1$, we decrement a_2 back to 0 and remove u from λ , which now becomes $(14, 13, 12, 12, 9, 5, 1, 1)$. On the next diagonal another 4-ribbon (of form $(1, 1, 0)$) can be removed, but since there are still no ribbons either in κ/μ or in κ/ν on this diagonal, and by now $a_2 = 0$, no action is taken here. Skipping another diagonal, on which no 4-ribbon can be removed from λ , we come to the diagonal of c . There a 4-ribbon can be removed from λ , but it is not c : it has form $(0, 0, 1)$ and height 1 and we shall call it c' . Since c occurs both in κ/μ and in κ/ν , we set $a_1 := 1$, and removing c' from λ it becomes $(14, 13, 11, 9, 9, 5, 1, 1)$. Three diagonals follow where no 4-ribbons can be removed from λ (nor could any be added) and we arrive at the diagonal of d . It occurs only in κ/μ so we remove it from λ and arrive at the diagonal of e and e' . The 4-ribbon that can be removed from λ on this diagonal has the form of e and therefore height 1; because of the occurrences of e and e' we raise $a_1 := 2$ and remove e from λ which now becomes $(14, 13, 11, 5, 5, 5, 1, 1)$. There follow a diagonal on which a 4-ribbon could be added to λ and one on which a 4-ribbon could be removed, but since the latter has height 2 while $a_2 = 0$ we do nothing here. Then comes f , which can be removed from λ ; since it has height 1 and occurs both in κ/μ and in κ/ν we further raise $a_1 := 3$ and remove f from λ . The next diagonal allows removal of a 4-ribbon v of height 1 from λ but has no ribbons in κ/μ or in κ/ν ; therefore we lower $a_1 := 2$ and remove the ribbon from λ leaving $(14, 13, 11, 5, 1, 1, 1, 1)$. The next two diagonals would allow adding 4-ribbons to λ so they are skipped; we come at the diagonal of g , on which the 4-ribbon g' that can be removed from λ is vertical (height 3), so we raise $a_3 := 1$ and remove g' from λ . No squares of λ remain on the current diagonal (let alone to its left) so we terminate with $\lambda = (14, 13, 11, 5)$ and $a = (0, 2, 0, 1)$.

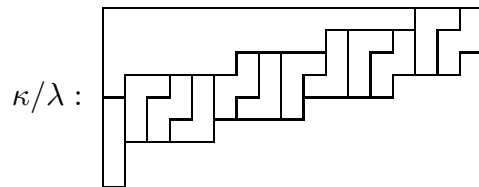
5.1 Informal description, examples

We summarise the result by displaying the standardisations of κ/λ , μ/λ , and ν/λ . For the two latter strips, each of which has only three ribbons, we have also included (in green when colours are available) standardisations of their complementary strips in κ/λ , to emphasise the relative locations of the ribbons.



We observe that $|\kappa/\mu|_4 - |\nu/\lambda|_4 = 6 - 3 = 3 = |a|$ shows that (6) indeed holds in this case, and similarly $\text{spin}(\kappa/\mu) + \text{spin}(\kappa/\nu) - \text{spin}(\mu/\lambda) - \text{spin}(\nu/\lambda) = \frac{7}{2} + \frac{7}{2} - \frac{3}{2} - \frac{1}{2} = 5 = n(a)$ shows that (20) holds. From the above description, the one for the opposite process (computing κ from the (final) values of λ and a) can be obtained by a reversal of its steps, replacing each action by its inverse; we leave it to the reader to work this out. We do note however that those diagonals where ribbons could be added to λ (a fact not really relevant to the initial computation) now get considered for action, but invariably get skipped nonetheless because the relevant component a_h is zero. This is for instance the case just after adding g' to λ and decrementing a_3 to 0 on the first diagonal considered: two occasions follow to add ribbons of height 3 and 2, respectively, but neither is used because now $a_3 = 0$, and $a_2 = 0$. Similarly, after adding ribbons v (on the diagonal of v') and f (because this time $a_1 = 3$ can be decremented) and skipping a diagonal where a ribbon could be removed, an occasion to add a ribbons of height 2 is not used because one (still) has $a_2 = 0$; the next occasion involves ribbon e of height 1, and it is used because $a_1 = 2$ there.

Let us use this example to make an informal observation: the algorithm turns out to be surprisingly sensitive the input. Suppose we change the data of the example slightly by including ribbon b into κ/μ . Then when this ribbon is reached, instead of leaving $a_2 = 1$, the initial procedure will increment a_2 . One might imagine that this just causes a_2 to be incremented in the final result, but in fact the whole process completely changes. One may check that the results will be $\lambda = (14, 10, 6, 1)$ and $a = (0, 0, 0, 1)$, so rather than an increase of a_2 we see a decrease of a_1 . Indeed the latter component never even gets to be incremented, which is related to the fact that κ/λ now has no ribbons of height less than 2 in its standardisation:



Here is how the statistics change: $|\kappa/\mu|_4 - |\nu/\lambda|_4 = 7 - 6 = 1 = |a| = 6 - 5 = |\kappa/\nu|_4 - |\mu/\lambda|_4$ and $\text{spin}(\kappa/\mu) + \text{spin}(\kappa/\nu) - \text{spin}(\mu/\lambda) - \text{spin}(\nu/\lambda) = \frac{9}{2} + \frac{7}{2} - \frac{4}{2} - \frac{6}{2} = 3 = n(a)$.

5.2. Formal description.

Our formal description will be directly in terms of edge sequences, since all the conditions and actions used are more easily expressed in terms of these. At the same time this allows us to generalise by considering arbitrary doubly infinite sequences of bits; the dependence on the limiting behaviour of these sequences will be made explicit. We therefore define a *bit sequence* to be any function $s: \mathbf{Z} \rightarrow \{0, 1\}$; its value at i will be denoted by s_i . Moreover, our description will be static, in the sense that there are no variables that are modified in time. Instead all values used are represented in a single structure, and the algorithms translate into relations between components of this structure that allow the whole to be recovered from partial information, much like growth diagrams can replace the algorithmic descriptions of Schensted and Knuth correspondences. An additional advantage of this method of description is that it simultaneously describes the forward and the reverse algorithms.

A first notion to formalise is that of horizontal r -ribbon strips. We have characterised the relation $\lambda \leftarrow_r \mu$ by the possibility to change λ into μ by adding certain r -ribbons in a left-to-right fashion. A static formulation of the relation $s \leftarrow_r t$ for bit sequences s, t could be given by requiring the existence of a set of intermediate bit sequences, and specifying the relations that force them to correspond to the mentioned type of transformation. There is however a more economical way to proceed, based on the observation that adding r -ribbons induces only a minimal change on the bit sequence: it will suffice to require the existence of a single “intermediate” bit sequence w , with a simple condition for the way it differs from s and t . At the same time w will allow reading off the height of the horizontal r -ribbon strip t/s . During the addition of the ribbons of a horizontal r -ribbon strip, a single bit s_i can get changed at most twice: once for a ribbon with its head on diagonal i , and once for a ribbon with its tail on diagonal $i + 1$, and therefore with its head on diagonal $i + r$; moreover, should both changes occur, then the former will precede the latter due to the left-to-right requirement. The bit w_i will describe the state of s_i after a possible modification of the first type, while t_i gives its final state after a possible modification of the second type. The principal condition on s , w , and t relates the two changes made by adding one ribbon: t_{i-r} differs from w_{i-r} if and only if w_i differs from s_i , in which case one must have $(w_{i-r}, t_{i-r}, s_i, w_i) = (1, 0, 0, 1)$. The bit sequence to which this ribbon is added consists of the bits t_j for $j < i - r$ followed by the bits w_j for $i - r \leq j < i$ followed by s_j for $j \geq i$; in particular, the height of the ribbon is given by $\sum_{j=1}^{r-1} w_{i-j}$. One additional condition is needed to characterise $s \leftarrow_r t$, namely that $(w_{i-r}, t_{i-r}, s_i, w_i) = (1, 0, 0, 1)$ occurs only for finitely many values of i (since horizontal r -ribbon strips must be finite); this means that $s_i = w_i = t_i$ for $i \ll 0$ and for $i \gg 0$. We state formally:

5.2.1. Definition. *Let s, t be bit sequences, then t/s is a horizontal r -ribbon strip (written $s \leftarrow_r t$) if there exists a bit sequence w such that $(w_{i-r}, t_{i-r}, s_i, w_i) \in \{(a, a, b, b) \mid a, b \in \{0, 1\}\} \cup \{(1, 0, 0, 1)\}$ for all $i \in \mathbf{Z}$, while $(w_{i-r}, t_{i-r}, s_i, w_i) = (1, 0, 0, 1)$ occurs only for finitely many i . We say that w is the witness for $s \leftarrow_r t$, that t/s has an r -ribbon at position i if $i \in I = \{i \in \mathbf{Z} \mid (w_{i-r}, t_{i-r}, s_i, w_i) = (1, 0, 0, 1)\}$, whose height is $\sum_{j=1}^{r-1} w_{i-j}$. We put $|t/s|_r = \#I$ and $\text{spin}(t/s) = \frac{1}{2} \sum_{i \in I} \sum_{j=1}^{r-1} w_{i-j}$.*

5.2 Formal description

The second part of this definition depends on the uniqueness of a witness for $s \prec_r t$, or equivalently the uniqueness of I ; by the required finiteness of I one has $\min(I) = \min(\{i \in \mathbf{Z} \mid s_i \neq t_i\}) + r$ and $\max(I) = \max(\{i \in \mathbf{Z} \mid s_i \neq t_i\})$ if $I \neq \emptyset$, and from this all of I (and of w) can be found by a simple induction. If $k = \min(I)$, one can split off the first r -ribbon: $s \prec_r s' \prec_r t$, with $(s'_{k-r}, s'_k) = (0, 1)$ and $s'_j = s_j$ for $j \notin \{k-r, k\}$, and the witness w' for $s' \prec_r t$ differing from w only in the bit $w'_{k-r} = 0$ (one has $w'_k = 1 = w_k$); any r -ribbon of s'/s or t/s' has the same height as the r -ribbon of t/s at the same position. We shall also need the following result, which deals with a situation in which ribbons are merged at interleaving positions.

5.2.2. Lemma. *Let $s, t, u: \mathbf{Z} \rightarrow \{0, 1\}$ be bit sequences such that $s \prec_r t \prec_r u$ and $s \prec_r u$ hold. Then the set of positions at which u/s has an r -ribbon is the disjoint union of the sets I, J of positions at which t/s and u/t , respectively, have r -ribbons. Moreover the sums of the heights of the ribbons of u/s with positions in each of the sets I and J exceed the sums of the heights of the ribbons of t/s and u/t , respectively, in both cases by the same number $c \in \mathbf{N}$.*

Proof. From a relation $p \prec_r q$ it follows that for the first difference $p_k \neq q_k$ (i.e., with minimal k), if any, one has $p_k = 1$ and $q_k = 0$. In the situation of the lemma this implies that $s = u$ can only happen if also $s = t$ (in which case the statement is trivial), and that otherwise the minimal position k at which u/s has an r -ribbon is either $\min(I)$ or $\min(J)$, but not both. We shall proceed by induction on $|u/s|_r$. The case $k = \min(I)$ is easy: splitting off an horizontal r -ribbon strip s'/s with a single r -ribbon at position k from either u/s or t/s , it suffices to apply the induction hypothesis to s', t, u .

Now assume $k = \min(J)$. As before we add an r -ribbon at position k to s , but since there is no such ribbon in t/s we must add a ribbon to t as well. So define bit sequences s', t' , and w' , that differ from s, t , and the witness w for $s \prec_r t$, respectively, by $(s_{k-r}, s_k) = (t'_{k-r}, t'_k) = (w'_{k-r}, w'_k) = (0, 1)$; then w' witnesses $s' \prec_r t'$, and t'/s' has r -ribbons at the same positions $i \in I$ as t/s . One also has $t' \prec_r u$ and $s' \prec_r u$, and each r -ribbon of u/t' or u/s' has the same height as the ribbon of u/t or u/s , respectively, at the same position. By induction the lemma holds for s', t', u with some value c_1 for c ; we shall complete the induction step by showing that $\text{spin}(s'/s) = \text{spin}(t'/t) + \frac{c_0}{2}$ and $\text{spin}(t'/s') = \text{spin}(t/s) + \frac{c_0}{2}$ for some $c_0 \in \mathbf{N}$, so that the lemma holds for s, t, u with $c = c_0 + c_1$.

Since $i > k$ for all $i \in I$, comparison of $\text{spin}(s'/s) = \frac{1}{2} \sum_{j=1}^{r-1} s_{k-r+j}$ with $\text{spin}(t'/t) = \frac{1}{2} \sum_{j=1}^{r-1} t_{k-r+j}$ shows that for each j either $s_{k-r+j} = t_{k-r+j}$ or $(s_{k-r+j}, t_{k-r+j}) = (1, 0)$, where the latter possibility occurs if $k+j \in I$; hence $c_0 = \#I_0$ where $I_0 = \{i \in I \mid k < i < k+r\}$. The same set I_0 occurs as the set of positions at which the r -ribbons of t'/s' differ in height from the r -ribbons of t/s at the same position, as $\sum_{j=1}^{r-1} w'_{i-j}$ includes the term $w'_k = 1 \neq w_k$ iff $i \in I_0$, proving $\text{spin}(t/s) = \text{spin}(t'/s') - \frac{c_0}{2}$. \square

5.2.3. Definition. *Let $l, m, n, k: \mathbf{Z} \rightarrow \{0, 1\}$ be bit sequences, and $a_{\ll}, a_{\gg} \in \mathbf{N}^r$. Then we call the data $((\binom{l}{n \ k}, a_{\ll}, a_{\gg}))$ a basic square configuration if the skew shapes m/l , n/l , k/m , k/n and k/l are all horizontal r -ribbon strips, and with w the witness for*

$l \leftarrow_r k$ there exists a map $a: \mathbf{Z} \rightarrow \mathbf{N}^r$ such that for all $i \in \mathbf{Z}$ the following statements hold, where $h = \sum_{j=1}^{r-1} w_{i-j}$:

- (1) If k/l has no r -ribbon at position i , then $a(i) = a(i+1)$.
- (2) If $w_{i-r} \neq w_i$, then $a(i)_h = 0$.
- (3) If k/l has an r -ribbon at position i (whose height is then h), put $d = 1$ if both m/l and n/l have r -ribbons at position i , put $d = -1$ if both k/m and k/n have r -ribbons at position i , and put $d = 0$ otherwise. Then $a(i+1) = a(i) + de_h$, i.e., $a(i+1)_h = a(i)_h + d$, and $a(i+1)_j = a(i)_j$ for all $j \neq h$.
- (4) If k/l has no r -ribbons at positions $j < i$, then $a(i) = a_{\ll}$.
- (5) If k/l has no r -ribbons at positions $j \geq i$, then $a(i) = a_{\gg}$.

We shall show that these conditions encode the rules of the mutually inverse algorithms described informally above, when l, m, n, k are the respective edge sequences of $\lambda, \mu, \nu, \kappa$, when a_{\ll} is the matrix entry a , and $a_{\gg} = 0$. The crucial condition (2) will also help us to prove the fact that the algorithms are indeed each others inverses (a basic square configuration can be completed in two directions from two different sets of partial information). It follows from the first condition that $a(i)$ becomes stationary for i at either side of a finite interval determined by k/l ; the two final conditions merely state that a_{\ll} and a_{\gg} give the limiting values. When dealing with edge sequences of partitions, one will always have $a_{\gg} = 0 \in \mathbf{N}^r$. It is nevertheless useful to allow both limiting values to be arbitrary: firstly because this gives greater generality while exhibiting the link with the asymptotic behaviour of the bit sequences, and secondly because it allows us to recognise a formal symmetry between the two directions of computation (in addition to the more obvious symmetry with respect to m and n).

5.2.4. Proposition. Let s^\diamond denote the reverse of a bit sequence s , defined by $s_i^\diamond = s_{-1-i}$. Then $((\binom{l}{n} \binom{m}{k}), a_{\ll}, a_{\gg})$ is a basic square configuration if and only if $((\binom{k^\diamond}{m^\diamond} \binom{n^\diamond}{l^\diamond}), a_{\gg}, a_{\ll})$ is one.

Proof. It is easy to see that $s \leftarrow_r u$ if and only if $u^\diamond \leftarrow_r s^\diamond$, since if w is the witness for the former relation, then w^\diamond is the witness for the latter. Let $a: \mathbf{Z} \rightarrow \mathbf{N}^r$ be the map used to establish that $((\binom{l}{n} \binom{m}{k}), a_{\ll}, a_{\gg})$ is a basic square configuration, then a direct verification then shows that $a': i \mapsto a(r-i)$ establishes the second basic square configuration. In this verification one uses the fact that when 5.2.3(2) applies, then so does 5.2.3(1), so $a(i) = a(i+1)$. The reverse implication follows by symmetry. \square

In the following (main) theorem, information about the asymptotic behaviour of the bit sequences is extracted by the “lim inf” operation. It is applied to expressions taking values in a finite set, namely the sum of r consecutive bits, so it returns the smallest value that is assumed infinitely often by this sum. As the bit sequences differ only at finitely many places, it does not matter of which of them the bits are considered. For the case of edge sequences one will have $h_{\ll} = r$ below, which leaves a_{\ll} unrestricted, as well as $h_{\gg} = 0$, which forces $a_{\gg} = 0 \in \mathbf{N}^r$. The previous proposition allows us to state the theorem in just one direction; the “insertion” direction is slightly easier to formulate.

5.2 Formal description

5.2.5. Theorem. Let $r \in \mathbf{N}_{>0}$ and let $l, m, n: \mathbf{Z} \rightarrow \{0, 1\}$ be bit sequences with $m \rightarrow_r l \leftarrow_r n$; let $a_{\ll} \in \mathbf{N}^r$ be such that $(a_{\ll})_h = 0$ for $h_{\ll} \leq h < r$, where $h_{\ll} = \liminf_{i \rightarrow -\infty} \sum_{j \in [r]} l_{i-j}$. Then there exist unique $k: \mathbf{Z} \rightarrow \{0, 1\}$ and $a_{\gg} \in \mathbf{N}^r$ such that $((\binom{l}{n} \binom{m}{k}), a_{\ll}, a_{\gg})$ is a basic square configuration; moreover, its values a, w satisfy for all $i \in \mathbf{Z}$ the condition $C_i: a(i)_h = 0$ for $h_i \leq h < r$ where $h_i = \sum_{j \in [r]} w_{i-1-j}$. Finally, one has $(a_{\gg})_h = 0$ for $h_{\gg} \leq h < r$ where $h_{\gg} = \liminf_{i \rightarrow +\infty} \sum_{j \in [r]} l_{i-j}$.

Proof. If $m = l = n$ while $(a_{\ll})_h = 0$ for $h_{\min} \leq h < r$ where $h_{\min} = \min_{i \in \mathbf{Z}} (\sum_{j \in [r]} l_{i-j})$, then one can take $k = l$ and for $a: \mathbf{Z} \rightarrow \mathbf{N}^r$ the constant function with value a_{\ll} , which then is also the value of a_{\gg} : in definition 5.2.3, conditions (1), (4), and (5) hold for all i , and one has $w = l$ so that in the cases where (2) applies, its value of h satisfies $h_{\min} \leq h = \min(h_i, h_{i+1}) < r$, whence $a(i)_h = (a_{\ll})_h = 0$ by assumption. That solution for this case is also unique, because if k/l should have any r -ribbon, then applying condition (3) for its ribbon at minimal position i , the conclusion $a(i+1) = a(i) - e_h$ is incompatible with the fact that for $h = \sum_{j=1}^{r-1} w_{i-j} = \sum_{j=0}^{r-1} l_{i-j} \geq h_{\min}$ one has $a(i)_h = (a_{\ll})_h = 0$.

In the remaining cases we can set i_0 to the smallest value such that either at least one of $\{m/l, n/l\}$ has an r -ribbon at position i_0 , or $h = \sum_{j \in [r]} l_{i_0-j}$ satisfies $h < r$ and $(a_{\ll})_h > 0$: due to the assumption about a_{\ll} , the set of such values is bounded below, and if it were empty, one would be in the case just treated. By an argument similar to the one above, k/l cannot have any r -ribbons at positions $i < i_0$, so that any solution should have $w_i = l_i$, $k_{i-r} = l_{i-r}$, and $a(i+1) = a_{\ll}$ for $i < i_0$; in particular C_i will be satisfied for $i \leq i_0$. We now show for any $i \in \mathbf{Z}$ that, when the values $a(i)$ and w_j, k_{j-r} for $j < i$ are known and C_i holds, one can uniquely determine w_i, k_{i-r} and $a(i+1)$ and deduce C_{i+1} ; this suffices to determine a unique solution. Suppose first that $(w_{i-r}, l_i) \neq (1, 0)$, so that k/l cannot have an r -ribbon at position i , which forces $(k_{i-r}, w_i) = (w_{i-r}, l_i)$, and $a(i+1) = a(i)$ by condition 5.2.3(1). In this case condition C_i implies C_{i+1} , since $h_{i+1} = h_i - w_{i-r} + w_i \geq h_i$. If in fact $(w_{i-r}, l_i) = (0, 1)$ one even has $a(i+1)_{h_i} = a(i+1)_{h_{i+1}-1} = 0$; this is not required for C_{i+1} , but it ensures condition 5.2.3(2).

Next assume that $(w_{i-r}, l_i) = (1, 0)$, and put $h = \sum_{j=1}^{r-1} w_{i-j}$. Then it will be possible to put $(k_{i-r}, w_i) = (0, 1)$, so that k/l has an r -ribbon at position i , unless 5.2.3(3) would then cause $a(i+1)_h$ to become negative, i.e., if neither m/l nor n/l has an r -ribbon at position i , and $a(i)_h = 0$. Another potential possibility is to put $(k_{i-r}, w_i) = (w_{i-r}, l_i) = (1, 0)$, so that k/l has no r -ribbon at position i . For this it is necessary that neither m/l nor n/l have such a ribbon, and 5.2.3(2) requires moreover that in this case $a(i)_h = 0$; one arrives at the same condition that was excluded for the first option. As a consequence there remains a unique possibility for the value of (k_{i-r}, w_i) in all cases. As for condition C_{i+1} , in the cases where k/l has an r -ribbon at position i it is equivalent to C_i , since $w_{i-r} = 1 = w_i$ implies $h_i = h_{i+1}$ while $h = h_i - 1 < h_{i+1}$ is the only index at which $a(i+1)$ may differ from $a(i)$. In the case where k/l has no r -ribbon at position i , one has $w_{i-r} = 1$ and $w_i = 0$, so $h_{i+1} = h_i - 1 = h$, requiring an additional component of $a(i+1) = a(i)$ to vanish, namely $a(i+1)_h$. But this is equal to the value $a(i)_h$ that was required to be zero in order to have this case

in the first place, so C_{i+1} still holds.

It remains to verify some statements after all of a and l have been determined. First, let $S \subseteq \mathbf{Z}$ be the set of positions i where it was decided that k/l has an r -ribbon, then for every $i \in S$ there either exists an r -ribbon of m/l or of n/l (which happens only finitely often), or one has $|a(i+1)| < |a(i)|$; this forces S to be finite, whence indeed $l \leftarrow_r k$. The positions i for which $a(i+1) \neq a(i)$ are a subset of S , whence $a(i)$ eventually becomes stationary at some value a_{\gg} as $i \rightarrow +\infty$; moreover one sees that 5.2.3(4) and (5) hold. When $i \notin S$ one has $w_i = l_i$, so under ‘lim inf’ in the expression defining h_{\gg} one may replace l by w , and the final statement of the theorem is a consequence of the validity of C_i for all $i \in \mathbf{Z}$. \square

Let us recapitulate somewhat less formally the roles played by condition 5.2.3(2) and condition C_i of the theorem. The primary function of 5.2.3(2) is to remove the ambiguity between (a) adding a ribbon at position i to k/l , leaving h_i unchanged as i is increased but decreasing $a(i)_{h_i-1}$, and (b) not adding such a ribbon, decreasing h_i while leaving $a(i)$ unchanged. By forbidding the decrease of h_i to some level h until $a(i)_h = 0$, option (b) is restricted to the case when option (a) is not available. By the symmetry of proposition 5.2.4, a similar ambiguity is removed in the opposite direction: in the case where k/l has a ribbon at position i but neither k/m nor k/n do, the data $(a(i+1), h_{i+1})$ can be distinguished from any that could arise for a case where h_{i+1} was increased from its previous level $h_i = h$, since the latter requires $a(i)_h = 0$. But that requirement could also block progress, because the case where h_i is increased is one where there is no choice: should it occur while $a(i)_{h_i} \neq 0$, then there is no way to satisfy definition 5.2.3. This is where condition C_i comes in, as it ensures that whenever the current level $h = h_i$ is not r one has $a(i)_h = 0$. It must require so as well for all higher levels to anticipate raising of the level, while as we saw condition (2) takes care of preserving C_i when lowering of the level.

By proposition 5.2.4 the final data (a_{\gg}, k) of theorem 5.2.5 can also be used to uniquely recover (a_{\ll}, l) , so that the construction establishes the following bijective correspondence.

5.2.6. Corollary. *Let $r \in \mathbf{N}_{>0}$, let m, n be a pair of bit sequences that differ only at finitely many positions, and put $h_{\ll} = \liminf_{i \rightarrow -\infty} \sum_{j \in [r]} m_{i-j}$ and $h_{\gg} = \liminf_{i \rightarrow -\infty} \sum_{j \in [r]} m_{i-j}$. Then there is a bijection between on one side the pairs (a_{\ll}, l) of a bit sequence l with $m \rightarrow_r l \leftarrow_r n$ and $a_{\ll} \in \mathbf{N}^r$ with $(a_{\ll})_h = 0$ for all $h_{\ll} \leq h < r$, and on the other side the pairs (a_{\gg}, k) of a bit sequence k with $m \leftarrow_r k \rightarrow_r n$ and $a_{\gg} \in \mathbf{N}^r$ with $(a_{\gg})_h = 0$ for all $h_{\gg} \leq h < r$; this bijection is determined by the fact that $((\binom{l \ m}{n \ k}), a_{\ll}, a_{\gg})$ is a basic square configuration. \square*

Taking for m, n the edge sequences of partitions μ, ν one has $h_{\ll} = r$ and $h_{\gg} = 0$, the latter of which forces $a_{\gg} = 0 \in \mathbf{N}^r$. Then in the corollary the bit sequences l, k will also be edge sequences of partitions, say of λ, κ respectively, and one has $\mu \rightarrow_r \lambda \leftarrow_r \nu$ and $\mu \leftarrow_r \kappa \rightarrow_r \nu$. Thus we arrive at our conclusion:

5.2.7. Theorem. *Let $r \in \mathbf{N}_{>0}$. There exists a shape datum $(b_{\mu, \nu})_{\mu, \nu \in \mathbf{Y}}$ for $(\mathbf{Y}, \leftarrow_r, \mathbf{N}^r)$ determined by the requirement that $(a, \lambda) = b_{\mu, \nu}(\kappa)$ holds if and only*

5.2 Formal description

if $((\begin{smallmatrix} \delta(\lambda) & \delta(\mu) \\ \delta(\nu) & \delta(\kappa) \end{smallmatrix}), a, 0)$ is a basic square configuration. Moreover, in that case one has $\text{spin}(\kappa/\mu) + \text{spin}(\kappa/\nu) = \text{spin}(\mu/\lambda) + \text{spin}(\nu/\lambda) + n(a)$.

Proof. The construction of the bijections $b_{\mu,\nu}$ is given by corollary 5.2.6; the only things left to check are the condition $|\kappa|_r - |\mu|_r - |\nu|_r + |\lambda|_r = |a|$ required by the notion of shape datum, and the one in the final statement of the theorem. The former follows by observing that for any basic square configuration $((\begin{smallmatrix} l & m \\ n & k \end{smallmatrix}), a_{\ll}, a_{\gg})$ one has $|a_{\ll}| = |a_{\gg}| + \sum_{i \in \mathbf{Z}} (|a(i)| - |a(i+1)|)$ (the sum is effectively finite), and that $|a(i)| - |a(i+1)|$ is equal (by 5.2.3(1,3) and the first part of lemma 5.2.2) to the contribution of ribbons at position i to the value $|\kappa/\mu|_r - |\nu/\lambda|_r = |\kappa/\nu|_r - |\mu/\lambda|_r$. For the final statement of the theorem one uses the second part of lemma 5.2.2. For a horizontal r -ribbon strip t/s let us denote by $I(t/s)$ the set of positions at which t/s has a ribbon, and for $i \in I(t/s)$ by $\text{spin}_i(t/s)$ the contribution of that ribbon to $\text{spin}(t/s)$ (i.e., $\frac{1}{2} \sum_{j=1}^{r-1} w_{i-j}$ where w is the witness for $s \leftarrow_r t$). Then the lemma says that

$$\sum_{i \in I(m/l)} (\text{spin}_i(k/l) - \text{spin}_i(m/l)) = \sum_{i \in I(k/m)} (\text{spin}_i(k/l) - \text{spin}_i(k/m))$$

and also

$$\sum_{i \in I(n/l)} (\text{spin}_i(k/l) - \text{spin}_i(n/l)) = \sum_{i \in I(k/n)} (\text{spin}_i(k/l) - \text{spin}_i(k/n))$$

for $(\begin{smallmatrix} l & m \\ n & k \end{smallmatrix}) = (\begin{smallmatrix} \delta(\lambda) & \delta(\mu) \\ \delta(\nu) & \delta(\kappa) \end{smallmatrix})$. Then the expression $\text{spin}(\kappa/\mu) + \text{spin}(\kappa/\nu) - \text{spin}(\mu/\lambda) - \text{spin}(\nu/\lambda)$ that we should prove equal to $n(a) = n(a_{\ll})$ can be rewritten as

$$\sum_{i \in I(k/m)} \text{spin}_i(k/l) + \sum_{i \in I(k/n)} \text{spin}_i(k/l) - \sum_{i \in I(m/l)} \text{spin}_i(k/l) - \sum_{i \in I(n/l)} \text{spin}_i(k/l).$$

It is readily checked that any ribbon considered in 5.2.3(3) contributes $-dh$ to this expression, which can be written as $-n(de_h)$; according to 5.2.3(3) this is equal to $n(a(i)) - n(a(i+1))$. Like above one has $n(a_{\ll}) = n(a_{\gg}) + \sum_{i \in \mathbf{Z}} (n(a(i)) - n(a(i+1)))$, while $n(a_{\gg}) = 0$; this gives the required identity. \square

Now by constructing Knuth-growths for this shape datum as described in §2.2, we obtain the spin preserving Knuth correspondence that we set out to find; it is symmetric since $b_{\mu,\nu} = b_{\nu,\mu}$ for all $\mu, \nu \in \mathbf{Y}$. We summarise its main characteristics.

5.2.8. Corollary. *Let $r \in \mathbf{N}_{>0}$, let γ be an r -core, and $\mathbf{RH}(\gamma)$ the associated r -rim hook lattice (the connected component of (\mathbf{Y}, \leq_r) containing γ). Then for any $m, n \in \mathbf{N}$ there exists a bijection between on one side $m \times n$ matrices A with entries in \mathbf{N}^r , and on the other side pairs (P, Q) of semistandard r -ribbon tableaux of equal shape λ/γ for some $\lambda \in \mathbf{RH}(\gamma)$, such that one has $\text{wt}(P) = (\sum_{i \in [m]} |A_{i,j}|)_{j \in [n]}$ and $\text{wt}(Q) = (\sum_{j \in [n]} |A_{i,j}|)_{i \in [m]}$, while also $\text{spin}(P) + \text{spin}(Q) = \sum_{i,j \in [m] \times [n]} n(A_{i,j})$. Moreover the correspondence is symmetric: if A is mapped to (P, Q) , then the transpose matrix A^t maps to (Q, P) . \square*

Theorem 5.2.7 bijectively proves equation (23). The full Knuth correspondence corresponds to the following obvious $q^{\frac{1}{2}}$ -analogue of (11). Define $G_{\lambda/\gamma}^{(r)}(q^{\frac{1}{2}}, X) = \sum_P q^{\text{spin}(P)} X^{\text{wt}(P)} \in \mathbf{Z}[q^{\frac{1}{2}}][[X]]$, where the sum is over all semistandard r -ribbon tableaux P of shape λ/γ , identifying tableaux that differ only by stationary steps at the end of the path (staying at the shape λ). Then one obtains from interpreting corollary 5.2.8, or directly by multiplying copies of (23):

$$\prod_{i,j \in \mathbf{N}} \prod_{k \in [r]} \frac{1}{1 - q^k X_i Y_j} = \sum_{\lambda \in \mathbf{RH}(\gamma)} G_{\lambda/\gamma}^{(r)}(q^{\frac{1}{2}}, X) G_{\lambda/\gamma}^{(r)}(q^{\frac{1}{2}}, Y). \tag{24}$$

This analogue of the Cauchy identity, and its variant (33) given below, have independently been derived in a different way (not using the commutation relations (23) and (32)) in an recent paper [Lam].

Returning to the questions considered in §1.1, we observe that corollary 5.2.6 applied to cases with $n = m$ for which $h_{\ll} = 0 = h_{\gg}$ provides a bijective proof of claim 1.1.2. For cases with $n = m$ for which $h_{\ll} = r$ and $h_{\gg} = 0$ it similarly proves claim 1.1.3, while for cases with $n = m$ for which $h_{\ll} = r = h_{\gg}$ it proves the equation obtained from that of claim 1.1.1 by multiplying both sides by $\prod_{k \in [r]} \frac{1}{1 - XY^k}$.

§6. Asymmetric correspondences.

In this section we shall define a generalisation to semistandard r -ribbon tableau of the asymmetric correspondence defined by Knuth, the one that proves the identity $\prod_{i,j} (1 + X_i Y_j) = \sum_{\lambda} s_{\lambda}(X) s_{\lambda^t}(Y)$; it will be spin preserving in the same sense as our symmetric Knuth correspondence. Much of the considerations will be similar to those of the symmetric case, but at the heart of the construction there is a fundamental difference, more substantial than the difference between Knuth’s symmetric and asymmetric constructions. In fact our construction will be a bit simpler for the asymmetric case, since it is based on a direct study of all possible values for λ and κ for given μ, ν .

First we need to give the context in which the construction applies, which involves less specialised instances of the constructions of [Fom5] than we have been considering so far. Specifically, we have been supposing that the horizontal neighbours in the grid, like λ, μ are related by the same relation ‘ \leftarrow_r ’ as the vertical neighbours like λ, ν . There is however nothing in the construction a Knuth correspondence that requires this, and the construction applies equally well when two different relations are used horizontally and vertically. Specifically, we shall consider cases where the two relations are each others transposes, for instance when requiring $\lambda \leftarrow \mu$ and $\nu \leftarrow \kappa$, we shall also require $\lambda^t \leftarrow \nu^t$ and $\mu^t \leftarrow \kappa^t$. Note that the relations ‘ \prec ’ and ‘ \prec_r ’ coincide with their transposes, which is why we shall not consider asymmetric Schensted correspondences (although for appropriate “dual graded graphs” they too can be constructed; see [Fom3] for examples). To facilitate notation we shall write $\lambda \leftarrow \nu$ instead of $\lambda^t \leftarrow \nu^t$, in which case we say ν/λ is a vertical strip, and similarly $\lambda \leftarrow_r \nu$ means $\lambda^t \leftarrow_r \nu^t$, and ν/λ is then said to be a vertical r -ribbon strip. Recall that for ‘ \leftarrow ’ and ‘ \leftarrow_r ’ we have chosen to form the symbols for the opposite relations by rotation rather than by reflection (which

6 Asymmetric correspondences

convention has the advantage that it can be extended without ambiguity to relations written vertically or diagonally in diagrams, even if this does not occur in our paper); by the same token $\nu \rightarrow \lambda$ will mean that ν/λ is a vertical strip. Standardisation of such strips is defined so as to commute with transposition: squares are added from top to bottom.


We shall use the relations ‘ \leftarrow ’ or ‘ \leftarrow_r ’ between shapes associated to horizontally adjacent points of our grid, and ‘ \leftarrow ’ or ‘ \leftarrow_r ’ for vertically adjacent grid points; this is easy to remember, but there is of course no connection between relative positions among squares of individual shapes and relative positions of those shapes as placed on the grid: growths with the opposite convention can equally well be defined. The notion of a shape datum is adapted to the asymmetric context, in the obvious way.

6.1. Definition. Let \mathcal{P} be a graded set equipped with two relations ‘ \leftarrow ’, ‘ \leftarrow_r ’, and S a graded set. An (asymmetric) shape datum for $(\mathcal{P}, \leftarrow, \leftarrow_r, S)$ consists of a family $(b_{\mu,\nu})_{\mu,\nu \in \mathcal{P}}$ of bijections

$$b_{\mu,\nu}: \{ \kappa \in \mathcal{P} \mid \mu \leftarrow \kappa \rightarrow \nu \} \rightarrow \{ (a, \lambda) \in S \times \mathcal{P} \mid \mu \rightarrow \lambda \leftarrow \nu \}, \quad (25)$$

such that $|\kappa| - |\nu| - |\mu| + |\lambda| = |a|$ holds whenever $(a, \lambda) = b_{\mu,\nu}(\kappa)$.

The basic asymmetric shape datum is for $(\mathbf{Y}, \leftarrow, \leftarrow_r, \{0, 1\})$, and is derived from Knuth’s asymmetric construction. In fact it comes in two different flavours, one for each of the natural 1-correspondences for (\mathbf{Y}, \prec) . As before they can be defined using standardisations of κ/μ and κ/ν and computing a partial Schensted-growth across a rectangle, by taking for a the sum of the matrix entries found, and for λ the shape at the top left corner. There will be at most one matrix entry 1, which occurs at the top right corner (where μ is placed) in the case of row insertion, and in the bottom left corner (with ν) in the case of column insertion. For $\mu = (6, 4, 2, 2)$, $\nu = (6, 3, 3, 1, 1)$, $\kappa = (7, 5, 3, 2, 1, 1)$, and using row insertion, one gets:



One cannot really say which flavour matches Knuth's original construction, since that construction does row insertion into a row-strict (transpose semistandard) tableau, with the recording tableau being column-strict (semistandard); with our conventions however the P -symbol will be column-strict and the Q -symbol row-strict. Two different symmetries can be used to map Knuth's construction to our conventions: one may replace (P, Q) either by (P^t, Q^t) or by (Q, P) ; this respectively results in the column-insertion and the row-insertion flavour, and shows that the two are essentially equivalent, unlike what we saw for the RSK and Burge correspondences.

Before considering the ribbon case, we shall reformulate in a more direct way these asymmetric shape data for $(\mathbf{Y}, \leftarrow, \leftarrow, \{0, 1\})$. By studying the sets at both sides of the bijection in (25), we can give a description involving hardly any algorithm at all. For given $\mu, \nu \in \mathbf{Y}$, it is clear that the diagram of any λ satisfying $\mu \rightarrow \lambda \leftarrow \nu$ will be contained in that of $\lambda' = \mu \cap \nu$, and that of any κ satisfying $\mu \leftarrow \kappa \rightarrow \nu$ will contain that of $\kappa' = \mu \cup \nu$. If the right hand side of (25) is to be non-empty one must have $\mu \rightarrow \lambda' \leftarrow \nu$, and if its left hand side is to be non-empty one must have $\mu \leftarrow \kappa' \rightarrow \nu$. Moreover these two conditions are equivalent since $\lambda' \leftarrow \mu \iff \nu \leftarrow \kappa'$ and $\lambda' \leftarrow \nu \iff \mu \leftarrow \kappa'$: the skew diagrams involved are identical. We assume henceforth that these conditions are satisfied, and denote by h the diagram of either of the horizontal strips μ/λ' and κ'/ν , and by v the diagram of either of the vertical strips ν/λ' and κ'/μ . Any squares of the diagram of λ'/λ must be in distinct rows and in distinct columns, nor can they share a column with h or a row with v ; the same must be true for κ/κ' . Let S (respectively T) be the set of squares that can be individually removed from λ' (respectively added to κ'), and that share no column with h , nor any row with v . For instance, in the example $\mu = (6, 4, 2, 2)$, $\nu = (6, 3, 3, 1, 1)$ above one has $\lambda' = (6, 3, 2, 1)$, $\kappa' = (6, 4, 3, 2, 1)$, $h = \{(3, 1), (1, 3)\}$, $v = \{(2, 2), (4, 0)\}$, and hence $S = \{(0, 5), (1, 2), (3, 0)\}$ and $T = \{(0, 6), (1, 4), (3, 2), (5, 0)\}$. Then the diagrams of the λ satisfying $\mu \rightarrow \lambda \leftarrow \nu$ can be obtained by removing from the diagram of λ' any subset of S , and the diagrams of the λ satisfying $\mu \leftarrow \kappa \rightarrow \nu$ can be obtained by adding to the diagram of κ' any subset of T . One easily shows that the squares of h , v , S and T are all on distinct diagonals; moreover, the squares of S and of T are perfectly interleaved in order of increasing diagonal numbers, with an element of T at either end. If one matches each square of S with the previous square of T (for the row-insertion flavour) or with the next square of T (for column-insertion) in this ordering, and matching the remaining square of T with the value of $a \in \{0, 1\}$, one discerns a shape datum is for $(\mathbf{Y}, \leftarrow, \leftarrow, \{0, 1\})$. Indeed, one checks in the example that $\lambda = (5, 3, 2)$ which corresponds to $\{(0, 5), (3, 0)\} \subseteq S$, and $a = 1$, together match $\{(0, 6), (1, 4), (5, 0)\} \subseteq T$, which corresponds to $\kappa = (7, 5, 3, 2, 1, 1)$.

The fact that the squares of S and T are interleaved can be shown by procedures that essentially trace bumping sequences in the Schensted-growth above, but it will be more useful to express S and T directly in terms of $\delta(\mu)$ and $\delta(\nu)$. Let us define for any $\lambda \in \mathbf{Y}$ a doubly infinite sequence $\Sigma(\lambda): \mathbf{Z} \rightarrow \mathbf{N}$, obtained by cumulation from $\delta(\lambda)$, by setting $\Sigma(\lambda)_k = \sum_{i \geq k} \delta(\lambda)_i$. This number can be interpreted graphically as the vertical coordinate of the point where the boundary of λ crosses the diagonal d_k . Then $\lambda \leftarrow \mu$ can be seen to be equivalent to $\Sigma(\lambda)_k \leq \Sigma(\mu)_k \leq \Sigma(\lambda)_{k+1} + 1$ for all $k \in \mathbf{Z}$, and

6 Asymmetric correspondences

$\lambda \leftarrow \nu$ similarly to $\Sigma(\lambda)_k \leq \Sigma(\nu)_k \leq \Sigma(\lambda)_{k-1}$. This can be expressed more elegantly by (termwise) comparison of sequences using the shifted sequences $\Sigma^-(\lambda), \Sigma^+(\lambda)$ defined by $\Sigma^-(\lambda)_{i-1} = \Sigma(\lambda)_i = \Sigma^+(\lambda)_{i+1}$:

$$\lambda \leftarrow \mu \iff \Sigma^+(\mu) - 1 \leq \Sigma(\lambda) \leq \Sigma(\mu) \leq \Sigma^-(\lambda) + 1 \quad (26)$$

$$\lambda \leftarrow \nu \iff \Sigma^-(\nu) \leq \Sigma(\lambda) \leq \Sigma(\nu) \leq \Sigma^+(\lambda) \quad (27)$$

where we have added redundant inequalities for symmetry. Now we see that $\mu \rightarrow \lambda \leftarrow \nu$ is equivalent to $\Sigma^+(\mu) - 1 \leq \Sigma(\lambda) \leq \Sigma(\nu) \leq \Sigma^+(\lambda) \leq \Sigma^+(\mu)$. Defining

$$\Delta(\mu, \nu) = \Sigma^+(\mu) - \Sigma(\nu), \quad (28)$$

that is $\Delta(\mu, \nu)_k = \Sigma(\mu)_{k-1} - \Sigma(\nu)_k$ for $k \in \mathbf{Z}$, we see that $\mu \rightarrow \lambda \leftarrow \nu$ yields a contradiction unless $\Delta(\mu, \nu)_k \in \{0, 1\}$ for all $k \in \mathbf{Z}$, and the same necessary condition is found for $\mu \leftarrow \kappa \rightarrow \nu$. Moreover, if $\Delta(\mu, \nu)_k = 1$ then $\Sigma(\lambda)_k = \Sigma(\nu)_k$, while if $\Delta(\mu, \nu)_{k+1} = 0$ then $\Sigma(\lambda)_k = \Sigma(\mu)_k$, so the term $\Sigma(\lambda)_k$ is completely determined by μ and ν , unless $\Delta(\mu, \nu)_k = 0$ and $\Delta(\mu, \nu)_{k+1} = 1$. In that remaining case one easily shows that $\Sigma(\lambda)_{k-1} = \Sigma(\mu)_{k-1} = \Sigma(\nu)_k$ is exactly one larger than $\Sigma(\lambda)_{k+1} = \Sigma(\nu)_{k+1} = \Sigma(\mu)_k - 1$, so that $\Sigma(\lambda)_k$ may indeed assume either of these values; in other words there is a square of S on the diagonal d_k . For κ , a similar reasoning shows that $\Sigma(\kappa)_k$ is determined by μ and ν unless $\Delta(\mu, \nu)_k = 1$ and $\Delta(\mu, \nu)_{k+1} = 0$, in which case it can assume either of the values $\Sigma(\kappa)_{k-1}$ and $\Sigma(\kappa)_{k+1}$, so that T has a square on the diagonal d_k . We conclude that the squares of S correspond to occurrences of ‘01’ in the bit sequence $\Delta(\mu, \nu)$, while the squares of T correspond to occurrences of ‘10’; the announced perfect interleaving follows by a familiar argument.

That was a bit technical, but it will allow us advance easily to the case of r -ribbons. To find a shape datum for $(\mathbf{Y}, \leftarrow_r, \leftarrow_r, \{0, 1\}^r)$, it suffices to consider only pairs μ, ν whose r -quotients $(\mu^i)_{i \in [r]}, (\nu^i)_{i \in [r]}$ satisfy $\Delta(\mu^i, \nu^i)_k \in \{0, 1\}$ for all $i \in [r]$ and $k \in \mathbf{Z}$, so that shapes λ, κ with $\mu \rightarrow_r \lambda \leftarrow_r \nu$ and $\mu \leftarrow_r \kappa \rightarrow_r \nu$ do exist. Then the sequence that will take the place of $\Delta(\mu, \nu)$ above is $\Delta^r(\mu, \nu)$, defined by

$$\Delta^r(\mu, \nu)_k = \Sigma(\mu)_{k-r} - \Sigma(\nu)_k. \quad (29)$$

To visualise this, imagine two points moving simultaneously, the first along the boundary of ν , the second along the boundary of μ , always keeping r diagonals behind the first (like the fire department ladder truck with independently steering rear wheels); then $\Delta^r(\mu, \nu)$ tracks the vertical distance between the two points as a function of time. From the defining relation $\delta(\lambda)_{i+jr} = \delta(\lambda^i)_j$ of r -quotients one deduces $\Sigma(\lambda)_k = \sum_{i \in [r]} \Sigma(\lambda^i)_{\lceil \frac{k-i}{r} \rceil}$, from which it follows that $\Delta^r(\mu, \nu)_k = \sum_{i \in [r]} \Delta(\mu^i, \nu^i)_{\lceil \frac{k-i}{r} \rceil}$. This shows that $0 \leq \Delta^r(\mu, \nu)_k \leq r$ for all k , and moreover that

$$\delta(\mu)_{k-r} - \delta(\nu)_k = \Delta^r(\mu, \nu)_k - \Delta^r(\mu, \nu)_{k+1} = \Delta(\mu^i, \nu^i)_j - \Delta(\mu^i, \nu^i)_{j+1}, \quad (30)$$

where $k = i + jr$. Let S be the set of indices k for which the members of this equation have the value -1 , and T the set of indices k for which they have the value $+1$. We

denote by $\lambda' = \mu \wedge \nu$ and $\kappa' = \mu \vee \nu$ the meet and join of the shapes μ, ν in (\mathbf{Y}, \leq_r) ; these operations are not the intersection and union of the shapes, but the components of their r -quotients are obtained in that way. Then the set of all shapes λ with $\mu \rightarrow_r \lambda \leftarrow_r \nu$ is in bijection with the set of subsets of S , the shape corresponding to $X \subseteq S$ being obtained by successively removing r -ribbons from λ' with their heads on the diagonals d_k with $k \in X$. Similarly the set of all shapes κ with $\mu \leftarrow_r \kappa \rightarrow_r \nu$ is in bijection with the set of subsets of T , the shape corresponding to $Y \subseteq T$ being obtained by successively adding r -ribbons to κ' with their heads on the diagonals d_k with $k \in Y$. By consideration of the r -quotients one can see that $\#T = \#S + r$ which is sufficient to obtain the existence of asymmetric shape data for $(\mathbf{Y}, \leftarrow_r, \leftarrow_r, \{0, 1\}^r)$. But we want more than that of course: we want a shape datum for which equation (20) holds, so that spin shall be preserved. It is for that purpose that the sequence $\Delta^r(\mu, \nu)$ will prove its real utility.

We must first define what exactly we mean by the spin of a vertical r -ribbon strip. When defining the standardisation of vertical strips, we have required commutation with transposition, and we do so for the standardisation of vertical r -ribbon strips as well, so that their ribbons are added from top right to bottom left (formally: by decreasing index of the diagonal containing the head of the ribbon). We note that skew shapes that happen to be simultaneously a horizontal r -ribbon strip and a vertical r -ribbon strip have two different standardisations (even without considering the possibility that they are an ordinary horizontal or vertical strip as well!). The context will always make clear which form of standardisation is intended, but for the spin, which is defined in terms of the standardisation, we wish to make the distinction clear in the notation, so we shall write $\text{spin}^t(\nu/\lambda)$ for the spin of ν/λ defined using its standardisation as a vertical r -ribbon strip; as for horizontal r -ribbon strip this spin is half the sum of the heights of the ribbons in the standardisation. Note that this is not the spin of the horizontal r -ribbon strip obtained by transposition (this already fails for single r -ribbons), although the two quantities are related. Generalising the notation used in the proof of theorem 5.2.7, we define for any $\lambda \leq_r \mu$ the set $I(\mu/\lambda) \subset \mathbf{Z}$ to be such that the heads of the ribbons of any r -ribbon tableau of shape μ/λ occur on the diagonals d_k for $k \in I(\mu/\lambda)$; this is easily seen to be independent of the r -ribbon tableau.

6.2. Lemma. *Let $\lambda, \mu, \nu, \kappa \in \mathbf{Y}$ with $\mu \rightarrow_r \lambda \leftarrow_r \nu$ and $\mu \leftarrow_r \kappa \rightarrow_r \nu$. Then*

$$\text{spin}^t(\kappa/\mu) + \text{spin}(\kappa/\nu) - \text{spin}(\mu/\lambda) - \text{spin}^t(\nu/\lambda) = \sum_{k \in Y} \Delta^r(\mu, \nu)_{k+1} - \sum_{k \in X} \Delta^r(\mu, \nu)_k. \tag{31}$$

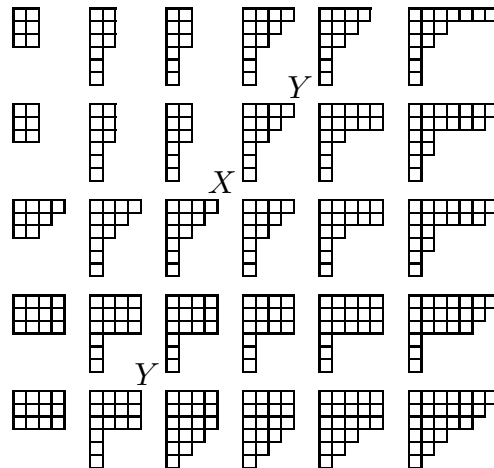
where $X = I(\mu/\lambda) \cap I(\nu/\lambda)$ and $Y = I(\kappa/\mu) \cap I(\kappa/\nu)$.

Proof. Put $\lambda' = \mu \wedge \nu$ and $\kappa' = \mu \vee \nu$, and then $h = I(\mu/\lambda') = I(\kappa'/\nu)$ and $v = I(\nu/\lambda') = I(\kappa'/\mu)$. It follows from consideration of r -quotients that X, h, v , and Y are all disjoint. We shall use a slight variation of a Schensted-growth to do the accounting for us. We consider a grid rectangle whose squares are indexed horizontally by $J = X \cup h \cup Y$ increasing from left to right, and vertically by $I = X \cup v \cup Y$ decreasing from top to bottom; shapes are assigned to its grid points according to the following rule. These shapes σ satisfy $\lambda \leq_r \sigma \leq_r \kappa$, so r -ribbon tableaux of shape σ/λ do not have more than

6 Asymmetric correspondences

one ribbon on any diagonal, and σ is determined by $I(\sigma/\lambda)$, which we shall specify by giving its intersections with X , h , v , and Y . For some grid point, let the rectangle to its top left consist of the set of squares $I' \times J'$ (so $J' \subseteq J$ is an order ideal and $I' \subseteq I$ a dual order ideal); the shape σ assigned to this grid point satisfies $I(\sigma/\lambda) \cap X = (I' \cup J') \cap X$, $I(\sigma/\lambda) \cap h = J' \cap h$, $I(\sigma/\lambda) \cap v = I' \cap v$, and $I(\sigma/\lambda) \cap Y = (I' \cap J') \cap Y$. Then the standardisations of μ/λ and of ν/λ appear along the top and left sides the top edge of the grid, but with repetition of the same shape for any step numbered by an element of Y , while the standardisations of κ/μ and of κ/ν similarly appear along the right and bottom sides, with shapes repeated for any step numbered by an element of X . That the given specification can be met, is proved in growth-diagram style by induction on the sizes of I' and J' , together with the following relation between a square (i, j) and the shapes σ , σ^\uparrow , and σ^\leftarrow , attached to its bottom right, top right, and bottom left corners, respectively. If $i \in v$ then always $\sigma^\uparrow \prec_r \sigma$, while $\sigma^\leftarrow \prec_r \sigma$ whenever $j \in h$. If $i \in X$ (respectively $j \in X$), then $\sigma^\uparrow \prec_r \sigma$ (respectively $\sigma^\leftarrow \prec_r \sigma$) for $j < i$, while $\sigma = \sigma^\uparrow$ (respectively $\sigma = \sigma^\leftarrow$) for $j \geq i$. If $i \in Y$ (respectively $j \in Y$), then $\sigma = \sigma^\uparrow$ (respectively $\sigma = \sigma^\leftarrow$) for $j < i$ while $\sigma^\uparrow \prec_r \sigma$ (respectively $\sigma \prec_r \sigma^\leftarrow$) for $j \geq i$.

Here is an example with $r = 3$, $\mu = (7, 3, 2, 1, 1, 1)$, $\nu = (4, 4, 4)$, so $\lambda' = (4, 3, 2)$, $\kappa' = (7, 4, 4, 1, 1, 1)$, and furthermore $\lambda = (2, 2, 2)$, $\kappa = (7, 6, 5, 3, 2, 1)$; then $X = \{3\}$, $h = \{-3, 6\}$, $v = \{2\}$ and $Y = \{-1, 4\}$, so $J = \{-3, -1, 3, 4, 6\}$ and $I = \{4, 3, 2, -1\}$. The grid squares (k, k) with $k \in X$ or $k \in Y$ are marked X respectively Y .



Around any grid square (i, j) with $i \neq j$ one has the configuration of a Schensted-growth for a square not involving the r -correspondence: opposite sides correspond to ribbons with their head on the same diagonal, or to no ribbon at all. Around a grid square (i, i) with $i \in X$ one has $\sigma = \sigma^\uparrow = \sigma^\leftarrow$, a shape obtained by adding a ribbon to the shape σ^\searrow at its top left corner; around a grid square (i, i) with $i \in Y$ one has $\sigma^\searrow = \sigma^\uparrow = \sigma^\leftarrow \prec_r \sigma$. We prove (31) by monitoring the changes to the “spin” along lattice paths from the bottom left corner of the rectangle to the top right corner, defined as half the sum of the heights of ribbons encountered (whether added or removed). Starting with the path passing through the top left corner, we transform it into the path passing through the bottom right corner; the left hand side of (31) measures the amount

added to the spin during the transformation. Moving the path across a grid square (i, j) with $i \neq j$, does not change the spin. We conclude the proof by showing that the change of the spin at any square numbered (k, k) is equal to the contribution of k to the right hand side of (31). That change is $-\text{ht}(\sigma/\sigma^{\setminus})$ if $k \in X$, and $+\text{ht}(\sigma/\sigma^{\setminus})$ if $k \in Y$. By definition $\text{ht}(\sigma/\sigma^{\setminus}) = \Sigma(\sigma)_{k-r+1} - \Sigma(\sigma)_k$, which we can write as $\Sigma(\sigma)_{k-r} - \Sigma(\sigma)_{k+1} - 1$ since $(\delta(\sigma)_{k-r}, \delta(\sigma)_k) = (0, 1)$. Now $\Sigma(\sigma)_{k-r} = \Sigma(\mu)_{k-r}$ since $\Sigma(\cdot)_{k-r}$ is unaffected by removal of r -ribbons with their head on diagonals d_i with $i \geq k$, and $I(\sigma/\sigma')$ and $I(\mu/\sigma')$ are both contained in $\{i \in \mathbf{Z} \mid i \geq k\}$, where σ' is the shape at the top of the grid straight above σ . Similarly the fact that $\Sigma(\cdot)_{k+1}$ is unaffected by removal of r -ribbons with their head on diagonals d_i with $i \leq k$ implies that $\Sigma(\sigma)_{k+1} = \Sigma(\nu)_{k+1}$, so that $\text{ht}(\sigma/\sigma^{\setminus}) = \Sigma(\mu)_{k-r} - \Sigma(\nu)_{k+1} - 1$. If $k \in X$ then $\Sigma(\nu)_{k+1} + 1 = \Sigma(\nu)_k$ since $\delta(\nu)_k = 1$ due to $X \subseteq S$ as defined below (30), so $\text{ht}(\sigma/\sigma^{\setminus}) = \Delta^r(\mu, \nu)_k$. If $k \in Y$ then $\Sigma(\mu)_{k-r} - 1 = \Sigma(\mu)_{k-r+1}$ since $\delta(\mu)_{k-r} = 1$ due to $Y \subseteq T$, so $\text{ht}(\sigma/\sigma^{\setminus}) = \Delta^r(\mu, \nu)_{k+1}$. Thus in both cases moving across the grid square (k, k) changes the value of the spin by the contribution of k to the right hand side of (31). \square

This lemma allows us to define a spin preserving asymmetric shape datum b^A for $(\mathbf{Y}, \leftarrow_r, \leftarrow_r, \{0, 1\}^r)$ as follows. We match elements of S corresponding to a rise of $\Delta^r(\mu, \nu)$ from h to $h+1$ with elements of T corresponding to a fall of $\Delta^r(\mu, \nu)$ from $h+1$ to h , while matching the remaining element of T of that form with the component a_h of $a \in \{0, 1\}^r$; this is possible since $\lim_{k \rightarrow -\infty} \Delta^r(\mu, \nu)_k = r$ and $\lim_{k \rightarrow +\infty} \Delta^r(\mu, \nu)_k = 0$. Then any κ in the domain of $b_{\mu, \nu}^A$, in other words with $\mu \leftarrow_r \kappa \rightarrow_r \nu$, determines a subset $Y \subseteq T$ which gives rise to a subset $X \subseteq S$ determining λ with $\mu \rightarrow_r \lambda \leftarrow_r \nu$, and to a value $a \in \{0, 1\}^r$. Then matching elements of S and T cancel each others contribution to the right hand side of (31), and the contribution of the remaining elements of T amounts to $n(a)$, whence one obtains the asymmetric counterpart of equation (20). We must still specify the precise matching of elements of S and of T ; of the two equally natural possibilities we choose the one which is most similar to the symmetric case: we match each rise of $\Delta^r(\mu, \nu)$ to the *next* descent back to the same level, leaving the very first descent from $h+1$ to h to match a_h .

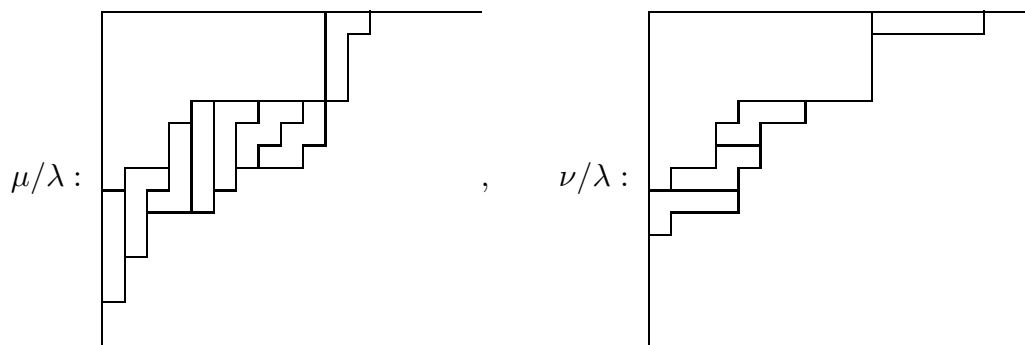
If one represents μ and ν by their edge sequences and the r -ribbon strips $\mu/\lambda, \nu/\lambda, \kappa/\mu, \kappa/\nu$ by the sets $I(\mu/\lambda), \dots, I(\kappa/\nu)$ giving the diagonals containing the heads of the ribbons in their standardisations, then there is a simple procedural description of this asymmetric shape datum. We shall describe it in the “insertion” direction, calculating $\kappa = (b_{\mu, \nu}^A)^{-1}(a, \lambda)$ by specifying the elements of $I(\kappa/\mu)$ and $I(\kappa/\nu)$. The procedure uses a as an initialised r -bit variable; a diagonal index k traverses a sufficiently large interval of \mathbf{Z} in increasing order, while an integer variable h taking values $0 \leq h \leq r$ keeps track of the current value of $\Delta^r(\mu, \nu)_k$. The initial value of k is taken sufficiently small so that $k \leq i$ for all $i \in I(\mu/\lambda) \cup I(\nu/\lambda)$ and $(\mu_{i-r}, \nu_i) = (1, 1)$ for all $i < k$; correspondingly h is initialised to r . For each k the following cases are distinguished. If $(\mu_{k-r}, \nu_k) = (1, 0)$ then necessarily $h > 0$ and $k \notin I(\mu/\lambda) \cup I(\nu/\lambda)$, and one starts by setting $h := h - 1$; if now $a_h = 1$ then it is established that $k \in I(\kappa/\mu)$ and that $k \in I(\kappa/\nu)$, and otherwise that $k \notin I(\kappa/\mu)$ and that $k \notin I(\kappa/\nu)$; finally one sets $a_h := 0$. If $(\mu_{k-r}, \nu_k) = (0, 1)$ then necessarily $h < r$ and $a_h = 0$; in this case it is established that $k \notin I(\kappa/\mu)$ and $k \notin I(\kappa/\nu)$, one sets $a_h := 1$ if $k \in I(\mu/\lambda) \cap I(\nu/\lambda)$ (otherwise $k \notin I(\mu/\lambda) \cup I(\nu/\lambda)$) and

6 Asymmetric correspondences

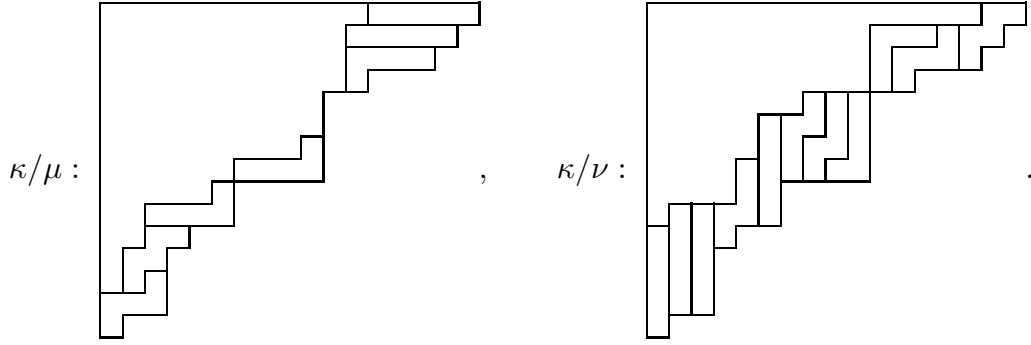
a_h stays 0), and finally one sets $h := h + 1$. In the remaining cases ($\mu_{k-r} = \nu_k$), both h and a are unchanged; it will be established that $k \in I(\kappa/\mu)$ if and only if $k \in I(\nu/\lambda)$, and that $k \in I(\kappa/\nu)$ if and only if $k \in I(\mu/\lambda)$ (at most one of these conditions can hold). Traversal may terminate when $k > i$ for all $i \in I(\mu/\lambda) \cup I(\nu/\lambda)$, and $h = 0$.

This procedure has some similarity to the one given for the symmetric shape datum, in particular if the latter is realised in such a way that the level h is maintained in an incremental fashion. Like we saw for that procedure, $a_i = 0$ here holds for all $h \leq i < r$ after each step, forcing $a = 0$ at the end. This condition also enables a step-by-step inverse procedure, but in the current case that is for a rather simple reason: the act of recording the presence of k in $I(\mu/\lambda)$ and in $I(\nu/\lambda)$ in the bit a_h can only be inverted if a_h previously had a known state (cleared). In other respects as well the current procedure is a lot simpler, notably because h evolves independently of the placement of ribbons.

Let us give an example of this asymmetric shape datum. We take $r = 5$, $\lambda = (10, 10, 10, 10, 4, 3, 3, 1)$, $a = (1, 0, 0, 1, 1)$, and μ, ν are such that the standardisations of μ/λ and ν/λ are



One can read off $\begin{pmatrix} \delta(\mu) \\ \delta(\nu) \end{pmatrix} = \begin{pmatrix} 11011011000101000101101110100000 \\ 11111010001101100100011100000100 \end{pmatrix}$ at indices $-15 \leq i < 17$, which we displayed shifted so as to align $\delta(\mu)_{i-5}$ with $\delta(\nu)_i$; from this one can easily determine the terms of $\Delta^5(\mu, \nu)$ at indices $-10 \leq k \leq 17$, namely $5, 4, 4, 4, 3, 2, 3, 3, 2, 3, 4, 4, 3, 4, 3, 3, 3, 4, 4, 5, 4, 3, 3, 2, 1, 1, 1, 0$. It then follows that $S = \{-5, -2, -1, 2, 6, 8\}$ and $T = \{-10, -7, -6, -3, 1, 3, 9, 10, 12, 13, 16\}$, and the injection $S \rightarrow T$ is given by $-5 \mapsto -3$, $-2 \mapsto 12$, $-1 \mapsto 1$, $2 \mapsto 3$, $6 \mapsto 10$, $8 \mapsto 9$, and the remaining elements of T that match a_4, \dots, a_0 are $-10, -7, -6, 13, 16$. One has $I(\mu/\lambda) = \{-8, -5, -2, 0, 2, 4, 5, 11\}$ and $I(\nu/\lambda) = \{-5, -2, 2, 14\}$, whose intersection gives $X = \{-5, -2, 2\}$; from this and the fact that the non-zero bits a_i are a_4, a_3, a_0 , one obtains $Y = \{-10, -7, -3, 3, 12, 16\}$, and it can be concluded that $I(\kappa/\mu) = \{-10, -7, -3, 3, 12, 14, 16\}$ and $I(\kappa/\nu) = \{-10, -8, -7, -3, 0, 3, 4, 5, 11, 12, 16\}$. Thus one finds $\kappa = (b_{\mu, \nu}^A)^{-1}(a, \lambda) = (17, 16, 15, 12, 10, 10, 10, 10, 6, 6, 4, 3, 3, 3, 1)$, with the following 5-ribbon strips



One checks from these results that $|\kappa/\mu|_r - |\nu/\lambda|_r = 7 - 4 = |a| = 11 - 8 = |\kappa/\nu|_r - |\mu/\lambda|_r$ and that $\text{spin}^t(\kappa/\mu) + \text{spin}(\kappa/\nu) - \text{spin}(\mu/\lambda) - \text{spin}^t(\nu/\lambda) = \frac{7}{2} + \frac{34}{2} - \frac{24}{2} - \frac{3}{2} = 7 = n(a)$ as it should be. The mentioned procedural description of the shape datum can be seen to do essentially the same computation, but in a more orderly left-to-right fashion. The reader may check that in fact everything can be done entirely in terms of edge sequences and in a single pass from left to right, finding $I(\mu/\lambda)$ and $I(\nu/\lambda)$ on the fly by transforming a copy of $\delta(\lambda)$ into $\delta(\mu)$ and a copy of $\delta(\nu)$ into $\delta(\lambda)$; the information obtained about $I(\kappa/\nu)$ can be used to simultaneously transform another copy of $\delta(\nu)$ into $\delta(\kappa)$. Such a description does not appear to give much new insight into correspondence however. We shall content ourselves here with summarising the result found.

6.3. Theorem. *For every $r > 0$ there exists an asymmetric shape datum b^A for $(\mathbf{Y}, \leftarrow_r, \leftarrow_r, \{0, 1\}^r)$ such that whenever $(a, \lambda) = b_{\mu, \nu}^A(\kappa)$, then $\text{spin}^t(\kappa/\mu) + \text{spin}(\kappa/\nu) = \text{spin}(\mu/\lambda) + \text{spin}^t(\nu/\lambda) + n(a)$. Hence*

$$\check{D}_{q^{1/2}, Y}^r \circ U_{q^{1/2}, X}^r = (U_{q^{1/2}, X}^r \circ \check{D}_{r, Y}) \prod_{i \in [r]} (1 + q^i XY), \tag{32}$$

where $U_{q^{1/2}, X}^r$ is the spin generating series for horizontal r -ribbon strips defined in equation (21), while $\check{D}_{q^{1/2}, Y}^r = \sum_{\mu \leftarrow_r \lambda} q^{\text{spin}^t(\lambda/\mu)} Y^{|\lambda/\mu|_r} \mu$ is the similar spin generating series for vertical r -ribbon strips. \square

6.4. Corollary. *For $r \in \mathbf{N}_{>0}$ there exists a spin preserving asymmetric Knuth correspondence for r -ribbon tableaux: for any r -core γ , and any $m, n \in \mathbf{N}$, there is a bijection between $m \times n$ matrices A with entries in $\{0, 1\}^r$, and pairs consisting of a semistandard r -ribbon tableau P and a transpose semistandard r -ribbon tableau Q of equal shape with $\text{wt}(P) = (\sum_{i \in [m]} |A_{i, j}|)_{j \in [n]}$ and $\text{wt}(Q) = (\sum_{j \in [n]} |A_{i, j}|)_{i \in [m]}$, which is such that in addition one has $\text{spin}(P) + \text{spin}^t(Q) = \sum_{i, j \in [m] \times [n]} n(A_{i, j})$. \square*

This corollary gives a counterpart of (24), which uses besides $G_{\lambda/\gamma}^{(r)}(q^{\frac{1}{2}}, X)$ the spin generating series $\check{G}_{\lambda/\gamma}^{(r)}(q^{\frac{1}{2}}, Y) = \sum_Q q^{\text{spin}^t(Q)} Y^{\text{wt}(Q)}$ of transpose semistandard r -ribbon tableaux of shape λ/γ ; it reads

$$\prod_{i, j \in \mathbf{N}} \prod_{k \in [r]} (1 + q^k X_i Y_j) = \sum_{\lambda \in \mathbf{RH}(\gamma)} G_{\lambda/\gamma}^{(r)}(q^{\frac{1}{2}}, X) \check{G}_{\lambda/\gamma}^{(r)}(q^{\frac{1}{2}}, Y). \tag{33}$$

§7. Conclusion.

We have defined constructions showing that symmetric and asymmetric Knuth correspondences for semistandard r -ribbon tableaux, of which trivial examples can be obtained using the r -quotient map, can be refined so as to satisfy the additional condition that the spins of tableaux be respected. This requires defining the basic ingredient (the shape datum) of the construction by a method radically different from the traditional one of applying a Schensted correspondence to standardisations. Instead, the method at the heart of the Shimozono-White r -correspondence, tracing level changes of a function derived from edge sequences, is extended to deal with many ribbons at the same time. This emphasises the importance of edge sequences for the understanding of r -ribbon tableaux, even when the r -quotient map is not used.

References.

- [Bur] W. H. Burge, “Four correspondences between graphs and generalized Young tableaux”, *J. Combin. Theory, Ser. A* **17**, (1974), 12–30.
- [CaLe] C. Carré and B. Leclerc, “Splitting the square of a Schur function into its symmetric and antisymmetric parts”, *J. of algebraic combinatorics* **4**, (1995), 201–231.
- [Fom1] S. V. Fomin, “Generalised Robinson-Schensted-Knuth correspondence”, *Journal of Soviet Mathematics* **41**, (1988), 979–991. (Translation from *Записки научных семинаров ЛОМИ* **155** (1986) 156–175).
- [Fom2] S. Fomin, *Dual graphs and Schensted correspondences*, Proceedings of the 4th International conference on Formal power series and Algebraic combinatorics, Montreal, (1992).
- [Fom3] S. Fomin, “Duality of graded graphs”, *J. Algebr. Combinatorics* **3**, (1994), 357–404.
- [Fom4] S. Fomin, “Schensted algorithms for dual graded graphs”, *J. Algebr. Combinatorics* **4**, (1995), 5–45.
- [Fom5] S. Fomin, “Schur operators and Knuth correspondences”, *J. Combin. Theory, Ser. A* **72**, (1995), 277–292.
- [FoSt] S. Fomin and D. Stanton, *Rim hook lattices*, Institut Mittag-Leffler report No. 23, (1991/92).
- [KMS] M. Kashiwara, T. Miwa, E. Stern, “Decomposition of q -deformed Fock spaces”, *Selecta Mathematica, New Series* **1 (no. 4)**, (1995), 787–805.
- [Knu] D. E. Knuth, “Permutations, matrices and generalized Young tableaux”, *Pacific J. of Math.* **34**, (1970), 709–727.
- [Lam] T. Lam, *Ribbon tableaux and the Heisenberg algebra*, arXiv:math.QA/0310250.
- [LLT] A. Lascoux, B. Leclerc, J.-Y. Thibon, “Ribbon tableaux, Hall-Littlewood functions, quantum affine algebras, and unipotent varieties”, *J. Math. Phys.* **38**, (1997), 1041–1068.
- [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, “Edge sequences, ribbon tableaux, and an action of affine permutations”, *European J. of Combinatorics* **20**, (1999), 397–426.
- [Rob] G. de B. Robinson, “On the representations of the symmetric group”, *American J. of Math.* **60**, (1938), 745–760.
- [SaSt] B. E. Sagan and R. P. Stanley, “Robinson-Schensted Algorithms for Skew Tableaux”, *J. Combin. Theory, Ser. A* **55**, (1990), 161–193.
- [Sche] C. Schensted, “Longest increasing and decreasing subsequences”, *Canadian J. of Math.* **13**, (1961), 179–191.
- [ShW1] M. Shimozono and D. E. White, “A Color-to-Spin Domino Schensted Algorithm”, *Electronic J. of Combinatorics* **8 (no. 1)**, R21, (2001), 50 pp.
- [ShW2] M. Shimozono and D. E. White, “Color-to-spin ribbon Schensted algorithms”, *Discrete Math.* **246**, (2002), 295–316.
- [Stan1] R. P. Stanley, “Differential posets”, *J. of the Am. Math. Soc.* **1**, (1988), 919–961.
- [Stan2] R. P. Stanley, *Enumerative Combinatorics Volume 2*, Cambridge Studies in Advanced Mathematics 62, Cambridge University Press, 1999.
- [StWh] D. W. Stanton and D. E. White, “A Schensted algorithm for rim hook tableaux”, *J. Combin. Theory, Ser. A* **40**, (1985), 211–247.