

Math 701 Fall 2026, Lecture 7: Skew Schur polynomials continued

website: <https://www.cip.ifi.lmu.de/~grinberg/t/26fs/>

(GPT-6 was used to create some of the examples below and to rewrite Remark 1.7.4.)

1.6.1. The Bender–Knuth involutions (continued)

We continue where we left off in Lecture 6:

Finishing the proof of Theorem 1.6.14. We continue with the notation from Lecture 6. In particular, fix $k \in [N - 1]$. We now define a map

$$\beta_k : \text{SSYT}(\lambda/\mu, N) \rightarrow \text{SSYT}(\lambda/\mu, N)$$

as follows: Given $T \in \text{SSYT}(\lambda/\mu, N)$, we perform the following procedure:

- For each row of T , if there are a free k 's and b free $(k + 1)$'s in this row, then we replace them by b many k 's followed by a many $(k + 1)$'s, from left to right in the same cells. All other entries remain unchanged.

We define $\beta_k(T)$ to be the result of this procedure. We still need to prove the following:

Observation 2: This map β_k is well-defined, i.e., the tableau $\beta_k(T)$ belongs to $\text{SSYT}(\lambda/\mu, N)$.

Before we prove this, let us give an example of β_k .

Example 1.6.17. Let $k = 2$ and $N = 4$. We use the tableau T from Example 1.6.16 in Lecture 6, of shape $(9, 8, 6) / (3, 1)$. The map β_2 acts as follows:

$$T = \begin{array}{cccccc} & & 1 & 1 & 1 & \mathbf{2} & \mathbf{3} & 4 \\ & 1 & \mathbf{2} & \mathbf{2} & \mathbf{2} & \mathbf{3} & \mathbf{3} & 4 \\ 1 & \mathbf{2} & \mathbf{3} & 4 & 4 & 4 & & \end{array} \mapsto \beta_2(T) = \begin{array}{cccccc} & & 1 & 1 & 1 & \mathbf{2} & \mathbf{2} & 4 \\ & 1 & \mathbf{2} & \mathbf{2} & \mathbf{3} & \mathbf{3} & \mathbf{3} & 4 \\ 1 & \mathbf{3} & \mathbf{3} & 4 & 4 & 4 & & \end{array}.$$

As before, the 2's and 3's are boldfaced, and the matched entries are shaded. The shaded entries stay unchanged. The free blocks change as follows:

$$\begin{array}{ccc} \text{row 1} & \text{row 2} & \text{row 3} \\ (3) \mapsto (2) & (2, 2, 3) \mapsto (2, 3, 3) & (2) \mapsto (3) \end{array}$$

All entries other than 2 and 3 stay unchanged as well. Thus, the five 2's and four 3's in T become four 2's and five 3's in $\beta_2(T)$. Notice that in the second row, we replace 223 by 233, not by 332: the new free block must still be weakly increasing.

Proof of Observation 2. The entries of $\beta_k(T)$ are clearly in $[N]$, since $k, k+1 \in [N]$. It remains to show that the tableau $\beta_k(T)$ is semistandard.

- Why are its rows weakly increasing? Consider a row of T . In $\beta_k(T)$, some of its free entries may have changed from k to $k+1$ or vice versa; no other entries changed. But Observation 1 from Lecture 6 shows that the matched k 's are to the left of the free k 's and the matched $(k+1)$'s are to the right of the free $(k+1)$'s; thus, all the free entries form a single contiguous block in our row. By definition, this block stays weakly increasing in $\beta_k(T)$. All the entries in this block are either k 's or $(k+1)$'s in both T and $\beta_k(T)$; all the entries to the left of this block are $\leq k$; all the entries to the right of this block are $\geq k+1$. Thus, the row is still weakly increasing in $\beta_k(T)$.
- Why are its columns strictly increasing? Because we only changed free k 's and free $(k+1)$'s, and there are no other k 's or $(k+1)$'s in their columns to mess up the inequalities. (For instance, changing a free k to a $k+1$ still keeps it smaller than the neighboring entry to its south (if there is one), since that entry cannot be a $k+1$ because the k was free.)

So $\beta_k(T)$ is semistandard, and Observation 2 is proved. $\square$

Observation 3: This map β_k is an involution, i.e., we have $\beta_k \circ \beta_k = \text{id}$.

Proof of Observation 3. We must show that

$$\beta_k(\beta_k(T)) = T \quad \text{for each } T \in \text{SSYT}(\lambda/\mu, N).$$

Fix $T \in \text{SSYT}(\lambda/\mu, N)$. To compute $\beta_k(T)$ from T , we identify the free and matched k 's and $(k+1)$'s, and flip the numbers of free k 's and free $(k+1)$'s in each row. Note that the matched entries of T are unchanged in $\beta_k(T)$ and thus, in particular, remain matched in $\beta_k(T)$ (since their "partners" also don't change). On the other hand, the free entries of T are replaced by new entries, which are again free (since all the other entries in their respective columns are $< k$ or $> k+1$, and thus don't change). Hence, if a row of T had a free k 's and b free $(k+1)$'s, then the same row of $\beta_k(T)$ has b free k 's and a free $(k+1)$'s.

Now, when we compute $\beta_k(\beta_k(T))$ from $\beta_k(T)$ (by the same procedure that gave us $\beta_k(T)$ from T), we flip the numbers of free k 's and free $(k+1)$'s in each row again. Hence, if a row of T had a free k 's and b free $(k+1)$'s, then the same row of $\beta_k(T)$ has b free k 's and a free $(k+1)$'s (as we saw above), and therefore the same row of $\beta_k(\beta_k(T))$ will have a free k 's and b free $(k+1)$'s again, which is the same distribution that T used to have. Since the other entries in that row don't change at all, this shows that the whole row is the same in $\beta_k(\beta_k(T))$ as it was in T . We have proved this for every row; thus, $\beta_k(\beta_k(T)) = T$. $\square$

Observation 3 shows that β_k is a bijection (since any involution is a bijection).

Observation 4: Let $T \in \text{SSYT}(\lambda/\mu, N)$. Then, we have the equalities

$$\begin{aligned} (\# \text{ of } k\text{'s in } \beta_k(T)) &= (\# \text{ of } (k+1)\text{'s in } T) & \text{and} \\ (\# \text{ of } (k+1)\text{'s in } \beta_k(T)) &= (\# \text{ of } k\text{'s in } T) \end{aligned}$$

and furthermore

$$(\# \text{ of } i\text{'s in } \beta_k(T)) = (\# \text{ of } i\text{'s in } T) \quad \text{for each } i \notin \{k, k+1\}.$$

Proof of Observation 4. In the proof of Observation 3, we observed that, as we transform T into $\beta_k(T)$, all matched entries stay matched and remain unchanged, while all free entries remain free (though they can change their values). Thus, the matched k 's in $\beta_k(T)$ are precisely the matched k 's in T . Hence,

$$\begin{aligned} (\# \text{ of matched } k\text{'s in } \beta_k(T)) &= (\# \text{ of matched } k\text{'s in } T) \\ &= (\# \text{ of matched } (k+1)\text{'s in } T) \end{aligned}$$

(as we saw above, back in Lecture 6).

Now let $i \geq 1$. Recall again from our proof of Observation 3 that when we transform T into $\beta_k(T)$, all free entries remain free (though they can change their values), while all matched entries remain matched. Hence, if the i -th row of T has a free k 's and b free $(k+1)$'s, then, in $\beta_k(T)$, these entries become b free k 's and a free $(k+1)$'s, and there are no further free k 's in the i -th row of $\beta_k(T)$ (since matched entries of T remain matched); therefore the i -th row of $\beta_k(T)$ has b free k 's. That is,

$$\begin{aligned} (\# \text{ of free } k\text{'s in the } i\text{-th row of } \beta_k(T)) \\ = (\# \text{ of free } (k+1)\text{'s in the } i\text{-th row of } T). \end{aligned}$$

Summing this equality over all $i \geq 1$, we obtain

$$(\# \text{ of free } k\text{'s in } \beta_k(T)) = (\# \text{ of free } (k+1)\text{'s in } T).$$

However, each k in $\beta_k(T)$ is either free or matched. Thus,

$$\begin{aligned} (\# \text{ of } k\text{'s in } \beta_k(T)) &= \underbrace{(\# \text{ of matched } k\text{'s in } \beta_k(T))}_{=(\# \text{ of matched } (k+1)\text{'s in } T)} + \underbrace{(\# \text{ of free } k\text{'s in } \beta_k(T))}_{=(\# \text{ of free } (k+1)\text{'s in } T)} \\ &= (\# \text{ of matched } (k+1)\text{'s in } T) + (\# \text{ of free } (k+1)\text{'s in } T) \\ &= (\# \text{ of } (k+1)\text{'s in } T). \end{aligned}$$

This proves the first equality of Observation 4. The second equality is similar; the third is trivial (since no entries other than k and $k+1$ change). $\square$

Observation 5: For each $T \in \text{SSYT}(\lambda/\mu, N)$, we have

$$x_{\beta_k(T)} = s_k \cdot x_T.$$

Proof of Observation 5. We have

$$x_T = \prod_{i=1}^N x_i^{(\# \text{ of } i\text{'s in } T)} \quad \text{and} \quad x_{\beta_k(T)} = \prod_{i=1}^N x_i^{(\# \text{ of } i\text{'s in } \beta_k(T))}.$$

These two monomials have largely the same exponents, by Observation 4; only the exponents on x_k and on x_{k+1} get swapped. This is precisely saying that one monomial turns into the other when we switch the variables x_k and x_{k+1} . In other words, the action of s_k takes one monomial into the other. This proves Observation 5. $\square$

Now, by the definition of $s_{\lambda/\mu}$, we have

$$\begin{aligned} s_{\lambda/\mu} &= \sum_{T \in \text{SSYT}(\lambda/\mu, N)} x_T = \sum_{T \in \text{SSYT}(\lambda/\mu, N)} x_{\beta_k(T)} \\ &\quad \left(\begin{array}{c} \text{here, we substituted } \beta_k(T) \text{ for } T, \text{ since } \beta_k \text{ is a bijection} \\ \text{(by Observation 3)} \end{array} \right) \\ &= \sum_{T \in \text{SSYT}(\lambda/\mu, N)} s_k \cdot x_T \quad (\text{by Observation 5}) \\ &= s_k \cdot \underbrace{\sum_{T \in \text{SSYT}(\lambda/\mu, N)} x_T}_{= s_{\lambda/\mu}} = s_k \cdot s_{\lambda/\mu}. \end{aligned}$$

That is, $s_k \cdot s_{\lambda/\mu} = s_{\lambda/\mu}$. Since we have proved this for each $k \in [N-1]$, we conclude (by Lemma 1.2.9 from Lecture 3) that $s_{\lambda/\mu}$ is symmetric. So Theorem 1.6.14 is proved. $\square$

As we mentioned before, if λ is an N -partition and $\mu = \emptyset = ()$, then $s_{\lambda/\mu} = s_\lambda$, so we conclude from Theorem 1.6.14 that s_λ is symmetric. Thus, Theorem 1.5.22 (a) is proved.

1.7. The Littlewood–Richardson rule

What happens when we multiply two Schur polynomials? Here are two examples:

$$\begin{aligned} s_{(2,1)} s_{(1,1)} &= s_{(2,1,1,1)} + s_{(2,2,1)} + s_{(3,1,1)} + s_{(3,2)}; \\ s_{(2,1)} s_{(2,1)} &= s_{(2,2,2)} + 2s_{(3,2,1)} + s_{(3,3)} + s_{(4,1,1)} \\ &\quad + s_{(4,2)} + s_{(2,2,1,1)} + s_{(3,1,1,1)}. \end{aligned}$$

In both cases, we have expanded the product as a $\mathbb{Z}$ -linear combination of Schur polynomials. The fact that such an expansion always exists is not too

surprising; indeed, the family $(s_\lambda)_\lambda$ is an N -partition is a basis of $\mathcal{S}$, as we will later show. That the same expansion works for every N is also unsurprising once we get to the $N \rightarrow \infty$ case (in the next chapter). But the fact that there are no negative coefficients in this expansion requires explanation! Historically, this expansion appeared in Schur's and Weyl's work on representations of $GL(N)$, which explained the nonnegativity of its coefficients in terms of representation theory. We shall instead explain the nonnegativity of the coefficients via an entirely combinatorial formula known as the **Littlewood–Richardson rule**. This rule was the holy grail of algebraic combinatorics for decades in the mid-20th century, but since Stembridge's proof in 2002 [Stembr02], it has become so accessible that we can prove it as one of our first properties of Schur polynomials.

To state it, we will need a bit more notation.

Definition 1.7.1. Let λ and μ be two partitions. Let $T \in \text{SSYT}(\lambda/\mu, N)$. We define the **content** of T to be the N -tuple $\text{cont } T = (a_1, a_2, \dots, a_N)$, where

$$a_i = (\# \text{ of } i\text{'s in } T) = |\{c \in Y(\lambda/\mu) \mid T(c) = i\}|.$$

For example, if $N = 5$, then

$$\text{cont } \begin{array}{|c|c|c|c|} \hline 1 & 1 & 2 & 4 \\ \hline 3 & 3 & & \\ \hline \end{array} = (2, 1, 2, 1, 0).$$

Note that

$$x_T = x^{\text{cont } T} \quad \text{for each } T \in \text{SSYT}(\lambda/\mu, N),$$

where, for any N -tuple $(a_1, a_2, \dots, a_N) \in \mathbb{N}^N$, we set $x^{(a_1, a_2, \dots, a_N)} = \prod_{i=1}^N x_i^{a_i}$.

Definition 1.7.2. Let λ/μ be a skew partition. Let T be a tableau of shape λ/μ . Let j be a positive integer. Then, $\text{col}_{\geq j} T$ means the restriction of T to columns $j, j+1, j+2, \dots$ (that is, the result of removing the first $j-1$ columns from T). The remaining cells retain their original coordinates. Thus,

$$\text{col}_{\geq j} T = T|_{\{(u,v) \in Y(\lambda/\mu) \mid v \geq j\}}.$$

This is again a skew tableau (of shape $\tilde{\lambda}/\tilde{\mu}$, where $\tilde{\lambda} = (\tilde{\lambda}_1, \tilde{\lambda}_2, \dots, \tilde{\lambda}_M)$ and $\tilde{\mu} = (\tilde{\mu}_1, \tilde{\mu}_2, \dots, \tilde{\mu}_M)$ are two partitions given by $M = \ell(\lambda)$ and $\tilde{\lambda}_i = \max\{j-1, \lambda_i\}$ and $\tilde{\mu}_i = \max\{j-1, \mu_i\}$ for all $1 \leq i \leq \ell(\lambda)$).

For example, $\text{col}_{\geq 1} T = T$, whereas $\text{col}_{\geq j} T$ is empty for any sufficiently large j .

For example,

$$\text{col}_{\geq 3} \begin{array}{|c|c|c|} \hline 1 & 1 & 2 \\ \hline 1 & 2 & 3 \\ \hline 3 & 4 & \\ \hline \end{array} = \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}.$$

The final ingredient for the Littlewood–Richardson rule is a property of some semistandard tableaux known as ν -**Yamanouchiness**:

Definition 1.7.3. Let λ/μ be a skew partition, and let ν be an N -partition. A tableau $T \in \text{SSYT}(\lambda/\mu, N)$ is said to be ν -**Yamanouchi** if it has the property that for each positive integer j ,

the N -tuple $\nu + \text{cont}(\text{col}_{\geq j} T)$ is an N -partition (i.e., weakly decreasing).

Remark 1.7.4. Here is a memorable metaphor for ν -Yamanouchiness:

Imagine an election with N candidates labelled $1, 2, \dots, N$. The entries of T are votes waiting to be counted, with each entry k representing one vote for candidate k . Some votes have already been counted: candidate i starts with ν_i votes.

We count the votes in T one column at a time, starting with the rightmost column and working our way to the left. For example, a column with entries 2 and 4 contributes one vote to candidate 2 and one to candidate 4. Each column contributes at most one vote to any given candidate, since its entries are strictly increasing. After counting all votes from any single column, we take a break before we move on to the next column.

The tableau T is ν -Yamanouchi exactly when, during each break, candidate 1 has at least as many votes as candidate 2, candidate 2 has at least as many as candidate 3, and so on. This ordering holds at the start because ν is an N -partition. (Ties are allowed.)

Example 1.7.5.

(a) Let $N = 3$ and $\nu = \emptyset = (0, 0, 0)$. Which of the following seven tableaux are ν -Yamanouchi?

$$T_1 = \begin{array}{|c|c|} \hline 1 & 1 \\ \hline 2 & 2 \\ \hline \end{array}$$

$$T_2 = \begin{array}{|c|c|} \hline 1 & 1 \\ \hline 2 & 3 \\ \hline \end{array}$$

$$T_3 = \begin{array}{|c|c|} \hline 1 & 1 \\ \hline 1 & 2 \\ \hline \end{array}$$

$$T_4 = \begin{array}{|c|} \hline 1 \\ \hline 1 \\ \hline 2 \\ \hline \end{array}$$

$$T_5 = \begin{array}{|c|} \hline 1 \\ \hline 1 \\ \hline 3 \\ \hline \end{array}$$

$$T_6 = \begin{array}{|c|c|c|} \hline & 1 & 1 \\ \hline & 1 & 2 & 2 \\ \hline 2 & 3 & & \end{array}$$

$$T_7 = \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 2 & 2 \\ \hline \end{array}$$

Answer: T_1, T_3, T_4 and T_6 are; T_2, T_5 and T_7 are not. Here is a column-

by-column check. For the sake of brevity, we set

$$C_j(T) := \nu + \text{cont}(\text{col}_{\geq j} T) = \text{cont}(\text{col}_{\geq j} T)$$

(since $\nu = \emptyset$). We can compute these easily by starting on the very right (none of our tableaux goes beyond column 4, so we have $C_j(T) = \text{cont}(\text{col}_{\geq j} T) = 0$ for all $j \geq 5$) and adding one column at a time (since $C_j(T) = C_{j+1}(T) + \text{cont}(\text{col}_j T)$, where $\text{col}_j T$ is the j -th column of T). Thus we obtain the following contents. The boldfaced tuples are the ones that are not weakly decreasing:

	$C_4(T)$	$C_3(T)$	$C_2(T)$	$C_1(T)$
T_1	(0,0,0)	(1,0,0)	(2,1,0)	(2,2,0)
T_2	(0,0,0)	(1,0,0)	(2,0,1)	(2,1,1)
T_3	(0,0,0)	(1,0,0)	(2,1,0)	(3,1,0)
T_4	(0,0,0)	(0,0,0)	(1,0,0)	(2,1,0)
T_5	(0,0,0)	(0,0,0)	(1,0,0)	(2,0,1)
T_6	(1,1,0)	(2,2,0)	(3,2,1)	(3,3,1)
T_7	(0,0,0)	(0,1,0)	(1,2,0)	(1,3,0)

For $j \geq 5$, all these contents are (0,0,0). Thus, the table checks every condition in the definition. In particular, T_2 fails at $j = 2$ because there is a 3 but no 2 in its last two columns; T_5 fails at $j = 1$ for similar reasons. The tableau T_7 already fails at $j = 3$, since its last column contains a 2 but no 1. Notice that the full content of T_2 is weakly decreasing; checking just the full content is therefore not enough.

Note that any semistandard tableau T that is $\emptyset$ -Yamanouchi will automatically be ν -Yamanouchi for any N -partition ν (because the entrywise sum of two N -partitions is an N -partition).

- (b) Now let us vary ν . What ν should you take to make T_2 and T_5 become ν -Yamanouchi? Try to find a ν that is as small in size as possible.

We can take $\nu = (1,1,0)$. After adding each nonempty column, from right to left, the tuples $\nu + \text{cont}(\text{col}_{\geq j} T)$ are

$$\begin{array}{l|l} T_2 & (2,1,0) \quad (3,1,1) \quad (3,2,1) \\ T_5 & (2,1,0) \quad (3,1,1) \end{array}$$

All these tuples are weakly decreasing, as is the initial tuple ν . This proves that both tableaux are ν -Yamanouchi.

Why is size 2 minimal? For T_2 at $j = 2$, and for T_5 at $j = 1$, the condition requires $\nu + (2,0,1)$ to be weakly decreasing. Thus, $\nu_2 \geq \nu_3 + 1 \geq 1$. Since ν itself is weakly decreasing, we also have $\nu_1 \geq \nu_2 \geq 1$. Therefore,

$|\nu| \geq 2$, and our choice $\nu = (1, 1, 0)$ attains this bound for each of the two tableaux.

To ensure that T_7 is ν -Yamanouchi, we can take $\nu = (2, 0, 0)$. The successive tuples $\nu + \text{cont}(\text{col}_{\geq j} T)$, again after adding each column from right to left, are

$$(2, 1, 0), \quad (3, 2, 0), \quad (3, 3, 0),$$

so this choice works. It also has the smallest possible size: the full content of T_7 is $(1, 3, 0)$, so any suitable ν must satisfy $\nu_1 + 1 \geq \nu_2 + 3$, or equivalently $\nu_1 \geq \nu_2 + 2$. Hence $|\nu| \geq 2$.

We can finally state the Littlewood–Richardson rule:

Theorem 1.7.6 (Zelevinsky’s generalized Littlewood–Richardson rule, in Yamanouchi form for finite N).

Let λ/μ be a skew partition and let ν be an N -partition. Then,

$$s_\nu \cdot s_{\lambda/\mu} = \sum_{\substack{T \in \text{SSYT}(\lambda/\mu, N); \\ T \text{ is } \nu\text{-Yamanouchi}}} s_{\nu + \text{cont } T}.$$

Some comments are warranted:

- The term $s_{\nu + \text{cont } T}$ in the above sum is always well-defined: Indeed, if T is ν -Yamanouchi, then

$$\nu + \text{cont } T = \nu + \text{cont}(\text{col}_{\geq 1} T)$$

is an N -partition.

- By picking $\mu = \emptyset$ in Theorem 1.7.6, we get a formula for $s_\nu \cdot s_\lambda$, that is, for any product of two Schur polynomials as a sum of Schur polynomials.

Before we get to the proof of Theorem 1.7.6, let us first get a simple fact out of our way:

Proposition 1.7.7. If λ is a partition of length $> N$, then $s_\lambda = 0$.

Proof. Let λ be a partition of length $> N$. Note that the length of λ is precisely the number of cells in the first column of $Y(\lambda)$. Hence, the first column of $Y(\lambda)$ has more than N cells (since λ has length $> N$). Therefore, any SSYT of shape λ must have more than N distinct entries in its first column (since the “strictly increasing down each column” condition in the definition of an SSYT forces all the entries in a given column to be distinct). But this is impossible if all entries

of the SSYT are to belong to $[N]$ (since we cannot find more than N distinct numbers in $[N]$). Thus, there exists no SSYT of shape λ whose entries belong to $[N]$. In other words, the set $\text{SSYT}(\lambda, N)$ is empty. Thus, s_λ is an empty sum (according to its definition) and equals 0. $\square$

In contrast, a skew Schur polynomial $s_{\lambda/\mu}$ can be nonzero even if λ and μ are arbitrarily long. For instance, for any positive integer k , setting $\lambda = (k, k-1, \dots, 2, 1)$ and $\mu = (k-1, k-2, \dots, 1)$, we can fill the k cells of λ/μ with 1's, and this always produces a tableau $T \in \text{SSYT}(\lambda/\mu, N)$ for any $N \geq 1$, so that $s_{\lambda/\mu}$ contains at least an x_1^k term.

References

- [21s] Darij Grinberg, *An Introduction to Algebraic Combinatorics (Math 531, Winter 2024 lecture notes)*, 22 July 2026.
<http://www.cip.ifi.lmu.de/~grinberg/t/21s/lecs.pdf>
- [Stembr02] John R. Stembridge, *A Concise Proof of the Littlewood-Richardson Rule*, *Electronic Journal of Combinatorics* **9** (2002), #N5
-