

This is the author's version of a work that was accepted for publication in the Journal of Algebraic Combinatorics. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in *J. Algebraic Combin.* **45** (2017) 1041–1082.

<http://dx.doi.org/10.1007/s10801-016-0734-2>.

The final publication is available at www.springerlink.com.

Dyck tilings and the homogeneous Garnir relations for graded Specht modules

Matthew Fayers

Queen Mary University of London, Mile End Road, London E1 4NS, U.K.

m.fayers@qmul.ac.uk

2010 Mathematics subject classification: 05E10, 20C30, 20C08

Abstract

Suppose λ and μ are integer partitions with $\lambda \supseteq \mu$. Kenyon and Wilson have introduced the notion of a *cover-inclusive Dyck tiling* of the skew Young diagram $\lambda \setminus \mu$, which has applications in the study of double-dimer models. We examine these tilings in more detail, giving various equivalent conditions and then proving a recurrence which we use to show that the entries of the transition matrix between two bases for a certain permutation module for the symmetric group are given by counting cover-inclusive Dyck tilings. We go on to consider the inverse of this matrix, showing that its entries are determined by what we call *cover-expansive Dyck tilings*. The fact that these two matrices are mutual inverses allows us to recover the main result of Kenyon and Wilson.

We then discuss the connections with recent results of Kim et al, who give a simple expression for the sum, over all μ , of the number of cover-inclusive Dyck tilings of $\lambda \setminus \mu$. Our results provide a new proof of this result. Finally, we show how to use our results to obtain simpler expressions for the *homogeneous Garnir relations* for the universal Specht modules introduced by Kleshchev, Mathas and Ram for the cyclotomic quiver Hecke algebras.

1 Introduction

The motivation for this paper is the study of the modular representation theory of the symmetric group, and more generally the representation theory of the cyclotomic Hecke algebra of type A . This area has recently been revolutionised by the discovery by Brundan and Kleshchev of new presentations for these algebras, which show in particular that the algebras are non-trivially $\mathbb{Z}$ -graded. The contribution in the present paper concerns the definition of the *Specht modules*, which play a central role in the representation theory of cyclotomic Hecke algebras. These modules have been studied within the graded setting by Brundan, Kleshchev and Wang, and developed further by Kleshchev, Mathas and Ram, who have given a presentation for each Specht module with a single generator and a set of homogeneous relations. These relations include homogeneous analogues of the classical Garnir relations for the symmetric group, which allow the Specht module to be expressed as a quotient of a ‘row permutation module’. Although the homogeneous Garnir relations are in some sense simpler than their classical counterparts, their statement in [KMR] is awkward in that the ‘Garnir elements’ involved are given as linear combinations of expressions in the standard generators $\psi_1, \dots, \psi_{n-1}$

for the row permutation module which are not always reduced. Our main result concerning Specht modules is an expression for each Garnir relation as a linear combination of reduced expressions; this simplifies calculations with Specht modules, both theoretically and computationally.

But this result is a by-product of the main work in this paper, which is to consider tilings of skew Young diagrams by Dyck tiles. These tilings were introduced by Kenyon and Wilson, who defined in particular the notion of a *cover-inclusive Dyck tiling*. They used these tilings to give a formula for the inverse of a certain matrix M arising in the study of double-dimer models. We re-interpret the entries of M in terms of what we call *cover-expansive Dyck tilings*, and then, by proving recurrence relations for the numbers of cover-inclusive and cover-expansive Dyck tilings, we show that (sign-modified versions of) M and M^{-1} are in fact transition matrices for two natural bases for a certain permutation representation of the symmetric group. The fact that the two transition matrices are obviously mutually inverse provides a new proof of Kenyon and Wilson's result. Along the way we give a result showing several different equivalent conditions to the cover-inclusive condition.

In order to derive our result on Garnir relations, we then express the sum of the elements of one of our two bases in terms of the other; this involves defining a certain function F on partitions, and proving a similar recurrence to the recurrence for cover-inclusive Dyck tilings. Combining this with our results on transition coefficients means that $F(\lambda)$ equals the sum over all partitions $\mu \subseteq \lambda$ of the number of cover-inclusive Dyck tilings of $\lambda \setminus \mu$. In fact, this had already been shown by Kim, and then by Kim, Mészáros, Panova and Wilson, verifying a conjecture of Kenyon and Wilson. As well as providing a new proof of this result, our results working directly with the function F allow us to derive our application to Garnir relations without using Dyck tilings.

We now describe the structure of this paper. Section 2 is devoted to definitions. In Section 3 we study cover-inclusive Dyck tilings, giving equivalent conditions for cover-inclusiveness and then proving several bijective results which allow us to deduce recurrences for the number $i_{\lambda\mu}$ of cover-inclusive Dyck tilings of $\lambda \setminus \mu$. In Section 4 we prove similar, though considerably simpler, recurrences for cover-expansive Dyck tilings. In Section 5 we recall the Young permutation module $\mathcal{M}^{(f,g)}$ for the symmetric group; we define our two bases for this module, and use Dyck tilings to describe the transition coefficients. We then introduce the function F and use it to express the sum of the elements of the first basis in terms of the second, before summarising the relationship between our work and that of Kenyon, Kim, Mészáros, Panova and Wilson. Finally in Section 6 we give the motivating application of this work, introducing the Specht modules in the modern setting, and using our earlier results to give a new expression for the homogeneous Garnir relations.

Since the first version of this paper was written, we have become aware of the paper [SZ] by Shigechi and Zinn-Justin, where cover-inclusive and cover-expansive Dyck tilings are used in the calculation of parabolic Kazhdan–Lusztig polynomials. (They do not use the same terminology – in their paper, Dyck tiles are called “Dyck strips”, and the cover-inclusive and cover-expansive conditions are “rules I and II”. In addition – following the more widespread convention for Dyck paths – their convention for diagrams is the top-to-bottom reflection of ours.) It is likely that some of our results can be deduced from the results in [SZ].

Acknowledgements. The author is greatly indebted to David Speyer and Philippe Nadeau for comments on the online forum MathOverflow which inspired the author to introduce Dyck tilings into this paper.

2 Definitions

2.1 Partitions and Young diagrams


As usual, a *partition* is a weakly decreasing sequence $\lambda = (\lambda_1, \lambda_2, \dots)$ of non-negative integers with finite sum. We write this sum as $|\lambda|$, and say that λ is a partition of $|\lambda|$. When writing partitions, we may group equal parts together with a superscript, and omit trailing zeroes, and we write the partition $(0, 0, \dots)$ as $\emptyset$.

The *Young diagram* of a partition λ is the set

$$\{(a, b) \in \mathbb{N}^2 \mid b \leq \lambda_a\}.$$

We may abuse notation by identifying λ with its Young diagram; for example, we may write $\lambda \supseteq \mu$ to mean that $\lambda_i \geq \mu_i$ for all i . If $\lambda \supseteq \mu$, then the *skew Young diagram* $\lambda \setminus \mu$ is simply the set difference between the Young diagrams for λ and μ .

We draw (skew) Young diagrams as arrays of boxes in the plane, and except in the final section of this paper we use the *Russian convention*, where a increases from south-east to north-west, and b increases from south-west to north-east. For example, the Young diagram of $(7^2, 4, 3, 2^2) \setminus (2, 1^2)$ is as follows.



The *conjugate partition* to λ is the partition λ' obtained by reflecting the Young diagram for λ left to right; thus $\lambda'_i = |\{j \geq 1 \mid \lambda_j \geq i\}|$ for all i .

We define a *node* to be an element of $\mathbb{N}^2$, and a node of λ to be an element of the Young diagram of λ . The *height* of the node (a, b) is $a + b$. The j th *column* of $\mathbb{N}^2$ is the set of all nodes (a, b) for which $b - a = j$.

We use compass directions to label the neighbours of a node; for example, if n is a node, then we write $\text{SW}(n) = n - (0, 1)$ and refer to this as the *SW neighbour* of n ; we also write $\text{N}(n) = n + (1, 1)$, and similarly for the other compass directions.

A node n of λ is *removable* if it can be removed from λ to leave the Young diagram of a partition (i.e. if neither $\text{NW}(n)$ nor $\text{NE}(n)$ is a node of λ), while a node n not in λ is an *addable node* of λ if it can be added to λ to leave the Young diagram of a partition.

2.2 Tiles and tilings

We define a *tile* to be a finite non-empty set t of nodes that can be ordered $n_1, \dots, n_r$ such that $n_{i+1} \in \{\text{NE}(n_i), \text{SE}(n_i)\}$ for each $i = 1, \dots, r - 1$. We say that t *starts* at its leftmost node, which we denote $\text{st}(t)$, and *ends* at its rightmost, which we denote $\text{en}(t)$. The *height* $\text{ht}(t)$ of t is defined to be $\max \{\text{ht}(n) \mid n \in t\}$, and we say that t is a *Dyck tile* if this maximum is achieved at the start and end nodes of t , i.e. $\text{ht}(t) = \text{ht}(\text{st}(t)) = \text{ht}(\text{en}(t))$. t is *big* if it contains more than one node, and is a *singleton* otherwise. The *depth* of a node $n \in t$ is

$$\text{dp}(n) := \text{ht}(t) - \text{ht}(n).$$

Now suppose λ and μ are partitions with $\lambda \supseteq \mu$. A Dyck tiling of $\lambda \setminus \mu$ is a partition of $\lambda \setminus \mu$ into Dyck tiles. Given a Dyck tiling T of $\lambda \setminus \mu$ and a node $n \in \lambda \setminus \mu$, we write $\text{tile}(n)$ for the tile containing n . We say that n is *attached to* $\text{NE}(n)$ if n and $\text{NE}(n)$ lie in the same tile in T , and similarly for $\text{SE}(n)$, $\text{NW}(n)$ and $\text{SW}(n)$. If t is a tile in T , the *NE neighbour* of t is the tile starting at $\text{NE}(\text{en}(t))$, if there is one, while the *SW neighbour* of t is the tile ending at $\text{SW}(\text{st}(t))$, if there is one. SE and NW neighbours of tiles are defined similarly.


Now we recast the main definition from [KW]. Say that a Dyck tiling T of $\lambda \setminus \mu$ is *left-cover-inclusive* if whenever a and $\text{N}(a)$ are nodes of $\lambda \setminus \mu$, $\text{st}(\text{tile}(\text{N}(a)))$ lies weakly to the left of $\text{st}(\text{tile}(a))$. Similarly, T is *right-cover-inclusive* if whenever a and $\text{N}(a)$ are nodes of $\lambda \setminus \mu$, $\text{en}(\text{tile}(\text{N}(a)))$ lies weakly to the right of $\text{en}(\text{tile}(a))$. Say that T is *cover-inclusive* if it is both left- and right-cover-inclusive.

We let $\mathcal{J}(\lambda, \mu)$ denote the set of cover-inclusive Dyck tilings of $\lambda \setminus \mu$ if $\lambda \supseteq \mu$, and set $\mathcal{J}(\lambda, \mu) = \emptyset$ otherwise. Let $i_{\lambda\mu} = |\mathcal{J}(\lambda, \mu)|$.

Next we make a definition which is in some sense dual to the notion of cover-inclusiveness and which appears heavily disguised in [KW]. Say that a Dyck tiling is *left-cover-expansive* if whenever a and $\text{SE}(a)$ are nodes of $\lambda \setminus \mu$, $\text{st}(\text{tile}(\text{SE}(a)))$ lies weakly to the left of $\text{st}(\text{tile}(a))$, and *right-cover-expansive* if whenever a and $\text{SW}(a)$ are nodes of $\lambda \setminus \mu$, $\text{en}(\text{tile}(\text{SW}(a)))$ lies weakly to the right of $\text{en}(\text{tile}(a))$. A Dyck tiling is *cover-expansive* if it is both left- and right-cover-expansive.


We write $e_{\lambda\mu}$ for the number of cover-expansive Dyck tilings of $\lambda \setminus \mu$, setting $e_{\lambda\mu} = 0$ if $\lambda \not\supseteq \mu$. We shall see later that $e_{\lambda\mu} \leq 1$ for all λ, μ .

Example. We illustrate four Dyck tilings of $(6^2, 4, 3, 1^2) \setminus (4, 1^2)$. Only the first one is cover-inclusive, and only the second is cover-expansive.



We end this section with some notation which we shall use repeatedly later. Suppose $j \in \mathbb{Z}$ is fixed, and λ is a partition with an addable node l in column j . We define X_λ to be the set of all integers x such that λ has a removable node n in column $j+x$, with $\text{ht}(n) = \text{ht}(l) - 1$, and $\text{ht}(p) < \text{ht}(l)$ for all nodes $p \in \lambda$ in all columns between j and $j+x$. We set X_λ^+ to be the set of positive elements of X_λ . Note that $x \in X_\lambda^+$ precisely when there is a Dyck tile $t \subset \lambda$, starting in column $j+1$ and ending in column $j+x$, which can be removed from λ to leave a smaller partition; we denote this smaller partition $\lambda^{[x]}$. We define $\lambda^{[x]}$ for $x \in X_\lambda^- = X_\lambda \setminus X_\lambda^+$ similarly.

Example. Take $\lambda = (6, 4^2, 3, 2^2)$. Then λ has an addable node in column 0, and for this node we have $X_\lambda = \{-1, 1, 5\}$. The partitions $\lambda^{[x]}$ are as follows.



3 Cover-inclusive Dyck tilings

In this section we examine cover-inclusive Dyck tilings in more detail. We give some equivalent conditions to the cover-inclusive condition, and then we prove a recurrence for the number $i_{\lambda\mu}$ of cover-inclusive Dyck tilings.

3.1 Equivalent conditions

Theorem 3.1. *Suppose λ and μ are partitions with $\lambda \supseteq \mu$, and T is a Dyck tiling of $\lambda \setminus \mu$. The following are equivalent.*

1. T is cover-inclusive.
2. If $\mathfrak{a}$ and $N(\mathfrak{a})$ are nodes of $\lambda \setminus \mu$, then $\text{dp}(N(\mathfrak{a})) \geq \text{dp}(\mathfrak{a})$.
3. If $\mathfrak{a}$ and $N(\mathfrak{a})$ are nodes of $\lambda \setminus \mu$, then $\text{tile}(\mathfrak{a}) + (1, 1) \subseteq \text{tile}(N(\mathfrak{a}))$.
4. If $\mathfrak{a}$ and $N(\mathfrak{a})$ are nodes of $\lambda \setminus \mu$ and $NW(\mathfrak{a})$ is attached to $\mathfrak{a}$, then $NW(N(\mathfrak{a}))$ is a node of $\lambda \setminus \mu$ and is attached to $N(\mathfrak{a})$.
5. If $\mathfrak{a}$ and $N(\mathfrak{a})$ are nodes of $\lambda \setminus \mu$ and $N(\mathfrak{a})$ is the end node of its tile, then $\mathfrak{a}$ is the end node of its tile.
6. T is right-cover-inclusive.
7. If $\mathfrak{a}$ and $N(\mathfrak{a})$ are nodes of $\lambda \setminus \mu$ and $NE(\mathfrak{a})$ is attached to $\mathfrak{a}$, then $NE(N(\mathfrak{a}))$ is a node of $\lambda \setminus \mu$ and is attached to $N(\mathfrak{a})$.
8. If $\mathfrak{a}$ and $N(\mathfrak{a})$ are nodes of $\lambda \setminus \mu$ and $N(\mathfrak{a})$ is the start node of its tile, then $\mathfrak{a}$ is the start node of its tile.
9. T is left-cover-inclusive.

Proof.

(1 $\Rightarrow$ 2) Take $\mathfrak{a} \in \lambda \setminus \mu$ such that $N(\mathfrak{a}) \in \lambda \setminus \mu$, and let $\mathfrak{c} = \text{st}(\text{tile}(\mathfrak{a}))$. Since T is cover-inclusive, there is a node $\mathfrak{b}$ in $\text{tile}(N(\mathfrak{a}))$ in the same column as $\mathfrak{c}$, and

$$\text{dp}(N(\mathfrak{a})) \geq \text{ht}(\mathfrak{b}) - \text{ht}(N(\mathfrak{a})) \geq \text{ht}(\mathfrak{c}) + 2 - (\text{ht}(\mathfrak{a}) + 2) = \text{dp}(\mathfrak{a}).$$

(2 $\Rightarrow$ 3) Suppose (3) is false, and take $\mathfrak{a} \in \lambda \setminus \mu$ such that $N(\mathfrak{a}) \in \lambda \setminus \mu$ and $\text{tile}(\mathfrak{a}) + (1, 1) \not\subseteq \text{tile}(N(\mathfrak{a}))$. There is a node $\mathfrak{b} \in \text{tile}(\mathfrak{a})$ such that $\mathfrak{b} + (1, 1) \notin \text{tile}(N(\mathfrak{a}))$, and we may assume that $\mathfrak{b}$ is attached to $\mathfrak{a}$; in fact, by symmetry, we may assume $\mathfrak{b}$ is either $NE(\mathfrak{a})$ or $SE(\mathfrak{a})$. If $\mathfrak{b} = NE(\mathfrak{a})$, then $N(\mathfrak{a})$ is attached to neither its NE nor its SE neighbour, so is the end node of its tile, and in particular $\text{dp}(N(\mathfrak{a})) = 0$. On the other hand $\text{dp}(\mathfrak{a}) > \text{dp}(\mathfrak{b}) \geq 0$, contradicting (2). If instead $\mathfrak{b} = SE(\mathfrak{a})$, then $NE(\mathfrak{a})$ is attached to neither its NW nor its SW neighbour, so is the start node of its tile, and has depth 0; but $\text{dp}(\mathfrak{b}) > \text{dp}(\mathfrak{a}) \geq 0$, and again (2) is contradicted.

(3 $\Rightarrow$ 4) This is trivial.

(4 $\Rightarrow$ 5) Suppose (5) is false, and take $a \in \lambda \setminus \mu$ as far to the left as possible such that $N(a) \in \lambda \setminus \mu$ and is the end node of its tile, while a is not the end node of its tile.

Since a is not the end node of its tile, it is attached to either $SE(a)$ or $NE(a)$. If a is attached to $SE(a)$, then $SE(a)$ is attached to its NW neighbour, but $N(SE(a)) = NE(a)$ is not attached to its NW neighbour, contradicting (4).

So assume that a is attached to $NE(a)$. This implies in particular that a has positive depth.

Claim 1. If $c \in \text{tile}(a)$ and $N(c) \in \text{tile}(N(a))$, then c is not the start node of $\text{tile}(a)$.

Proof. Since $\text{tile}(N(a))$ is a Dyck tile and $N(a)$ is its end node, we have $\text{ht}(N(c)) \leq \text{ht}(N(a))$. Hence $\text{ht}(c) \leq \text{ht}(a)$, so $\text{dp}(c) \geq \text{dp}(a) > 0$.

Claim 2. For every node $b \in \text{tile}(N(a))$ we have $S(b) \in \text{tile}(a)$.

Proof. If the claim is false, let b be the rightmost counterexample. Obviously $b \neq N(a)$, and in particular b is not the end node of its tile, so b is attached to either $NE(b)$ or $SE(b)$. In the first case, the choice of b means that we have $SE(b) \in \text{tile}(a)$; neither b nor $S(b)$ is in $\text{tile}(a)$, so $SE(b)$ is the start node of $\text{tile}(a)$, contradicting Claim 1.

So we can assume b is attached to $SE(b)$. The choice of b means that $S(SE(b)) \in \text{tile}(a)$ and is not attached to $S(b)$. By Claim 1 $S(SE(b))$ is not the start node of $\text{tile}(a)$, so is attached to $S(S(b))$. But now $S(b)$ is not attached to either its NE or SE neighbour, so is the end node of its tile. $S(S(b))$ is not the end node of its tile, and this contradicts the choice of a .

Now let $b = \text{st}(\text{tile}(N(a)))$. Then by Claim 2, $c := S(b) \in \text{tile}(a)$, and by Claim 1 c is not the start node of $\text{tile}(a)$. So c is attached to either $NW(c)$ or $SW(c)$. The first possibility contradicts (4), since b is not attached to $NW(b)$, so assume that c is attached to $SW(c)$. But then $NW(c)$ is attached to neither b nor c , so is the end node of its tile, while $SW(c)$ is not the end node of its tile, and this contradicts the choice of a .

(5 $\Rightarrow$ 6) Suppose T is not right-cover-inclusive, and take $a \in \lambda \setminus \mu$ such that $N(a) \in \lambda \setminus \mu$ and $\text{tile}(a)$ ends strictly to the right of $\text{tile}(N(a))$. Let $b = \text{en}(\text{tile}(N(a)))$; then there is a node in $\text{tile}(a)$ in the same column as b , which we can write as $b - (h, h)$ for some $h > 0$. If we let $i \in \{1, \dots, h\}$ be minimal such that $b - (i, i)$ is not the end node of its tile, then $N(b - (i, i))$ is the end node of its tile, contradicting (5).

(6 $\Rightarrow$ 7) Suppose $a, N(a) \in \lambda \setminus \mu$ and a is attached to $NE(a)$. Then $\text{en}(\text{tile}(a))$ lies to the right of a , so by (6) $\text{en}(\text{tile}(N(a)))$ does too. So $N(a)$ is attached to either $NE(N(a))$ or $SE(N(a))$. But $SE(N(a)) = NE(a)$ is attached to a , so $N(a)$ is attached to $NE(N(a))$.

(7 $\Rightarrow$ 8) This is symmetrical to the argument that 4 $\Rightarrow$ 5.

(8 $\Rightarrow$ 9) This is symmetrical to the argument that 5 $\Rightarrow$ 6.

(9 $\Rightarrow$ 1) Since 6 $\Rightarrow$ 7 $\Rightarrow$ 8 $\Rightarrow$ 9, right-cover-inclusive implies left-cover-inclusive. Symmetrically, left-cover-inclusive implies right-cover-inclusive, and hence cover-inclusive. $\square$

We shall use these equivalent definitions of cover-inclusiveness, often without comment, in what follows. We also observe the following property of cover-inclusive Dyck tilings, which we shall use without comment.

Lemma 3.2. *Suppose T is a cover-inclusive Dyck tiling of $\lambda \setminus \mu$. If $n \in \lambda \setminus \mu$ is the highest node in its column in $\lambda \setminus \mu$, then every node in $\text{tile}(n)$ is the highest node in its column in $\lambda \setminus \mu$.*

Proof. Suppose not, and take $m \in \text{tile}(n)$ which is not the highest node in its column. Without loss of generality we may assume m lies in the column immediately to the right of n , i.e. m is either $\text{NE}(n)$ or $\text{SE}(n)$. But if $m = \text{NE}(n)$, then the assumption $N(m) \in \lambda \setminus \mu$ means that $N(n) \in \lambda \setminus \mu$ (otherwise $\lambda \setminus \mu$ would not be a skew Young diagram), a contradiction. So assume $m = \text{SE}(n)$. But now m is attached to $\text{NW}(m)$ and $N(m) \in \lambda \setminus \mu$, so by Theorem 3.1(4) $N(n) \in \lambda \setminus \mu$, contradiction. $\square$

3.2 Recurrences

Now we consider recurrences. We start with a simple result.

Proposition 3.3. *Suppose λ and μ are partitions and $j \in \mathbb{Z}$, and that μ has an addable node m in column j , but λ does not have an addable node in column j . Let $\mu^+ = \mu \cup \{m\}$. Then $i_{\lambda\mu} = i_{\lambda\mu^+}$.*

Proof. The fact that λ does not have an addable node in column j implies that $\lambda \supseteq \mu$ if and only if $\lambda \supseteq \mu^+$, so we may as well assume that both of these conditions hold. Given a cover-inclusive Dyck tiling of $\lambda \setminus \mu^+$, we can obtain a tiling of $\lambda \setminus \mu$ simply by adding the singleton tile $\{m\}$, and it is clear that this tiling is a cover-inclusive Dyck tiling.

In the other direction, suppose T is a cover-inclusive Dyck tiling of $\lambda \setminus \mu$; then we claim that m forms a singleton tile. If we let n denote the highest node in column j of $\lambda \setminus \mu$, then (since λ does not have an addable node in column j) $\text{NE}(n)$ and $\text{NW}(n)$ are not both nodes of $\lambda \setminus \mu$; suppose without loss that $\text{NE}(n) \notin \lambda \setminus \mu$. Then in particular n is not attached to $\text{NE}(n)$ in T , and so (using Theorem 3.1(7)) m is not attached to $\text{NE}(m)$. m cannot be attached to $\text{SE}(m)$ or $\text{SW}(m)$ (since these are not nodes of $\lambda \setminus \mu$), and a node in a Dyck tile cannot be attached only to its NW neighbour. So m is not attached to any of its neighbours, i.e. $\{m\}$ is a singleton tile as claimed. We can remove this tile, and we clearly obtain a cover-inclusive Dyck tiling of $\lambda \setminus \mu^+$. So we have a bijection between $\mathcal{J}(\lambda, \mu)$ and $\mathcal{J}(\lambda, \mu^+)$. $\square$

Now we prove a more complicated result, for which it will help us to fix some notation.

Notation in force for the remainder of Section 3:

j is a fixed integer, and λ, μ are partitions such that λ has an addable node l in column j , while μ has an addable node m in column j . We define $\lambda^+ = \lambda \cup \{l\}$ and $\mu^+ = \mu \cup \{m\}$.

If $T \in \mathcal{J}(\lambda, \mu^+)$; then $\overrightarrow{T}$ denotes the size of the highest tile starting in column $j+1$, and $\overleftarrow{T}$ denotes the size of the highest tile ending in column $j-1$.

Note that if $T \in \mathcal{J}(\lambda, \mu^+)$, then there must be at least one tile starting in column $j+1$, since there are more nodes in column $j+1$ of $\lambda \setminus \mu^+$ than in column j ; so $\overrightarrow{T}$ is well-defined, and similarly $\overleftarrow{T}$ is well-defined.

To prove our main recurrence result for the numbers $i_{\lambda\mu}$, we give three results in which we construct bijections between sets of cover-inclusive Dyck tilings. The first of these is as follows; recall the definition of X_λ from Section 2.2.


Proposition 3.4.

$$\left| \left\{ T \in \mathcal{J}(\lambda, \mu^+) \mid \overrightarrow{T} \notin X_\lambda \right\} \right| = \left| \left\{ T \in \mathcal{J}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is a big tile in } T \\ \text{starting in column } j \end{array} \right\} \right|.$$

Proof. Note first that if $\lambda^+ \not\supseteq \mu^+$ then $\lambda \not\supseteq \mu^+$, so both sides equal zero. Conversely, if $\lambda \not\supseteq \mu^+$, then either $\lambda^+ \not\supseteq \mu^+$ or there are no nodes in column j of $\lambda^+ \setminus \mu^+$, so again both sides are zero. So we may assume that $\lambda \supseteq \mu^+$. We'll construct a bijection

$$\phi_1 : \left\{ T \in \mathcal{T}(\lambda, \mu^+) \mid \vec{T} \notin X_\lambda \right\} \longrightarrow \left\{ T \in \mathcal{T}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is a big tile in } T \\ \text{starting in column } j \end{array} \right\}.$$

We begin by giving an example so that the reader can see how the bijection works; in this example, the tiling on the left lies in $\left\{ T \in \mathcal{T}(\lambda, \mu^+) \mid \vec{T} \notin X_\lambda \right\}$, and the tiling on the right is its image under ϕ_1 . The nodes l and m are marked, the nodes in column j are shaded and the tile t (introduced below) is marked with dots.



Given T in $\mathcal{T}(\lambda, \mu^+)$ with $\vec{T} \notin X_\lambda$, let t be the highest tile in T starting in column $j+1$, and let A be the set of nodes in column j which are higher than $\text{st}(t)$. No node in column $j+1$ can be attached to its NW neighbour, because the highest node in this column (namely $\text{SE}(l)$) is not. So (by the choice of t) every node b higher than $\text{st}(t)$ in column $j+1$ is attached to $\text{SW}(b)$, and hence each $a \in A$ is attached to $\text{NE}(a)$. In a Dyck tiling a node cannot be attached only to its NE neighbour, so each $a \in A$ is attached to either its NW or SW neighbour. But a node in column $j-1$ cannot be attached to its NE neighbour (because the highest node in column $j-1$ is not), so every $a \in A$ is attached to $\text{NW}(a)$.

Now we consider columns x and $x+1$, where $x = |\vec{T}| + j = |t| + j$. Let $d = \text{SE}(l)$. Then $d = \text{st}(t) + (h, h)$ for some $h \geq 0$, so by Theorem 3.1(3) $\text{tile}(d)$ contains $t + (h, h)$. In particular, $\text{tile}(d)$ includes the node $e = \text{en}(t) + (h, h)$ in column x , and (if $\text{st}(t) \neq \text{en}(t)$) includes $\text{SW}(e)$. Since d is the highest node in its column, the same is true for every node in $\text{tile}(d)$, and in particular neither $\text{N}(e)$ nor $\text{NW}(e)$ is a node of λ . However, e cannot be a removable node of λ , since then $\vec{T} = |t|$ would lie in X_λ . So $\text{NE}(e)$ is a node of $\lambda \setminus \mu$. Now $\text{NE}(e)$ is not attached to its NW neighbour (since this is $\text{N}(e) \notin \lambda$), so no node in column $x+1$ is attached to its NW neighbour, and so no node in column x is attached to its SE neighbour. This implies that any node in column x of positive depth must be attached to its NE neighbour, while any node in column x of depth 0 is the end node of its tile, and this tile has a NE neighbour. In particular, t has a NE neighbour.

Now suppose $a \in A$ has depth 1, and write $a = \text{st}(t) + (i, i-1)$ for $i > 0$. Then $\text{tile}(a)$ contains $\text{st}(t) + (i, i)$, and hence contains $\text{en}(t) + (i, i)$. Since a has depth 1, $\text{en}(t) + (i, i)$ has depth 0, and so from the last paragraph $\text{en}(t) + (i, i)$ is the end node of $\text{tile}(a)$, and $\text{tile}(a)$ has a NE neighbour.

Now we can construct $\phi_1(T)$ as follows.

- For each $a \in A$ of depth at least 2, change $\text{tile}(a)$ by replacing a with $\text{N}(a)$.

- For each $a \in A$ of depth 1, let $u = \text{tile}(a)$. Then u has a NE neighbour v ; replace u and v with the following tiles:
 - the tile consisting of the portion of u ending at $\text{NW}(a)$;
 - the tile obtained by combining $\text{N}(a)$, the portion of u starting at $\text{NE}(a)$, and v .
- Finally let w denote the NE neighbour of t , and replace t and w with the tile obtained by joining $\text{NW}(\text{st}(t))$ to t and w .

Clearly all these new tiles are Dyck tiles. It's also clear that the last tile mentioned is a big tile starting in column j . So to see that $\phi_1(T)$ lies in the codomain, all that remains is to check that $\phi_1(T)$ is cover-inclusive. But this is easy using (for example) Theorem 3.1(4).

To show that ϕ_1 is a bijection, we construct the inverse map

$$\psi_1 : \left\{ T \in \mathcal{J}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is a big tile in } T \\ \text{starting in column } j \end{array} \right\} \longrightarrow \left\{ T \in \mathcal{J}(\lambda, \mu^+) \mid \vec{T} \notin X_\lambda \right\}.$$

Suppose T is a cover-inclusive Dyck tiling of $\lambda^+ \setminus \mu^+$ in which there is at least one big tile starting in column j . Since the highest node in column j (namely $\mathfrak{l}$) is not attached to its NE or NW neighbour, no node in column j is attached to its NE or NW neighbour. So from top to bottom, column j of T consists of:

- (possibly) some nodes attached to both their SE and SW neighbours;
- at least one node attached only to its SE neighbour;
- (possibly) some singleton tiles.

Construct $\psi_1(T)$ as follows.

- For each node a in column j which is attached to its SE and SW neighbours, change $\text{tile}(a)$ by replacing a with $\text{S}(a)$.
- For each node a in column j which is the start of a big tile, let b be the first node of $\text{tile}(a)$ to the right of column j which has the same height as a . Replace $\text{tile}(a)$ and $\text{tile}(\text{SW}(a))$ with the following tiles:
 - the tile consisting of the portion of $\text{tile}(a)$ starting at b ;
 - (if $\text{tile}(a)$ is the lowest big tile of T starting in column j) the tile consisting of the portion of $\text{tile}(a)$ running from $\text{SE}(a)$ to $\text{SW}(b)$;
 - (if $\text{tile}(a)$ is not the lowest big tile of T starting in column j) the tile obtained by joining together $\text{tile}(\text{SW}(a))$, the node $\text{S}(a)$ and the portion of $\text{tile}(a)$ running from $\text{SE}(a)$ to $\text{SW}(b)$.

Again, it is easy to check that the new tiles are all Dyck tiles, and checking that $\psi_1(T)$ is cover-inclusive is straightforward using Theorem 3.1(4).

So $\psi_1(T)$ is a cover-inclusive Dyck tiling. It remains to check that $\vec{\psi_1(T)} \notin X_\lambda$. Let t be the highest tile in $\psi_1(T)$ starting in column $j+1$; then t starts at $\text{SE}(a)$, where a is the lowest node in column j which is the start of a big tile in T , and ends at $\text{SW}(b)$, where b is the first node in $\text{tile}(a)$

to the right of column j with the same height as a . In T , $\text{SW}(b)$ is attached to b , and so every node above $\text{SW}(b)$ in the same column is attached to its NE neighbour in T . In particular, the highest node in this column has a NE neighbour in λ , and so is not removable; so $\vec{\psi}_1(\vec{T}) \notin X_\lambda$.

So our two maps ϕ_1, ψ_1 really do map between the specified sets. It is easy to see from the construction that they are mutual inverses. $\square$

Symmetrically, we have the following result.

Proposition 3.5.

$$\left| \left\{ T \in \mathcal{J}(\lambda, \mu^+) \mid -\overleftarrow{T} \notin X_\lambda \right\} \right| = \left| \left\{ T \in \mathcal{J}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is a big tile in } T \\ \text{ending in column } j \end{array} \right\} \right|.$$

Our next bijective result is the following.


Proposition 3.6. For each $x \in X_\lambda^+$,

$$\left| \left\{ T \in \mathcal{J}(\lambda, \mu^+) \mid \vec{T} = x \right\} \right| = i_{\lambda^{[x]}\mu^+}.$$

Proof. Fix $x \in X_\lambda^+$. Our aim is to define a bijection

$$\phi_2 : \left\{ T \in \mathcal{J}(\lambda, \mu^+) \mid \vec{T} = x \right\} \longrightarrow \mathcal{J}(\lambda^{[x]}, \mu^+).$$

Again we begin with an example to illustrate the bijection. In this example $x = 5$, and the nodes in columns j and $j + x + 1$ are shaded, and the tiles ρ and t (introduced below) are marked with dots.



Since $x \in X_\lambda^+$, there is a Dyck tile $\rho \subset \lambda$ consisting of the highest nodes in columns $j + 1, \dots, j + x$ of λ . Take $T \in \mathcal{J}(\lambda, \mu^+)$ with $\vec{T} = x$, and let t be the highest tile in T starting in column $j + 1$. Then every northward translate of t is an interval in a tile; in particular, $\text{tile}(\text{SE}(t))$ contains a translate $t + (h, h)$ for some $h \geq 0$; since $\text{SE}(t)$ is the highest node in its column, the same is true for every node in $\text{tile}(\text{SE}(t))$, so the translate $t + (h, h)$ coincides with ρ .

Consider the nodes in column j of $\lambda \setminus \mu^+$. Arguing as in the proof of Proposition 3.4, every node a in column j which is higher than t is attached to $\text{NE}(a)$ and $\text{NW}(a)$.

Next we consider nodes in column $j + x$. The highest node in column $j + x$ of $\lambda \setminus \mu^+$ is the end node of ρ , and so is not attached to its NE neighbour (since this node is not in λ); hence no node in column $j + x$ is attached to its NE neighbour. So no node in column $j + x + 1$ is attached to its SW neighbour. So from top to bottom, the nodes in column $j + x + 1$ higher than $\text{en}(t)$ comprise:

- (possibly) some nodes attached to their NW neighbours;
- (possibly) some nodes attached to neither their NW nor SW neighbours.

Now we can construct $\phi_2(T)$ by doing the following, for each pair (a, d) of nodes with a in column j , d in column $j + x + 1$ and $\text{ht}(a) = \text{ht}(d) > \text{ht}(t)$.

- If d is attached to $\text{NW}(d)$, then a and d lie in the same tile in T ; change this tile by moving the portion lying in columns $j + 1$ to $j + x$ south one step.
- If d is not attached to $\text{NW}(d)$, then $\text{en}(\text{tile}(a)) = \text{NW}(d)$ and hence $\text{NW}(a)$ has depth 0. Replace $\text{tile}(a)$ and $\text{tile}(d)$ with the following two tiles:
 - the portion of $\text{tile}(a)$ ending at $\text{NW}(a)$;
 - the tile obtained by joining the portion of $\text{tile}(\text{SE}(a))$ between columns $j + 1$ and $j + x$ to a and $\text{tile}(d)$.

Finally remove t . It is easy to see that the new tiles are Dyck tiles, and the cover-inclusive property follows very easily from the cover-inclusive property for T .

Now we construct the inverse map

$$\psi_2 : \mathcal{J}(\lambda^{[x]}, \mu^+) \longrightarrow \{ T \in \mathcal{J}(\lambda, \mu^+) \mid \vec{T} = x \}.$$

Suppose T is a cover-inclusive Dyck tiling of $\lambda^{[x]} \setminus \mu^+$. No node in column $j - 1$ or column j of $\lambda^{[x]}$ can be attached to its NE neighbour (since the NE neighbours of the highest nodes in these columns are not nodes of $\lambda^{[x]}$). Hence from top to bottom the nodes in column j comprise:

- (possibly) some nodes attached to their NW and SE neighbours;
- (possibly) some nodes attached only to their SE neighbours;
- (possibly) some singleton nodes.

Note that if a is a node in column j which is not attached to $\text{NW}(a)$, then $\text{NW}(a)$ is the end node of its tile; in particular, $\text{tile}(a)$ has a NW neighbour.

Claim. Suppose $r \geq 1$, and the node $a := l - (r, r)$ is attached to $\text{SE}(a)$. Then $\text{tile}(a)$ reaches column $j + x + 1$ and includes all the nodes in $\rho - (r, r)$.

Proof. We proceed by induction on r . First suppose $r = 1$. Since $x \in X_\lambda$, every node in columns $j + 1, \dots, j + x$ of λ has height less than $\text{ht}(l)$. Hence every node in columns $j + 1, \dots, j + x$ of $\lambda^{[x]}$ has height less than $\text{ht}(l) - 2 = \text{ht}(a)$; since $\text{tile}(a)$ is a Dyck tile, it must reach $\text{ht}(a)$ at some point to the right of $\text{SE}(a)$, and so must reach column $x + j + 1$. Furthermore, a is the highest node in its column, so every node in $\text{tile}(a)$ is the highest in its column, and in particular the portion of $\text{tile}(a)$ between columns $j + 1$ and $j + x$ must consist of the highest nodes in these columns, i.e. the nodes in $\rho - (1, 1)$.

Now suppose $r > 1$. By induction $\text{tile}(a)$ cannot include any of the nodes in $\rho - (v, v)$ for any $v < r$; all the remaining nodes in columns $j + 1, \dots, j + r$ have height less than $\text{ht}(a)$, and so (since $\text{tile}(a)$ must reach $\text{ht}(a)$ at some point to the right of $\text{SE}(a)$) $\text{tile}(a)$ must reach column $j + x + 1$. The translate $\text{tile}(a) + (1, 1)$ is contained in $\text{tile}(\text{N}(a))$, and by induction this includes the nodes in $\rho - (r - 1, r - 1)$; so $\text{tile}(a)$ includes the nodes in $\rho - (r, r)$.

Now we consider the nodes in column $j + x + 1$. If $\mathfrak{d}$ is a node in column $j + x + 1$, then by similar arguments to those used above, $\mathfrak{d}$ cannot be attached to $\text{NW}(\mathfrak{d})$, and if $\mathfrak{d}$ is attached to $\text{SW}(\mathfrak{d})$, then $\text{tile}(\mathfrak{d})$ includes a node in column j and also includes all the nodes in $\rho - (r, r)$, where $r = \frac{1}{2}(\text{ht}(\mathfrak{l}) - \text{ht}(\mathfrak{d}))$.

So we find that if $\mathfrak{a}$ and $\mathfrak{d}$ are nodes in columns j and $j + x + 1$ respectively with the same height, then $\mathfrak{a}$ is attached to $\text{SE}(\mathfrak{a})$ if and only if $\mathfrak{d}$ is attached to $\text{SW}(\mathfrak{d})$, and that in this case $\text{tile}(\mathfrak{a}) = \text{tile}(\mathfrak{d})$. Say that such a pair $(\mathfrak{a}, \mathfrak{d})$ is a *connected* pair. Now we can construct $\psi_2(T)$ as follows: for each connected pair $(\mathfrak{a}, \mathfrak{d})$, write $\mathfrak{a} = \mathfrak{l} - (r, r)$ and then:

- if $\mathfrak{a}$ is attached to $\text{NW}(\mathfrak{a})$, change $\text{tile}(\mathfrak{a})$ by replacing the portion $\rho - (r, r)$ with $\rho - (r - 1, r - 1)$;
- if $\mathfrak{a}$ is not attached to $\text{NW}(\mathfrak{a})$, replace $\text{tile}(\text{NW}(\mathfrak{a}))$ and $\text{tile}(\mathfrak{a})$ with the following two tiles:
 - the tile comprising $\mathfrak{a}$, $\text{tile}(\text{NW}(\mathfrak{a}))$ and $\rho - (r - 1, r - 1)$;
 - the portion of $\text{tile}(\mathfrak{a})$ starting at $\mathfrak{d}$.

Finally, add the tile $\rho - (r, r)$, where r is the number of connected pairs.

Once more, we see that the new tiles are all Dyck tiles, and the cover-inclusive property is easy to check from the construction. Furthermore, the last tile mentioned above is the highest tile starting in column $j + 1$, and has size x .

So ψ_2 really does map $\mathcal{J}(\lambda^{[x]}, \mu^+)$ to $\{T \in \mathcal{J}(\lambda, \mu^+) \mid \overrightarrow{T} = x\}$. And it is easy to see that ϕ_2 and ψ_2 are mutual inverses. $\square$

Symmetrically, we have the following.

Proposition 3.7. *For each $x \in X_\lambda^-$,*

$$|\{T \in \mathcal{J}(\lambda, \mu^+) \mid \overleftarrow{T} = -x\}| = i_{\lambda^{[x]}\mu^+}.$$

Our third bijection involves the partition λ^+ .

Proposition 3.8.


$$\left| \left\{ T \in \mathcal{J}(\lambda, \mu) \mid \begin{array}{l} \mathfrak{m} \text{ does not lie in a} \\ \text{singleton tile of } T \end{array} \right\} \right| = \left| \left\{ T \in \mathcal{J}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is no big tile in } T \text{ start-} \\ \text{ing or ending in column } j \end{array} \right\} \right|.$$

Proof. First observe that $\lambda \supseteq \mu$ if and only if $\lambda^+ \supseteq \mu^+$, so we may as well assume that both these conditions hold. Next note that if $\mathfrak{l} = \mathfrak{m}$, then $\lambda \setminus \mu = \lambda^+ \setminus \mu^+$, and this skew Young diagram contains no nodes in column j ; in particular, it does not include $\mathfrak{m}$. So in this case, Proposition 3.8 amounts to the trivial statement $i_{\lambda\mu} = i_{\lambda^+\mu^+}$.

So we assume that $\lambda \supset \mu$ and $\mathfrak{l} \neq \mathfrak{m}$; so $\mathfrak{m}$ is a node of $\lambda \setminus \mu$. Now we want to construct a bijection

$$\phi_3 : \left\{ T \in \mathcal{J}(\lambda, \mu) \mid \begin{array}{l} \mathfrak{m} \text{ lies in a} \\ \text{big tile in } T \end{array} \right\} \longrightarrow \left\{ T \in \mathcal{J}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is no big tile in } T \text{ start-} \\ \text{ing or ending in column } j \end{array} \right\}.$$

The following diagrams gives an illustrative example of how our bijection works; the nodes in the same column as $\mathfrak{l}$ and $\mathfrak{m}$ are shaded to make it easier to see the effect on these nodes.



Suppose we have $T \in \mathcal{J}(\lambda, \mu)$ with m lying in a big tile. Then m is attached to both $NE(m)$ and $NW(m)$, and hence every node a in column j is attached to both $NE(a)$ and $NW(a)$. In particular, every node in column j has depth at least 1. Now we construct $\phi_3(T)$ from T as follows. For each node a in column j :

- if a has depth greater than 1, then change $\text{tile}(a)$ by replacing a with $N(a)$;
- if a has depth 1, then replace $\text{tile}(a)$ with the following three tiles:
 - the portion of $\text{tile}(a)$ ending in column $j - 1$;
 - the portion of $\text{tile}(a)$ starting in column $j + 1$;
 - the singleton tile $\{N(a)\}$.

Clearly all these new tiles are Dyck tiles. The cover-inclusive property follows easily from the corresponding property of T , and in particular the fact (Theorem 3.1(2)) that depth weakly decreases down columns in a cover-inclusive tiling. Moreover, $\phi_3(T)$ has no big tile starting or ending in column j , so lies in the codomain.

To show that ϕ_3 is a bijection, we construct its inverse

$$\psi_3 : \left\{ T \in \mathcal{J}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is no big tile in } T \text{ start-} \\ \text{ing or ending in column } j \end{array} \right\} \longrightarrow \{ T \in \mathcal{J}(\lambda, \mu) \mid m \text{ lies in a big tile in } T \}.$$

Suppose $T \in \mathcal{J}(\lambda^+, \mu^+)$, and that there is no big tile in T starting or ending in column j . The highest node in column j of $\lambda^+ \setminus \mu^+$, namely l , is not attached to $NE(l)$ or $NW(l)$ (since these are not nodes of λ^+) and so no node in column j is attached to its NE or NW neighbour. So every node in column j is either a singleton or attached to both its SE and SW neighbours. Now we can construct $\psi_3(T)$ from T as follows: for each node a in column j , replace the tiles $\text{tile}(a)$, $\text{tile}(SW(a))$, $\text{tile}(SE(a))$ (which might or might not coincide) with the tile obtained by taking the union of these three tiles and replacing a with $S(a)$.

Again, it is easy to check that $\psi_3(T)$ is a cover-inclusive Dyck tiling, and clearly m lies in a big tile in $\psi_3(T)$. Furthermore, it is easy to see that ϕ_3 and ψ_3 are mutually inverse. $\square$

Now we combine Propositions 3.4–3.8 to prove our main recurrence. We retain the assumptions and notation from above.

Our main result is as follows. Recall that $i_{\lambda\mu}$ denotes the total number of cover-inclusive Dyck tilings of $\lambda \setminus \mu$.

Proposition 3.9. *With notation as above,*

$$i_{\lambda\mu} + i_{\lambda\mu^+} = i_{\lambda^+\mu^+} + \sum_{x \in X_\lambda} i_{\lambda^{[x]}\mu^+}.$$

Proof. By Proposition 3.6, we have

$$|\{T \in \mathcal{T}(\lambda, \mu^+) \mid \vec{T} \in X_\lambda\}| = \sum_{x \in X_\lambda^+} i_{\lambda[x]\mu^+}.$$

So

$$\begin{aligned} i_{\lambda\mu^+} &= |\{T \in \mathcal{T}(\lambda, \mu^+) \mid \vec{T} \in X_\lambda\}| + |\{T \in \mathcal{T}(\lambda, \mu^+) \mid \vec{T} \notin X_\lambda\}| \\ &= \sum_{x \in X_\lambda^+} i_{\lambda[x]\mu^+} + \left| \left\{ T \in \mathcal{T}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is a big tile in } T \\ \text{starting in column } j \end{array} \right\} \right| \end{aligned}$$

by Proposition 3.4.

Symmetrically, we have

$$i_{\lambda\mu^+} = \sum_{x \in X_\lambda^-} i_{\lambda[x]\mu^+} + \left| \left\{ T \in \mathcal{T}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is a big tile in } T \\ \text{ending in column } j \end{array} \right\} \right|.$$

Now consider $i_{\lambda\mu}$. Obviously we have

$$|\{T \in \mathcal{T}(\lambda, \mu) \mid \mathbf{m} \text{ lies in a singleton tile in } T\}| = i_{\lambda\mu^+},$$

and by Proposition 3.8

$$\left| \left\{ T \in \mathcal{T}(\lambda, \mu) \mid \begin{array}{l} \mathbf{m} \text{ does not lie in a} \\ \text{singleton tile of } T \end{array} \right\} \right| = \left| \left\{ T \in \mathcal{T}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is no big tile in } T \text{ start-} \\ \text{ing or ending in column } j \end{array} \right\} \right|.$$

Hence

$$\begin{aligned} i_{\lambda\mu} + i_{\lambda\mu^+} &= |\{T \in \mathcal{T}(\lambda, \mu) \mid \mathbf{m} \text{ lies in a singleton tile in } T\}| + \left| \left\{ T \in \mathcal{T}(\lambda, \mu) \mid \begin{array}{l} \mathbf{m} \text{ does not lie in a} \\ \text{singleton tile of } T \end{array} \right\} \right| + i_{\lambda\mu^+} \\ &= \left| \left\{ T \in \mathcal{T}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is no big tile in } T \text{ start-} \\ \text{ing or ending in column } j \end{array} \right\} \right| + 2i_{\lambda\mu^+} \\ &= \sum_{x \in X_\lambda} i_{\lambda[x]\mu^+} + \left| \left\{ T \in \mathcal{T}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is a big tile in } T \\ \text{starting in column } j \end{array} \right\} \right| \\ &\quad + \left| \left\{ T \in \mathcal{T}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is a big tile in } T \\ \text{ending in column } j \end{array} \right\} \right| \\ &\quad + \left| \left\{ T \in \mathcal{T}(\lambda^+, \mu^+) \mid \begin{array}{l} \text{there is no big tile in } T \text{ start-} \\ \text{ing or ending in column } j \end{array} \right\} \right|. \end{aligned}$$

Since a cover-inclusive Dyck tiling cannot contain big tiles starting and ending in the same column, the sum of the last three terms is $i_{\lambda^+\mu^+}$, and we are done. $\square$

4 Cover-expansive Dyck tilings

In this section we consider cover-expansive Dyck tilings, proving similar (though considerably simpler) recurrences to those in Section 3.

4.1 Basic properties

We begin by studying left- and right-cover-expansive Dyck tilings. In contrast to cover-inclusive Dyck tilings, it is not the case that the left-cover-expansive and right-cover-expansive conditions are equivalent. For example, the unique Dyck tiling of $(2)\backslash\emptyset$ is left- but not right-cover-expansive.

We begin with equivalent conditions to the left- and right-cover-expansive conditions.

Proposition 4.1. *Suppose λ and μ are partitions with $\lambda \supseteq \mu$, and T is a Dyck tiling of $\lambda \backslash \mu$.*

1. *T is left-cover-expansive if and only if for every tile t in T , we have $\text{NW}(\text{st}(t)) \notin \lambda \backslash \mu$.*
2. *T is right-cover-expansive if and only if for every tile t in T , we have $\text{NE}(\text{en}(t)) \notin \lambda \backslash \mu$.*

Proof. We prove (1); the proof of (2) is similar.

Suppose T is left-cover-expansive. Given a tile t , let $\alpha = \text{NW}(\text{st}(t))$. If $\alpha \in \lambda \backslash \mu$, then by the left-cover-expansive property $\text{st}(\text{tile}(\alpha))$ lies weakly to the right of $\text{st}(t)$; but α lies strictly to the left of $\text{st}(t)$, a contradiction.

Conversely, suppose the given property holds, and $\alpha, \text{SE}(\alpha)$ are nodes of $\lambda \backslash \mu$. Let $\beta = \text{st}(\text{tile}(\text{SE}(\alpha)))$, and suppose β lies in column i . Then $\text{NW}(\beta) \notin \lambda \backslash \mu$, and hence there are no nodes in column $i - 1$ of $\lambda \backslash \mu$ higher than β (otherwise $\lambda \backslash \mu$ would not be a skew Young diagram). In particular, there are no nodes of $\text{tile}(\alpha)$ in column $i - 1$, and hence $\text{st}(\text{tile}(\alpha))$ lies weakly to the right of β . $\square$

Lemma 4.2. *Suppose λ, μ are partitions with $\lambda \supseteq \mu$, and T is a left-cover-expansive Dyck tiling of $\lambda \backslash \mu$. If $\alpha \in \lambda \backslash \mu$ is the lowest node in its column, then every node in $\text{tile}(\alpha)$ to the right of α is the lowest node in its column.*

Proof. Suppose α lies in column i , and proceed by induction on the number of nodes to the right of α in $\text{tile}(\alpha)$. Assuming α is not the end node of its tile, there is a node β in $\text{tile}(\alpha)$ in column $i + 1$, which must be either $\text{NE}(\alpha)$ or $\text{SE}(\alpha)$. The only way β can fail to be the lowest node in column $i + 1$ is if $\beta = \text{NE}(\alpha)$ and $\text{SE}(\alpha) \in \lambda \backslash \mu$. But in this case, $\text{SE}(\alpha)$ is not attached to α or to $\text{S}(\alpha)$ (which is not a node of $\lambda \backslash \mu$), and so $\text{SE}(\alpha)$ is the start of its tile; but this contradicts Proposition 4.1.

So β is the lowest node in its column. By induction every node to the right of β in the same tile is the lowest node in its column, and we are done. $\square$

Proposition 4.3. *Suppose λ, μ are partitions with $\lambda \supseteq \mu$. Then $\lambda \backslash \mu$ admits at most one left-cover-expansive Dyck tiling, and at most one right-cover-expansive Dyck tiling. If $\lambda \backslash \mu$ admits both a left- and a right-cover-expansive Dyck tiling, then these tilings coincide.*

Proof. We use induction on $|\lambda \backslash \mu|$. If $\lambda = \mu$ then the result is trivial, so assume $\lambda \supset \mu$. Let α be the unique leftmost node of $\lambda \backslash \mu$, and suppose α lies in column i . Let $m \geq 0$ be maximal such that columns $i, \dots, i + m$ each contain a node of height at most $\text{ht}(\alpha)$. Suppose T is a left-cover-expansive Dyck tiling of $\lambda \backslash \mu$.

Claim 1. $\text{tile}(\alpha)$ consists of the lowest nodes in columns $i, \dots, i + m$.

Proof. Since α is the start node of its tile, every node in $\text{tile}(\alpha)$ has height at most $\text{ht}(\alpha)$, and in particular $\text{tile}(\alpha)$ cannot contain a node in column $i + m + 1$ or further to the right.

Now we prove by induction on l that $\text{tile}(\mathfrak{a})$ contains the lowest node in column l , for $i \leq l \leq i + m$. Let $\mathfrak{b}$ be the lowest node in column l ; then (assuming $l > i$) the lowest node in column $l - 1$ is either $\text{NW}(\mathfrak{b})$ or $\text{SW}(\mathfrak{b})$. In the first case, $\mathfrak{b}$ cannot be the start of its tile, by Proposition 4.1; $\mathfrak{b}$ cannot be attached to $\text{SW}(\mathfrak{b})$, since this is not a node of $\lambda \setminus \mu$, and so $\mathfrak{b}$ is attached to $\text{NW}(\mathfrak{b})$, and hence lies in $\text{tile}(\mathfrak{a})$. In the second case, the fact that $\text{ht}(\text{SW}(\mathfrak{b})) < \text{ht}(\mathfrak{b}) \leq \text{ht}(\mathfrak{a})$ means that $\text{SW}(\mathfrak{b})$ cannot be the end node of $\text{tile}(\mathfrak{a})$, so is attached to either $\mathfrak{b}$ or $\text{S}(\mathfrak{b})$; but $\text{S}(\mathfrak{b})$ is not a node of $\lambda \setminus \mu$, and hence $\text{SW}(\mathfrak{b})$ is attached to $\mathfrak{b}$, i.e. $\mathfrak{b} \in \text{tile}(\mathfrak{a})$.

The definition of m means that the nodes in $\text{tile}(\mathfrak{a})$ can be removed to leave a smaller skew Young diagram $\lambda \setminus \nu$, and the fact that T is left-cover-expansive means that $T \setminus \{\text{tile}(\mathfrak{a})\}$ is a left-cover-expansive Dyck tiling of $\lambda \setminus \nu$. By induction on $|\lambda \setminus \mu|$ there is at most one such tiling, and so T is uniquely determined.

So there is at most one left-cover-expansive Dyck tiling of $\lambda \setminus \mu$, and similarly at most one right-cover-expansive Dyck tiling. To prove the final statement, we continue to assume that $\lambda \setminus \mu$ is non-empty; we choose a connected component $\mathcal{C}$ of $\lambda \setminus \mu$, and let $\mathfrak{a}$ denote the unique leftmost node of $\mathcal{C}$, and $\mathfrak{c}$ the unique rightmost node of $\mathcal{C}$. (So there is no node in the column to the left of the column containing $\mathfrak{a}$ or in the column to the right of the column containing $\mathfrak{c}$, but there are nodes in all columns in between.)

Claim 2. Suppose there exists a left-cover-expansive Dyck tiling T of $\lambda \setminus \mu$. Then $\text{ht}(\mathfrak{a}) \leq \text{ht}(\mathfrak{c})$, and if equality occurs then $\mathfrak{a}$ and $\mathfrak{c}$ lie in the same tile in T .

Proof. Let $\mathfrak{b}_1, \dots, \mathfrak{b}_r$ be the nodes in $\mathcal{C}$ which are both the end nodes of their tiles and the lowest nodes in their columns, numbering them so that they appear in order from left to right. Then $\mathfrak{b}_1 = \text{en}(\text{tile}(\mathfrak{a}))$ by Claim 1, and $\mathfrak{b}_r = \mathfrak{c}$. We claim that $\text{ht}(\mathfrak{b}_1) < \dots < \text{ht}(\mathfrak{b}_r)$. Given $1 \leq l < r$, let $\mathfrak{d}$ be the lowest node in the column to the right of $\mathfrak{b}_l$. Then $\mathfrak{d}$ is either $\text{NE}(\mathfrak{b}_l)$ or $\text{SE}(\mathfrak{b}_l)$; but in the latter case, $\mathfrak{d}$ must be the start node of its tile, and this contradicts Theorem 3.1. So $\mathfrak{d} = \text{NE}(\mathfrak{b}_l)$, and in particular $\text{ht}(\mathfrak{d}) > \text{ht}(\mathfrak{b}_l)$. By Lemma 4.2 $\text{en}(\text{tile}(\mathfrak{d})) = \mathfrak{b}_{l+1}$, and hence $\text{ht}(\mathfrak{b}_{l+1}) = \text{ht}(\mathfrak{d}) > \text{ht}(\mathfrak{b}_l)$.

So we have $\text{ht}(\mathfrak{a}) = \text{ht}(\mathfrak{b}_1) < \dots < \text{ht}(\mathfrak{b}_r) = \text{ht}(\mathfrak{c})$, so $\text{ht}(\mathfrak{a}) \leq \text{ht}(\mathfrak{c})$, with equality if and only if $r = 1$, in which case $\mathfrak{c} = \mathfrak{b}_1 = \text{en}(\text{tile}(\mathfrak{a}))$.

Now we can complete the proof. Assume there is a left-cover-expansive Dyck tiling T of $\lambda \setminus \mu$ and a right-cover-expansive Dyck tiling U . By Claim 2, we have $\text{ht}(\mathfrak{a}) \leq \text{ht}(\mathfrak{c})$, and symmetrically (since U exists) we have $\text{ht}(\mathfrak{a}) \geq \text{ht}(\mathfrak{c})$. Hence $\text{ht}(\mathfrak{a}) = \text{ht}(\mathfrak{c})$, and so $\mathfrak{a}, \mathfrak{c}$ lie in the same tile $t \in T$, which consists of the lowest node in every column of $\mathcal{C}$. Similarly, t is a tile in U . Removing t from $\lambda \setminus \mu$ yields a smaller skew Young diagram, and $T \setminus \{t\}$ is a left-cover-expansive Dyck tiling of this diagram, while $U \setminus \{t\}$ is a right-cover-expansive Dyck tiling. By induction $T \setminus \{t\} = U \setminus \{t\}$, and hence $T = U$. $\square$

Now we restrict attention to cover-expansive Dyck tilings. We shall need the following lemma, which examines the effect of the cover-expansive property on depths of nodes.

Lemma 4.4. Suppose $\lambda \supseteq \mu$, T is a cover-expansive Dyck tiling of $\lambda \setminus \mu$, and $\mathfrak{a}, \text{N}(\mathfrak{a}) \in \lambda \setminus \mu$. Then $\text{dp}(\text{N}(\mathfrak{a})) < \text{dp}(\mathfrak{a})$.

Proof. Suppose the lemma is false, and take $a \in \lambda \setminus \mu$ as far to the left as possible such that $N(a) \in \lambda \setminus \mu$ and $\text{dp}(N(a)) \geq \text{dp}(a)$. We have $NW(a) \in \lambda \setminus \mu$, and $NW(a)$ cannot be the end node of its tile by Proposition 4.1, so is attached to either a or $N(a)$.

If $NW(a)$ is attached to $N(a)$, then a must be attached to $SW(a)$ since it cannot be the start node of its tile (again by Proposition 4.1). But now $\text{dp}(NW(a)) - \text{dp}(SW(a)) = (\text{dp}(N(a)) + 1) - (\text{dp}(a) + 1) \geq 0$, contradicting the choice of a .

So suppose instead that $NW(a)$ is attached to a . Then $N(a)$ cannot be the start node of its tile, since $\text{dp}(N(a)) \geq \text{dp}(a) = \text{dp}(NW(a)) + 1 > 0$. So $N(a)$ is attached to $NW(N(a))$. But now $\text{dp}(NW(N(a))) - \text{dp}(NW(a)) = (\text{dp}(N(a)) - 1) - (\text{dp}(a) - 1) \geq 0$, and again the choice of a is contradicted. $\square$

We remark that, in contrast to the similar condition in Theorem 3.1(2) for cover-inclusive Dyck tilings, the condition in Lemma 4.4 does not imply the cover-expansive condition. For example, the unique Dyck tiling of $(2) \setminus \emptyset$ satisfies this condition (trivially) but is not right-cover-expansive.

The following lemma will be useful in the next section.

Lemma 4.5. *Suppose $\lambda \supseteq \mu$, and that μ has an addable node in column j . If there exists a cover-expansive Dyck tiling of $\lambda \setminus \mu$, then λ has either an addable or a removable node in column j .*

Proof. Suppose not, and let T be the cover-expansive Dyck tiling of $\lambda \setminus \mu$. There must be at least one node in column j of $\lambda \setminus \mu$ (otherwise the addable node of μ would also be an addable node of λ). Let a be the highest node in column j , and consider the nodes attached to a in T .

The lowest node in column j of $\lambda \setminus \mu$, namely m , is not attached to its SW neighbour (since this is not a node of $\lambda \setminus \mu$). Now let b be the highest node in column j which is not attached to its SW neighbour, and suppose $b \neq a$. Then the choice of b means that $NW(b)$ is attached to $N(b)$, so b is the start node of its tile; but $NW(b) \in \lambda \setminus \mu$, and this contradicts Proposition 4.1. So $b = a$.

So a is not attached to $SW(a)$, and symmetrically is not attached to $SE(a)$. If a is attached to either $NW(a)$ or $NE(a)$, then (since T is a Dyck tiling) it is attached to both $NW(a)$ and $NE(a)$, and hence both of these nodes belong to λ , so λ has an addable node in column j . The remaining possibility is that a is a singleton tile. But now Proposition 4.1 implies that neither $NW(a)$ nor $NE(a)$ is a node of $\lambda \setminus \mu$, so λ has a removable node in column j . $\square$

4.2 Recurrences

We now reassume the notation in Section 3: λ, μ are partitions with addable nodes l, m respectively in column j , and λ^+, μ^+ are the partitions obtained by adding these nodes.


Proposition 4.6. *Suppose $\lambda \supseteq \mu^+$, and let $a = S(l)$. Then:*

1. *there exists a cover-expansive Dyck tiling of $\lambda \setminus \mu$ in which a has depth 1 if and only if there exists a cover-expansive Dyck tiling of $\lambda \setminus \mu^+$;*
2. *there exists a cover-expansive Dyck tiling of $\lambda \setminus \mu$ in which a has depth greater than 1 if and only if there exists a cover-expansive Dyck tiling of $\lambda^+ \setminus \mu^+$.*

Proof. Suppose there is a cover-expansive Dyck tiling T of $\lambda \setminus \mu$. Arguing as in the proof of Lemma 4.5, a cannot be attached to either $SE(a)$ or $SW(a)$ (since μ has an addable node in column j), and $\{a\}$ cannot be a singleton tile (since λ has an addable node in column j). So a is

attached to both $NW(a)$ and $NE(a)$. So by the cover-expansive conditions, every node in column j is attached to its NE and NW neighbours.

1. Suppose a has depth 1 in T . We construct a cover-expansive Dyck tiling U of $\lambda \setminus \mu^+$ from T as illustrated in the following diagram.



Formally, to construct a cover-expansive Dyck tiling of $\lambda \setminus \mu^+$:

- replace $\text{tile}(a)$ with the two tiles comprising $\text{tile}(a) \setminus \{a\}$;
- for every node $b \neq a$ in column j , change $\text{tile}(b)$ by replacing b with $N(b)$.

It is easy to check that U really does give a cover-expansive Dyck tiling.

For the other direction, suppose T is a cover-expansive Dyck tiling of $\lambda \setminus \mu^+$. We claim that every node in column j must be attached to its SW and SE neighbours. If there are no nodes in column j (i.e. if $m = S(l)$) then this statement is trivial, so suppose otherwise; then $a \in \lambda \setminus \mu^+$. a cannot be attached to $NW(a)$ in T , since then every node in column $j-1$ would be attached to its SE neighbour; but the SE neighbour of the bottom node in column $j-1$ is not a node of $\lambda \setminus \mu^+$. Since a is not attached to $NW(a)$, it must be attached to $SW(a)$. Similarly a is attached to $SE(a)$, and the cover-expansive property implies that every node in column j is attached to its SE and SW neighbours as claimed. Now we construct a cover-expansive Dyck tiling of $\lambda \setminus \mu$ as follows:

- for each node b in column j , change $\text{tile}(b)$ by replacing b with $S(b)$;
- replace the tiles $\text{tile}(NW(a))$ and $\text{tile}(NE(a))$ with the tile obtained by joining these two tiles to a .

Again, it is easy to see that we have a cover-expansive Dyck tiling. Moreover, a has depth 1 in this tiling, since $NW(a)$ is the end node of its tile in T , and so has depth 0 (in both tilings).

2. Suppose a has depth greater than 1 in T . Construct a cover-expansive Dyck tiling U of $\lambda^+ \setminus \mu^+$ from T as follows: for each node b in column j , change $\text{tile}(b)$ by replacing b with $N(b)$.

Again, U is a Dyck tiling, since every node in column j of $\lambda \setminus \mu$ has depth greater than 1 in T . And the cover-expansive property for U follows from that for T .

The other direction is very similar. □

Corollary 4.7. $e_{\lambda\mu} = e_{\lambda\mu^+} + e_{\lambda^+\mu^+}$.

Proof. First suppose $\lambda \supseteq \mu^+$. The first paragraph of the proof of Proposition 4.6 shows that the node $a = S(l)$ must have positive depth in a cover-expansive Dyck tiling of $\lambda \setminus \mu$. Hence the result follows from Proposition 4.6.

Alternatively, suppose $\lambda \not\supseteq \mu^+$. Then we have $\lambda \supseteq \mu$ if and only if $\lambda^+ \supseteq \mu^+$, and if these conditions hold then $\lambda \setminus \mu = \lambda^+ \setminus \mu^+$; so $e_{\lambda\mu} = e_{\lambda^+\mu^+}$. □

Proposition 4.8. $e_{\lambda^+\mu} = e_{\lambda\mu}$.

Proof. Suppose T is a cover-expansive Dyck tiling of $\lambda^+\mu$. Then $\mathfrak{l}$ cannot be attached to $\text{SE}(\mathfrak{l})$, since then every node in column j would be attached to its SE neighbour; but $\text{SE}(\mathfrak{m})$ is not a node of $\lambda^+\mu$. Similarly $\mathfrak{l}$ is not attached to $\text{SW}(\mathfrak{l})$, so forms a singleton tile in T . Removing this tile yields a cover-expansive Dyck tiling of $\lambda\mu$.

In the other direction, suppose T is a cover-expansive Dyck tiling of $\lambda\mu$. From the proof of Proposition 4.6, $\mathfrak{S}(\mathfrak{l})$ is attached to both $\text{SE}(\mathfrak{l})$ and $\text{SW}(\mathfrak{l})$. Hence if we add $\mathfrak{l}$ as a singleton tile, the resulting tiling is a cover-expansive Dyck tiling of $\lambda^+\mu$. $\square$

5 Young permutation modules for two-part compositions

In this section we apply our results on Dyck tilings to compute transition coefficients for two bases for a certain module for the symmetric group.

5.1 The Young permutation module $\mathcal{M}^{(f,g)}$

Suppose $k \geq 0$, and let $\mathfrak{S}_k$ denote the symmetric group of degree k , with $\{t_1, \dots, t_{k-1}\}$ the set of Coxeter generators (so t_i is the transposition $(i, i+1)$). Let $\mathbb{F}$ be any field, and consider the group algebra $\mathbb{F}\mathfrak{S}_k$. Define $s_i = t_i - 1 \in \mathbb{F}\mathfrak{S}_k$ for $i = 1, \dots, k-1$.

Now write $k = f + g$ with f, g non-negative integers. Consider the Young permutation module $\mathcal{M}^{(f,g)}$ for $\mathbb{F}\mathfrak{S}_k$ indexed by the composition (f, g) . This is just the permutation module on the set of cosets of the maximal Young subgroup $\mathfrak{S}_f \times \mathfrak{S}_g$, and has the following presentation:

$$\mathcal{M}^{(f,g)} = \langle m \mid t_i m = m \text{ for } i \neq f \rangle.$$

Lemma 5.1. For $1 \leq i \leq k-2$, the element $s_i s_{i+1} s_i - s_i = t_i s_{i+1} s_i - s_{i+1} s_i - s_i$ annihilates $\mathcal{M}^{(f,g)}$.

Proof. In terms of permutations, the given element is

$$(i, i+2) - (i, i+1, i+2) - (i, i+2, i+1) + (i, i+1) + (i+1, i+2) - 1.$$

$\mathcal{M}^{(f,g)}$ may be viewed as the permutation module on the set of f -subsets of the set $\{1, \dots, k\}$; it is easily seen that the given element kills any such subset. $\square$

$\mathcal{M}^{(f,g)}$ has a basis indexed by the set of minimal left coset representatives of $\mathfrak{S}_f \times \mathfrak{S}_g$ in $\mathfrak{S}_k$; this set is in one-to-one correspondence with the set $\mathcal{P}_{f,g}$ of partitions λ for which $\lambda_1 \leq f$ and $\lambda'_1 \leq g$; given $\lambda \in \mathcal{P}_{f,g}$, we write the corresponding basis element as $t_\lambda m$, where

$$t_\lambda = (t_{\lambda_1+g} t_{\lambda_1+g+1} \dots t_{f+g-1}) (t_{\lambda_2+g-1} t_{\lambda_2+g} \dots t_{f+g-2}) \dots (t_{\lambda_g+1} t_{\lambda_g+2} \dots t_f).$$

Our objective here is to study the elements

$$s_\lambda = (s_{\lambda_1+g} s_{\lambda_1+g+1} \dots s_{f+g-1}) (s_{\lambda_2+g-1} s_{\lambda_2+g} \dots s_{f+g-2}) \dots (s_{\lambda_g+1} s_{\lambda_g+2} \dots s_f)$$

for $\lambda \in \mathcal{P}_{f,g}$. It is easy to see that $\{s_\lambda m \mid \lambda \in \mathcal{P}_{f,g}\}$ is also a basis for $\mathcal{M}^{(f,g)}$, since when each $s_\lambda m$ is expressed as a linear combination of the elements $t_\mu m$, the matrix of coefficients is unitriangular with respect to a suitable ordering.

We shall use cover-expansive Dyck tilings to describe this transition matrix explicitly, and then describe its inverse using cover-inclusive Dyck tilings. We also give a simple expression for the sum $\sum_{\lambda \in \mathcal{P}_{f,g}} t_\lambda m$ as a linear combination of the elements $s_\lambda m$, which will be useful in Section 6.

5.2 Change of basis

Our first result on transition coefficients is the following.

Theorem 5.2. *Suppose f, g are non-negative integers, and $\mu \in \mathcal{P}_{f,g}$. Then*

$$s_\mu m = \sum_{\lambda \in \mathcal{P}_{f,g}} (-1)^{|\lambda|+|\mu|} e_{\lambda\mu} t_\lambda m.$$

We begin with some simple observations concerning the actions of the generators $t_1, \dots, t_{k-1}$ on the basis elements $t_\lambda m$. We continue to use the Russian convention for Young diagrams.

Lemma 5.3. *Suppose $\lambda \in \mathcal{P}_{f,g}$ and $1 \leq i < k$. Then:*

- *if λ has an addable node $\mathfrak{l}$ in column $i - g$, then $t_i t_\lambda m = t_{\lambda^+} m$, where λ^+ is the partition obtained by adding $\mathfrak{l}$ to λ ;*
- *if λ has a removable node $\mathfrak{l}$ in column $i - g$, then $t_i t_\lambda m = t_{\lambda^-} m$, where λ^- is the partition obtained by removing $\mathfrak{l}$ from λ ;*
- *if λ has neither an addable nor a removable node in column $i - g$, then $t_i t_\lambda m = t_\lambda m$.*

Proof. This is an easy consequence of the definitions and the Coxeter relations. $\square$

Proof of Theorem 5.2. We proceed by downwards induction on $|\mu|$. If $\mu = (f^g)$, then $s_\mu = t_\mu = 1$ and the result follows. Assuming $\mu \neq (f^g)$, μ has an addable node in column j , for some $-g < j < f$. We let μ^+ denote the partition obtained by adding this addable node; then we have $s_\mu = s_{j+g} s_{\mu^+}$, and by induction

$$s_{\mu^+} m = \sum_{\lambda \in \mathcal{P}_{f,g}} (-1)^{|\lambda|+|\mu^+|} e_{\lambda\mu^+} t_\lambda m.$$

Let $\mathcal{P}_{f,g}^+$ denote the set of $\lambda \in \mathcal{P}_{f,g}$ having an addable node in column j , and for $\lambda \in \mathcal{P}_{f,g}^+$ let λ^+ denote the partition obtained by adding this addable node; similarly, let $\mathcal{P}_{f,g}^-$ denote the set of partitions in $\mathcal{P}_{f,g}$ having a removable node in column j , and for each $\lambda \in \mathcal{P}_{f,g}^-$ let λ^- denote the partition obtained by removing this removable node. Note that the functions $\lambda \mapsto \lambda^+$ and $\lambda \mapsto \lambda^-$ define mutually inverse bijections between $\mathcal{P}_{f,g}^+$ and $\mathcal{P}_{f,g}^-$. Now

$$\begin{aligned} s_\mu m &= (t_{j+g} - 1) \sum_{\lambda \in \mathcal{P}_{f,g}} (-1)^{|\lambda|+|\mu^+|} e_{\lambda\mu^+} t_\lambda m \\ &= \sum_{\lambda \in \mathcal{P}_{f,g}^+} (-1)^{|\lambda|+|\mu^+|} e_{\lambda\mu^+} (t_{\lambda^+} - t_\lambda) m + \sum_{\lambda \in \mathcal{P}_{f,g}^-} (-1)^{|\lambda|+|\mu^+|} e_{\lambda\mu^+} (t_{\lambda^-} - t_\lambda) m \quad \text{by Lemma 5.3} \\ &= \sum_{\lambda \in \mathcal{P}_{f,g}^+} (-1)^{|\lambda|+|\mu|} (e_{\lambda\mu^+} + e_{\lambda^+\mu^+}) t_\lambda m + \sum_{\lambda \in \mathcal{P}_{f,g}^-} (-1)^{|\lambda|+|\mu|} (e_{\lambda\mu^+} + e_{\lambda^-\mu^+}) t_\lambda m \\ &= \sum_{\lambda \in \mathcal{P}_{f,g}^+} (-1)^{|\lambda|+|\mu|} e_{\lambda\mu} t_\lambda m + \sum_{\lambda \in \mathcal{P}_{f,g}^-} (-1)^{|\lambda|+|\mu|} e_{\lambda^-\mu} t_\lambda m \quad \text{by Proposition 4.6} \\ &= \sum_{\lambda \in \mathcal{P}_{f,g}^+ \cup \mathcal{P}_{f,g}^-} (-1)^{|\lambda|+|\mu|} e_{\lambda\mu} t_\lambda m \quad \text{by Proposition 4.8} \\ &= \sum_{\lambda \in \mathcal{P}_{f,g}} (-1)^{|\lambda|+|\mu|} e_{\lambda\mu} t_\lambda m \quad \text{by Lemma 4.5.} \end{aligned}$$

So Theorem 5.2 follows by induction. $\square$

Now we give our second main result on transition coefficients.

Theorem 5.4. *Suppose f, g are non-negative integers, and $\mu \in \mathcal{P}_{f,g}$. Then*

$$t_\mu m = \sum_{\lambda \in \mathcal{P}_{f,g}} i_{\lambda\mu} s_\lambda m.$$

The proof of this result is rather more difficult. To begin with, we compute the actions of $s_1, \dots, s_{k-1}$ on the basis elements $s_\lambda m$.

Proposition 5.5. *Suppose $\mu \in \mathcal{P}_{f,g}$ and $1 \leq i < k$. Then exactly one of the following occurs.*

1. μ has an addable node in column $i - g$. In this case $s_i s_\mu m = -2s_\mu m$.
2. μ has a removable node in column $i - g$. In this case $s_i s_\mu m = s_{\mu^-} m$, where μ^- denotes the partition obtained by removing this node.
3. For some $0 \leq a \leq g$ we have $\mu_a > i - g + a > \mu_{a+1}$ (where the left-hand inequality is regarded as automatically true in the case $a = 0$).

(a) If $\mu_w < i - g + 2a - w$ for all $a < w \leq g$, then $s_i s_\mu m = 0$.

(b) Otherwise, let $w > a$ be minimal such that $\mu_w = i - g + 2a - w$, and set

$$\mu^{a,w} = (\mu_1, \dots, \mu_a, i - g + a, \mu_{a+1} + 1, \dots, \mu_{w-1} + 1, \mu_{w+1}, \dots, \mu_g).$$

Then $s_i s_\mu m = s_{\mu^{a,w}} m$.

4. For some $1 \leq a < g$ we have $\mu_a = \mu_{a+1} = i - g + a$.

(a) If $i + 2a > k$ and $\mu_w < i - g + 2a - w$ for $w = 1, \dots, a - 1$, then $s_i s_\mu m = 0$.

(b) Otherwise, let $w < a$ be maximal such that $\mu_w \geq i - g + 2a - w$ (taking $w = 0$ if there is no such w), and define

$$\mu^{w,a} = (\mu_1, \dots, \mu_w, i - g + 2a - w, \mu_{w+1} + 1, \dots, \mu_{a-1} + 1, \mu_{a+1}, \dots, \mu_g).$$

Then $s_i s_\mu m = s_{\mu^{w,a}} m$.

Examples.

1. Suppose $f = 4, g = 5$ and $\mu = (4, 2^4)$. Then

$$s_\mu = s_6 s_7 s_5 s_6 s_4 s_5 s_3 s_4.$$

Taking $i = 3$, we find that μ satisfies condition 4(a) of Proposition 5.5, with $a = 4$. And indeed

$$\begin{aligned} s_3 s_\mu m &= s_3 s_6 s_7 s_5 s_6 s_4 s_5 s_3 s_4 m \\ &= s_6 s_7 s_5 s_6 (s_3 s_4 s_3) s_5 s_4 m \\ &= s_6 s_7 s_5 s_6 s_3 s_5 s_4 m && \text{by Lemma 5.1} \\ &= s_6 s_7 (s_5 s_6 s_5) s_3 s_4 m \\ &= s_6 s_7 s_5 s_3 s_4 m && \text{by Lemma 5.1} \\ &= s_6 s_5 s_3 s_4 s_7 m \\ &= 0 && \text{since } s_7 m = 0. \end{aligned}$$

2. Now suppose $f = 4, g = 5$ and $\mu = (4, 2, 1^3)$. Then

$$s_\mu = s_6 s_7 s_4 s_5 s_6 s_3 s_4 s_5 s_2 s_3 s_4.$$

Taking $i = 2$, we find that μ satisfies condition 4(b) of Proposition 5.5, with $a = 4$ and $w = 1$, giving $\mu^{w,a} = (4^2, 3, 2, 1)$. Now we have

$$\begin{aligned} s_2 s_\mu m &= s_2 s_6 s_7 s_4 s_5 s_6 s_3 s_4 s_5 s_2 s_3 s_4 m \\ &= s_6 s_7 s_4 s_5 s_6 (s_2 s_3 s_2) s_4 s_5 s_3 s_4 m \\ &= s_6 s_7 s_4 s_5 s_6 s_2 s_4 s_5 s_3 s_4 m \\ &= s_6 s_7 (s_4 s_5 s_4) s_6 s_2 s_5 s_3 s_4 m \\ &= s_6 s_7 s_4 s_6 s_2 s_5 s_3 s_4 m \\ &= (s_6 s_7 s_6) s_4 s_2 s_5 s_3 s_4 m \\ &= s_6 s_4 s_2 s_5 s_3 s_4 m \\ &= s_{(4^2, 3, 2, 1)} m. \end{aligned}$$

Proof of Proposition 5.5. To see that μ satisfies exactly one of the four conditions, let l denote the lowest node in column j which is not a node of μ , and consider whether $SE(l)$ and $SW(l)$ lie in μ ; for the purposes of this argument we regard all nodes of the form $(a, 0)$ or $(0, b)$ as lying in μ . If both $SE(l)$ and $SW(l)$ lie in μ , then l is an addable node and we are in case (1). If neither lies in μ , then $S(l)$ is a removable node, and we are in case (2). If $SE(l) \in \mu \not\subset SW(l)$, then we are in case (3), while if $SE(l) \notin \mu \supset SW(l)$ then we are in case (4).

Now we analyse the four cases in Proposition 5.5.

1. Since μ has an addable node in column $i - g$, s_μ can be written in the form $s_i s_{\mu^+}$. We have $s_i^2 = -2s_i$, and the result follows.
2. By definition (and the fact that s_i and s_j commute when $j \neq i \pm 1$) $s_{\mu^-} = s_i s_\mu$.
3. (a) We proceed by induction on i . Consider the case $a = g$. In this case s_μ only involves terms s_j for $j \geq i + 2$, so (since $s_i m = 0$) we get $s_i s_\mu m = 0$.
Now suppose $a < g$, and set

$$\bar{\mu} = (\mu_1, \dots, \mu_a, i - g + a, \mu_{a+2}, \dots, \mu_g).$$

Then $\bar{\mu}$ satisfies the inductive hypothesis, with a replaced by $a + 1$ and i replaced by $i - 2$. Furthermore, we can write

$$s_\mu = s_{\mu_{a+1}+g-a} \dots s_{i-1} s_{\bar{\mu}}.$$

$\bar{\mu}$ has an addable node $(a + 1, i - g + a + 1)$ in column $i - g$, so $s_{\bar{\mu}} m$ can be written in the form $s_i x$ for some $x \in \mathcal{M}^{(f,g)}$. So we have

$$\begin{aligned} s_i s_\mu m &= s_i s_{\mu_{a+1}+g-a} \dots s_{i-1} s_i x \\ &= s_{\mu_{a+1}+g-a} \dots s_{i-2} s_i s_{i-1} s_i x \\ &= s_{\mu_{a+1}+g-a} \dots s_{i-2} s_i x && \text{by Lemma 5.1} \\ &= s_{\mu_{a+1}+g-a} \dots s_{i-2} s_{\bar{\mu}} m \\ &= 0 && \text{by induction.} \end{aligned}$$

- (b) We use induction on $w - a$. If $w = a + 1$, then we have $s_\mu = s_{i-1}s_{\mu^{a,w}}$, and $\mu^{a,w}$ has an addable node $(a + 1, i - g + a + 1)$ in column $i - g$, so $s_{\mu^{a,w}}m$ can be written as $s_i x$ for some $x \in \mathcal{M}^{(f,g)}$. So by Lemma 5.1

$$s_i s_\mu m = s_i s_{i+1} s_i x = s_i x = s_{\mu^{a,w}} m.$$

Now assume that $w > a + 1$, and as above set

$$\bar{\mu} = (\mu_1, \dots, \mu_a, i - g + a, \mu_{a+2}, \dots, \mu_g).$$

Then $\bar{\mu}$ satisfies the inductive hypothesis, with a replaced by $a + 1$ and i replaced by $i - 2$, and with the same value of w , yielding

$$\bar{\mu}^{a+1,w} = (\mu_1, \dots, \mu_a, i - g + a, i - g + a - 1, \mu_{a+2} + 1, \dots, \mu_{w-1} + 1, \mu_{w+1}, \dots, \mu_g).$$

As above we write $s_{\bar{\mu}} = s_i x$, and obtain

$$\begin{aligned} s_i s_\mu m &= s_{\mu_{a+1}+g-a} \dots s_{i-2} s_{\bar{\mu}} m \\ &= s_{\mu_{a+1}+g-a} \dots s_{i-3} s_{\bar{\mu}^{a+1,w}} m && \text{by induction} \\ &= s_{\mu^{a,w}} m. \end{aligned}$$

4. This is symmetrical to case (3), replacing partitions with their conjugates and swapping f and g . $\square$

We now seek to re-phrase Proposition 5.5 so that for a given λ, μ we can write down the coefficient of $s_\lambda m$ in $s_i s_\mu m$. Fix i , and suppose that as in part (3) of Proposition 5.5 we have $\mu_a > i - g + a > \mu_{a+1}$ for some $1 \leq a \leq g$. Let $w > a$ be minimal such that $\mu_w = i - g + 2a - w$, assuming there is such a w , and write $\mu \xrightarrow{i} \mu^{a,w}$, where $\mu^{a,w}$ is as defined in Proposition 5.5.

Now recall some notation from Section 2: if j is fixed and λ is a partition with an addable node $\mathfrak{l}$ in column j , then X_λ denotes the set of all integers x such that λ has a removable node $\mathfrak{n}$ in column $j + x$, with $\text{ht}(\mathfrak{n}) = \text{ht}(\mathfrak{l}) - 1$, and $\text{ht}(\mathfrak{p}) < \text{ht}(\mathfrak{l})$ for all nodes $\mathfrak{p}$ in all columns between j and $j + x$. X_λ^+ denotes the set of positive elements of X_λ , and X_λ^- the set of negative elements. Given $x \in X_\lambda^+$, $\lambda^{[x]}$ is the partition obtained from λ by removing the highest node in each column from $j + 1$ to $j + x$; $\lambda^{[x]}$ is defined similarly for $x \in X_\lambda^-$.

Lemma 5.6. Suppose $\lambda, \mu \in \mathcal{P}_{f,g}$ and $1 \leq i < k$, and set $j = i - g$. Then $\mu \xrightarrow{i} \lambda$ if and only if λ has an addable node in column j and $\mu = \lambda^{[x]}$ for some $x \in X_\lambda^-$.

Proof.

($\Rightarrow$) If $\mu \xrightarrow{i} \lambda$ then we have

$$\lambda = \mu^{a,w} = (\mu_1, \dots, \mu_a, j + a, \mu_{a+1} + 1, \dots, \mu_{w-1} + 1, \mu_{w+1}, \dots, \mu_g),$$

where $\mu_a > j + a > \mu_{a+1}$ and $w > a$ is minimal such that $\mu_w \geq j + 2a - w$. Observe that λ has an addable node $\mathfrak{l} = (a + 1, j + a + 1)$ in column j . The choice of w implies that $\mu_{w-1} = \mu_w = j + 2a - w$, and so λ has a removable node $\mathfrak{n} = (w, j + 2a - w + 1)$ in column $j + x$, where $x = 2(a - w) + 1$. Furthermore, $\text{ht}(\mathfrak{n}) = j + 2a = \text{ht}(\mathfrak{l}) - 1$, and we claim that all the nodes in columns $j + x + 1, \dots, j - 1$ of λ have height at most $j + 2a$: if this

condition fails, then it must fail for a node of the form (x, λ_x) with $a + 2 \leq x \leq w - 1$, but for x in this range we have $\text{ht}((x, \lambda_x)) = x + \lambda_x = x + \mu_{x-1} + 1 \leq x + j + 2a - x$. So $x \in X_\lambda^-$.

To construct μ from λ we remove the nodes

$$\begin{aligned} & (w, \mu_w + 1), (w, \mu_w + 2), \dots, (w, \mu_{w-1} + 1), \\ & (w - 1, \mu_{w-1} + 1), (w - 1, \mu_{w-1} + 2), \dots, (w - 1, \mu_{w-2} + 1), \\ & \dots \\ & (a + 2, \mu_{a+2} + 1), (a + 2, \mu_{a+2} + 2), \dots, (a + 2, \mu_{a+1} + 1), \\ & (a + 1, \mu_{a+1} + 1), (a + 1, \mu_{a+1} + 2), \dots, (a + 1, j + a), \end{aligned}$$

which lie in columns $j + x, \dots, j - 1$ respectively. So $\mu = \lambda^{[x]}$.

($\Leftarrow$) Suppose λ has an addable node in column j . Writing this node as $(a + 1, \lambda_{a+1} + 1)$, we have $a = 0$ or $\lambda_a > \lambda_{a+1}$, and $\lambda_{a+1} - a = j$.

Now suppose $x \in X_\lambda^-$. Then λ has a removable node $\mathbf{n} = (w, \lambda_w)$ in column $j + x$, and $w > a$ since $x < 0$. Then $\lambda_w > \lambda_{w+1}$ and $\lambda_w - w = j + x$, and (since $\text{ht}(\mathbf{n}) = \text{ht}(\mathbf{l}) - 1$) $w + \lambda_w = a + 1 + \lambda_{a+1}$. The fact that $\text{ht}(\mathbf{p}) < \text{ht}(\mathbf{l})$ for all nodes $\mathbf{p}$ in columns $j + x + 1, \dots, j - 1$ implies in particular that $\text{ht}((x, \lambda_x)) < \text{ht}(\mathbf{l})$ for $a + 2 \leq x \leq w - 1$, i.e. $x + \lambda_x \leq a + 1 + \lambda_{a+1}$.

Constructing $\lambda^{[x]}$ involves removing the nodes

$$\begin{aligned} & (w, \lambda_w), (w - 1, \lambda_w), (w - 1, \lambda_w + 1), \dots, (w - 1, \lambda_{w-1}), \\ & (w - 2, \lambda_{w-1}), (w - 2, \lambda_{w-1} + 1), \dots, (w - 2, \lambda_{w-2}), \\ & \dots \\ & (a + 1, \lambda_{a+2}), (a + 1, \lambda_{a+2} + 1), \dots, (a + 1, \lambda_{a+1}), \end{aligned}$$

so

$$\lambda^{[x]} = (\lambda_1, \dots, \lambda_a, \lambda_{a+2} - 1, \dots, \lambda_w - 1, \lambda_w - 1, \lambda_{w+1}, \dots, \lambda_g).$$

Setting $\mu = \lambda^{[x]}$, the (in)equalities observed above for λ give $\mu_a > j + a > \mu_{a+1}$, $\mu_x < j + 2a - x$ for $a < x < w$ and $\mu_w = j + 2a - w$. Furthermore, λ equals the partition $\mu^{a,w}$, so $\mu \xrightarrow{i} \lambda$. $\square$

We now note a counterpart to this for part (4) of Proposition 5.5, which follows by conjugating partitions. Suppose that $\mu \in \mathcal{P}_{f,g}$ satisfies the conditions of (4), and that $\mu^{w,a}$ is defined, and write $\mu \xrightarrow{i} \mu^{w,a}$.

Lemma 5.7. *Suppose $\lambda, \mu \in \mathcal{P}_{f,g}$ and $1 \leq i < k$, and set $j = i - g$. Then $\mu \xrightarrow{i} \lambda$ if and only if λ has a removable node in column j and $\mu = \lambda^{[x]}$ for some $x \in X_\lambda^+$.*

The following result is now immediate from Proposition 5.5 and Lemmas 5.6 and 5.7.

Corollary 5.8. *Suppose $\mu \in \mathcal{P}_{f,g}$, and write $s_i s_\mu m$ as a linear combination $\sum_{\lambda \in \mathcal{P}_{f,g}} a_\lambda s_\lambda m$. Then a_λ equals:*

- -2 , if λ has an addable node in column $i - g$ and $\mu = \lambda$;

- 1, if λ has an addable node in column $i - g$, and
 - μ is the partition obtained by adding this node, or
 - $\mu = \lambda^{[x]}$ for some $x \in X_\lambda$;
- 0, otherwise.

Proof of Theorem 5.4. As with Theorem 5.2, we proceed by downwards induction on $|\mu|$. If $\mu = (f^g)$, then $s_\mu = t_\mu$ and the result follows, since clearly $i_{\lambda\lambda} = 1$.

Assuming $\mu \neq (f^g)$, μ has an addable node in column j for some $-g < j < f$. We let μ^+ denote the partition obtained by adding this addable node. Then $t_\mu = t_{j+g}t_{\mu^+}$, and by induction

$$t_{\mu^+}m = \sum_{\lambda \in \mathcal{P}_{f,g}} i_{\lambda\mu^+} s_\lambda m.$$

As before, we write $\mathcal{P}_{f,g}^+$ for the set of $\lambda \in \mathcal{P}_{f,g}$ having an addable node in column j , and for $\lambda \in \mathcal{P}_{f,g}^+$ we let λ^+ denote the partition obtained by adding this addable node. We have

$$\begin{aligned} t_\mu m &= \sum_{\lambda \in \mathcal{P}_{f,g}} i_{\lambda\mu^+}(s_{j+g} + 1)s_\lambda m \\ &= \sum_{\lambda \in \mathcal{P}_{f,g}} i_{\lambda\mu^+} s_\lambda m + \sum_{\lambda \in \mathcal{P}_{f,g}^+} \left(i_{\lambda^+\mu^+} - 2i_{\lambda\mu^+} + \sum_{x \in X_\lambda} i_{\lambda^{[x]}\mu^+} \right) s_\lambda m && \text{by Corollary 5.8} \\ &= \sum_{\lambda \in \mathcal{P}_{f,g} \setminus \mathcal{P}_{f,g}^+} i_{\lambda\mu^+} s_\lambda m + \sum_{\lambda \in \mathcal{P}_{f,g}^+} \left(i_{\lambda^+\mu^+} - i_{\lambda\mu^+} + \sum_{x \in X_\lambda} i_{\lambda^{[x]}\mu^+} \right) s_\lambda m \\ &= \sum_{\lambda \in \mathcal{P}_{f,g} \setminus \mathcal{P}_{f,g}^+} i_{\lambda\mu} s_\lambda m + \sum_{\lambda \in \mathcal{P}_{f,g}^+} i_{\lambda\mu} s_\lambda m && \text{by Propositions 3.3 and 3.9} \\ &= \sum_{\lambda \in \mathcal{P}_{f,g}} i_{\lambda\mu} s_\lambda m. \end{aligned}$$

Theorem 5.4 follows by induction. □

5.3 The sum of the elements $t_\mu m$

It will be critical in the next section to be able to express $\sum_{\lambda \in \mathcal{P}_{f,g}} t_\lambda m$ in terms of the basis elements $s_\lambda m$. To enable us to do this, we define an integer-valued function F on partitions.

If λ is a partition, we let $\partial^a \lambda$ denote the partition obtained by removing the first a parts of λ , i.e. the partition $(\lambda_{a+1}, \lambda_{a+2}, \dots)$, and let $\partial_b \lambda$ denote the partition obtained by reducing all the parts by b (and deleting all negative parts), i.e. the partition $(\max\{\lambda_1 - b, 0\}, \max\{\lambda_2 - b, 0\}, \dots)$. We write $\partial_b^a \lambda = \partial^a \partial_b \lambda$.

Now we can define F recursively. Set $F(\emptyset) = 1$. Given a partition $\lambda \neq \emptyset$, take a node (a, b) of λ for which $a + b$ is maximal (call this a *highest* node of λ) and define

$$\sigma = \partial^a \lambda, \quad \tau = \partial_b \lambda.$$

Now set

$$F(\lambda) = \binom{a+b}{a} F(\sigma) F(\tau).$$

Lemma 5.9. $F(\lambda)$ is well-defined, i.e. does not depend on the choice of (a, b) .

Proof. Suppose (a, b) and $(\hat{a}, \hat{b})$ are both highest nodes of λ ; in particular, $a + b = \hat{a} + \hat{b}$. Assume without loss that $a < \hat{a}$. Let $\sigma = \partial^a \lambda$ and $\tau = \partial_b \lambda$ as above, and define $\hat{\sigma}, \hat{\tau}$ correspondingly. By induction $F(\sigma), F(\tau), F(\hat{\sigma}), F(\hat{\tau})$ are well-defined, and we must show that

$$F(\sigma)F(\tau) \binom{a+b}{a} = F(\hat{\sigma})F(\hat{\tau}) \binom{\hat{a}+\hat{b}}{\hat{a}}.$$

Now $(\hat{a} - a, \hat{b})$ is a highest node of σ , so we have

$$F(\sigma) = F(\hat{\sigma})F(\xi) \binom{\hat{a} - a + \hat{b}}{\hat{a} - a},$$

where $\xi = \partial_{\hat{b}}^a \lambda$. Also, $(a, b - \hat{b})$ is a highest node of $\hat{\tau}$, so that

$$F(\hat{\tau}) = F(\xi)F(\tau) \binom{a + b - \hat{b}}{a}.$$

So (since $a + b = \hat{a} + \hat{b}$) we obtain

$$F(\sigma)F(\tau)F(\xi) \binom{\hat{a}}{a} = F(\hat{\sigma})F(\hat{\tau})F(\xi) \binom{b}{\hat{b}}.$$

Dividing both sides by $F(\xi)\hat{a}!b!$ (which is obviously positive) and multiplying by $(a+b)!(\hat{a}-a)! = (\hat{a}+\hat{b})!(\hat{b}-b)!$ gives the result. $\square$

Remark. In the notation of the above lemma, we have

$$F(\lambda) = F(\hat{\sigma})F(\xi)F(\tau)(a, \hat{a} - a, \hat{b})!,$$

where we use $(x, y, z)!$ to denote the trinomial coefficient $\frac{(x+y+z)!}{x!y!z!}$. We will refer to this as ‘factorising using the nodes (a, b) and $(\hat{a}, \hat{b})$ ’. This readily generalises to factorising using any number of highest nodes, with a corresponding multinomial coefficient.

Our main objective is to prove the following statement.

Theorem 5.10. Suppose f, g are non-negative integers. Then

$$\sum_{\mu \in \mathcal{P}_{f,g}} t_{\mu} m = \sum_{\lambda \in \mathcal{P}_{f,g}} F(\lambda) s_{\lambda} m.$$

We prove this using the following recurrence, in which we again use the notation X_{λ} and $\lambda^{[x]}$ from Section 2.

Proposition 5.11. Assume that λ is a partition with an addable node $\mathfrak{l}$ in column j , and let λ^+ denote the partition obtained by adding this node. Then

$$F(\lambda^+) + \sum_{x \in X_{\lambda}} F(\lambda^{[x]}) = 2F(\lambda).$$

Proof. Write l as (a, b) . Suppose first of all that l is not the unique highest node of λ^+ . Take a highest node $m = (\hat{a}, \hat{b}) \neq l$, and suppose without loss that $\hat{a} < a$. Since the height of m is at least the height of l , $\hat{b} - \hat{a} - j$ is strictly larger than any element of X_λ . So m is a highest node of $\lambda^{[x]}$ for every $x \in X_\lambda$, as well as of λ and λ^+ . By induction the result holds for the partition $(\lambda_{\hat{a}+1}, \lambda_{\hat{a}+2}, \dots)$, and so it holds for λ , factorising $F(\lambda)$ (and also $F(\lambda^+)$ and $F(\lambda^{[x]})$ for each x) using the node $(\hat{a}, \hat{b})$.

So we can assume that l is the unique highest node of λ^+ . Let $(a_1, b_1), \dots, (a_r, b_r)$ be the nodes of λ to the right of l for which $a_x + b_x = a + b - 1$; order these nodes so that $a_1 > \dots > a_r$. Note that if $r \geq 1$, then automatically $(a_1, b_1) = (a - 1, b)$. Similarly, let $(c_1, d_1), \dots, (c_s, d_s)$ be the nodes to the left of l for which $c_x + d_x = a + b - 1$.

We compute $F(\lambda^+)$ by first factorising using the node l . We may also factorise $F(\lambda)$ using the nodes $(a - 1, b)$ and $(a, b - 1)$ (omitting either or both of these, if a or b equals 1), and we readily obtain

$$F(\lambda^+) = \frac{a+b}{ab} F(\lambda). \quad (\dagger)$$

Now we consider the partitions $\lambda^{[x]}$, for $x \in X_\lambda^+$. From our assumptions so far, X_λ^+ is the set of integers $x_i = b_i - a_i - b + a$, for $i = 1, \dots, r$. Choose such an i , and set $\kappa = \lambda^{[x_i]}$. We wish to compare $F(\kappa)$ with $F(\lambda)$. To do this, we define

$$\gamma = \partial_{b-1}^{a_{i+1}} \lambda, \quad \delta = \partial_{b-1}^{a_{i+1}} \kappa$$

(where we put $a_{i+1} = 0$ if $i = r$), and we claim that $F(\kappa)/F(\lambda) = F(\delta)/F(\gamma)$. If $s = 0$ and $i = r$, then this is trivial since $\gamma = \lambda$ and $\delta = \kappa$; assuming instead that $s \geq 1$ or $i < r$, factorising both $F(\kappa)$ and $F(\lambda)$ at the nodes $(c_1, d_1), \dots, (c_s, d_s), (a_{i+1}, b_{i+1}), \dots, (a_r, b_r)$ (i.e. at all the highest nodes of κ) gives the result.

Now γ has highest nodes $(a - 1 - a_{i+1}, 1)$ and $(a_i - a_{i+1}, b_i - b + 1)$ (which coincide if $i = 1$), and we factorise $F(\gamma)$ at these nodes, obtaining

$$F(\gamma) = F(\mu)F(\nu) \frac{(a - a_{i+1})!}{(a - 1 - a_i)!(a_i - a_{i+1})!},$$

where $\mu = \partial_1^{a_i - a_{i+1}} \gamma$ and $\nu = \partial_{b_i - b + 1} \gamma$.

Claim.
$$\frac{F(\delta)}{F(\gamma)} = \frac{a_i - a_{i+1}}{(a - a_i)(a - a_{i+1})}.$$

Proof. In the case where $a_i = a_{i+1} + 1$, we have $\mu = \delta$ and $\nu = \emptyset$, so the factorisation of $F(\gamma)$ above gives the result. If $a_i > a_{i+1} + 1$, then δ has a highest node $(a_i - a_{i+1} - 1, b_i - b + 1)$, and factorising at this node we get

$$F(\delta) = F(\mu)F(\nu) \binom{a_i - a_{i+1} + b_i - b}{a_i - a_{i+1} - 1}.$$

The claim follows from this factorisation and the factorisation of $F(\gamma)$ above, using the fact that $a_i + b_i = a + b - 1$.

This claim yields

$$\frac{F(\kappa)}{F(\lambda)} = \frac{F(\delta)}{F(\gamma)} = \frac{1}{a - a_i} - \frac{1}{a - a_{i+1}},$$

and summing over i we obtain

$$\frac{\sum_{x \in X_\lambda^+} F(\lambda^{[x]})}{F(\lambda)} = \frac{1}{a - a_1} - \frac{1}{a - a_{r+1}} = 1 - \frac{1}{a}.$$

A symmetrical calculation applies to partitions $\lambda^{[x]}$ for $x \in X_\lambda^-$, yielding

$$\frac{\sum_{x \in X_\lambda^-} F(\lambda^{[x]})}{F(\lambda)} = 1 - \frac{1}{b}.$$

Adding these two results together and combining with (†) gives the result. $\square$

Corollary 5.12. *For $1 \leq i < k$ we have*

$$s_i \sum_{\lambda \in \mathcal{P}_{f,g}} F(\lambda) s_\lambda m = 0.$$

Proof. As above, write $\mathcal{P}_{f,g}^+$ for the set of $\lambda \in \mathcal{P}_{f,g}$ having an addable node in column $i - g$, and for each such λ let λ^+ denote the partition obtained by adding this addable node. Then by Corollary 5.8 we have

$$s_i \sum_{\lambda \in \mathcal{P}_{f,g}} F(\lambda) s_\lambda m = \sum_{\lambda \in \mathcal{P}_{f,g}^+} \left(-2F(\lambda) + F(\lambda^+) - \sum_{x \in X_\lambda} F(\lambda^{[x]}) \right) s_\lambda m,$$

and each summand on the right-hand side is zero by Proposition 5.11. $\square$


Proof of Theorem 5.10. Corollary 5.12 shows that $\sum_\lambda F(\lambda) s_\lambda m$ is killed by all of $s_1, \dots, s_{k-1}$, and hence lies in a trivial submodule of $\mathcal{M}^{(f,g)}$. But $\mathcal{M}^{(f,g)}$ is a permutation module for a transitive action of $\mathfrak{S}_k$, and so has a unique trivial submodule, spanned by the sum of the elements of the permutation basis, i.e. $\sum_\lambda t_\lambda m$. So the two sides agree up to multiplication by a scalar. On the other hand, it is clear that when the right-hand side is expressed as a linear combination of the $t_\lambda m$, the coefficient of $t_\emptyset m$ is $F(\emptyset) = 1$, so in fact the two sides are equal. $\square$

5.4 Results of Kenyon, Kim, Mészáros, Panova and Wilson

As we have mentioned, cover-inclusive Dyck tilings were introduced by Kenyon and Wilson in [KW], and in fact the main result in their paper has a bearing on the results in the present paper.

First we explain how cover-expansive Dyck tilings appear in disguise in [KW]. Given a partition λ , we define a sequence of parentheses (and), by ‘reading along the boundary’ of λ , as follows. Working from left to right, for each node α such that $\text{NW}(\alpha)$ is not a node of λ , we write a), and for each node α such that $\text{NE}(\alpha)$ is not a node of λ we write a (. (If neither $\text{NW}(\alpha)$ nor $\text{NE}(\alpha)$ is a node of λ , then we write)(.) We then append an infinite string of (s at the start of the sequence, and an infinite sequence of)s at the end. The resulting doubly infinite sequence is called the *parenthesis sequence* of λ .

Example. Let $\lambda = (5, 3^2, 1)$. Then the parenthesis sequence of λ is $\cdots (((()())())()) \cdots$, as we see from the Young diagram, in which we mark the boundary in bold, extending it infinitely far to the north-west and north-east; segments $\backslash$ contribute a $($ to the parenthesis sequence, while segments $/$ contribute a $)$.



Now partition the parenthesis expression for λ into pairs, in the usual way for pairing up parentheses: each $($ is paired with the first subsequent $)$ for which there are an equal number of $)$ s and $($ s in between.

Example. Continuing the last example, we illustrate the pairs by numbering the $)$ s in increasing order, and numbering each $($ with the same number as its corresponding $)$:

$$\begin{array}{cccccccccccccccc} \cdots & (& (& (& (&) & (&) &) & (& (&) &) & (&) &) & \cdots \\ \cdots & 8 & 7 & 3 & 1 & 1 & 2 & 2 & 3 & 5 & 4 & 4 & 5 & 6 & 6 & 7 & 8 & \cdots \end{array}$$

Remark. We remark that in [KW] and elsewhere, finite parenthesis expressions are used: instead of appending infinite strings of $($ and $)$ at the start and end of the expression obtained from the partition, one appends sufficiently long finite strings that the resulting expression is *balanced*, meaning that it contains equally many $($ s and $)$ s, and that any initial segment contains at least as many $($ s as $)$ s. We find our convention more straightforward in the present context, and translation between the two conventions is very easy.

Now we recall a definition from [KW]; it is phrased there in terms of finite Dyck paths, but it is more convenient for us to phrase it in terms of partitions. Given two partitions λ, μ , write $\lambda \stackrel{\Omega}{\leftarrow} \mu$ if the parenthesis expression for λ can be obtained from that for μ by taking some of the pairs $(\cdots)$ and reversing them to get $)\cdots($.

With this definition, we can describe the connection to cover-expansive Dyck tilings. The following proposition is not hard to prove, and we leave it as an exercise for the reader.

Proposition 5.13. *Suppose λ, μ are partitions. Then $\lambda \stackrel{\Omega}{\leftarrow} \mu$ if and only if $\lambda \supseteq \mu$ and there is a cover-expansive Dyck tiling of $\lambda \setminus \mu$.*

Now we can describe the relationship between our results and those of Kenyon and Wilson. Suppose f, g are positive integers as in Section 5.1, and define matrices N, P with rows and columns indexed by $\mathcal{P}_{f,g}$, and with

$$N_{\lambda\mu} = (-1)^{|\lambda|+|\mu|} e_{\lambda\mu}, \quad P_{\lambda\mu} = i_{\lambda\mu}.$$

Then Theorems 5.2 and 5.4 show that N and P are mutual inverses. In fact, Kenyon and Wilson prove this result, but in a slightly different form. They define a matrix M with $M_{\lambda\mu} = 1$ if $\lambda \stackrel{\Omega}{\leftarrow} \mu$

and 0 otherwise, and show (using a different recursion for $i_{\lambda\mu}$ from ours) that $(M^{-1})_{\lambda\mu} = (-1)^{|\lambda|+|\mu|} i_{\lambda\mu}$ [KW, Theorem 1.5]. In view of Proposition 5.13, we have $M_{\lambda\mu} = e_{\lambda\mu}$, and so the main result of [KW] shows that Theorems 5.2 and 5.4 are equivalent. (There is another minor difference with our results, in that instead of the set $\mathcal{P}_{f,g}$ of partitions lying inside a rectangle, they consider the set of partitions lying inside a partition of the form $(c, c-1, \dots, 1)$; however, one can easily show the equivalence of the two results by taking very large bounding partitions and appropriate submatrices.)

Of course, using this result we could remove a lot of the work in this paper, but we prefer to keep our proofs, thereby keeping our paper self-contained and providing a new (albeit longer) proof of the main result of [KW].


We remark that Kim [Ki, §7] observes a q -analogue of the above result (for which he also credits Konvalinka); in this, the integer $i_{\lambda\mu}$ is replaced with the polynomial

$$i_{\lambda\mu}(q) = \sum_{T \in \mathcal{T}(\lambda, \mu)} q^{|T|},$$

where $|T|$ denotes the number of tiles in the tiling T , and similarly for $e_{\lambda\mu}$. It should be possible to give a new proof of this result using our techniques: in our bijective results (Propositions 3.3–3.8 and 4.6) it is easy to keep track of how the number of tiles in a tiling changes, so q -analogues of our recursive results can easily be proved. Replacing the group algebra of the symmetric group with the Hecke algebra of type A should enable q -analogues of Theorems 5.2 and 5.4 to be proved, yielding a new proof of Kim’s and Konvalinka’s result. We leave the details to the proverbial interested reader.

Another result from the literature relates to Theorem 5.10. To describe this relationship, we give an alternative characterisation of the function F , for which we need another definition. Given a partition λ , construct the parenthesis sequence of λ and divide it into pairs as above, and define a partial order on the set of pairs by saying that the pair $(\dots)$ is larger than the pair $(\dots)$ if the former is nested inside the latter: $(\dots (\overset{a}{\dots}) \overset{a}{\dots})$ $\overset{b}{\dots} \overset{b}{\dots}$. Call this partially ordered set $\mathcal{P}(\lambda)$.

Example. Continuing the last example and retaining the labelling for the pairs in the parenthesis sequence for λ , we have the following Hasse diagram for $\mathcal{P}(\lambda)$.



Now we have the following.


Proposition 5.14. *Suppose λ is a partition. Then $F(\lambda)$ is the number of linear extensions of the poset $\mathcal{P}(\lambda)$.*

A proof of this proposition was sketched by David Speyer in a response to the author's MathOverflow question [S]. In a comment to the same answer, Philippe Nadeau pointed out that this provides a non-recursive expression for $F(\lambda)$. To state this, we let $\mathcal{P}_N(\lambda)$ denote the poset obtained by taking the N largest elements of $\mathcal{P}(\lambda)$. Then for large enough N the number of linear extensions of $\mathcal{P}(\lambda)$ equals the number of linear extensions of $\mathcal{P}_N(\lambda)$. Given a pair $p \in \mathcal{P}(\lambda)$, define its *length* $l(p)$ to be 1 plus the number of intervening pairs. Now an easy exercise [Kn, p.70, Exercise 20] gives the following.

Proposition 5.15. *Suppose λ is a partition, and $N \gg 0$. Then the number of linear extensions of the poset $\mathcal{P}_N(\lambda)$ equals*

$$\frac{N!}{\prod_{p \in \mathcal{P}_N(\lambda)} l(p)}.$$

Example. Taking $N = 8$ in the last example and labelling each pair by its length, we have



So $F(\lambda)$ is the number of linear extensions of $\mathcal{P}_8(\lambda)$, which is

$$\frac{8!}{2 \times 3 \times 7 \times 8} = 120.$$

The author is very grateful to David Speyer and Philippe Nadeau for these comments (and also to Gjergji Zaimi for similar comments in the same thread), which inspired the introduction of Dyck tilings in the present work.

Given Theorem 5.4 and Proposition 5.14, Theorem 5.10 is equivalent to the following theorem.

Theorem 5.16. *Suppose λ is a partition and $N \gg 0$. Then*

$$\sum_{\mu \subseteq \lambda} i_{\lambda\mu} = \frac{N!}{\prod_{p \in \mathcal{P}_N(\lambda)} l(p)}.$$

This statement (in fact, a q -analogue) was conjectured by Kenyon and Wilson [KW, Conjecture 1]. It was proved inductively by Kim [Ki, Theorem 1.1], and then a bijective proof was given by Kim, Mészáros, Panova and Wilson [KMPW, Theorem 1.1]. We retain our proof of Theorem 5.10 to keep the paper self-contained (and so that Theorem 5.10 can be proved without reference to tilings), and we note in passing that this yields a new proof of Theorem 5.16.

6 The homogeneous Garnir relations

In this section we apply our earlier results to the study of representations of the (cyclo-tomic) quiver Hecke algebras (also known as *KLR algebras*) of type A , and in particular the homogeneous Garnir relations for the *universal graded Specht modules* of Kleshchev, Mathas and Ram [KMR].

6.1 The quiver Hecke algebra

Suppose $\mathbb{F}$ is a field of characteristic p , $e \in \{2, 3, 4, \dots\}$ and $n \in \mathbb{N}$. The *quiver Hecke algebra* $\mathcal{R}_n$ of type A is a unital associative $\mathbb{F}$ -algebra with a generating set

$$\{y_1, \dots, y_n\} \cup \{\psi_1, \dots, \psi_{n-1}\} \cup \{e(\mathbf{i}) \mid \mathbf{i} \in (\mathbb{Z}/e\mathbb{Z})^n\}$$

and a somewhat complicated set of defining relations, which may be found in [BK] or [KMR], for example. These relations allow one to write down a basis for $\mathcal{R}_n$: to do this, choose and fix a reduced expression $w = t_{i_1} \dots t_{i_r}$ for each $w \in \mathfrak{S}_n$, and set $\psi_w = \psi_{i_1} \dots \psi_{i_r}$. Then the set

$$\{\psi_w y_1^{c_1} \dots y_n^{c_n} e(\mathbf{i}) \mid w \in \mathfrak{S}_n, c_1, \dots, c_n \geq 0, \mathbf{i} \in (\mathbb{Z}/e\mathbb{Z})^n\}$$

is a basis for $\mathcal{R}_n$.

The quiver Hecke algebra has attracted considerable attention in recent years, thanks to the astonishing result of Brundan and Kleshchev [BK, Main Theorem] that when p does not divide e , a certain finite-dimensional ‘cyclotomic’ quotient of $\mathcal{R}_n$ is isomorphic to the cyclotomic Hecke algebra of type A (as introduced by Ariki–Koike [AK] and Broué–Malle [BM]) defined at a primitive e th root of unity in $\mathbb{F}$; when $e = p$, there is a corresponding isomorphism to the degenerate cyclotomic Hecke algebra (which includes the group algebra $\mathbb{F}\mathfrak{S}_n$ as a special case). This in particular shows that these Hecke algebras (which include the group algebra of the symmetric group) are non-trivially $\mathbb{Z}$ -graded, and has initiated the study of the graded representation theory of these algebras.

A crucial role in the representation theory of Hecke algebras is played by the *Specht modules*. Brundan, Kleshchev and Wang [BKW] showed how to work with the Specht modules in the quiver Hecke algebra setting, demonstrating in particular that these modules are graded. Kleshchev, Mathas and Ram [KMR] gave a presentation for each Specht module with a single generator and a set of homogeneous relations. These relations include ‘homogeneous Garnir relations’; although these are in general simpler than the classical Garnir relations for ungraded Specht modules (which go back to [G] in the symmetric group case), the expressions given for these relations in [KMR] are quite complicated. The purpose of this section is to use the results of the previous section to give a simpler expression for each homogeneous Garnir relation. In computations with graded Specht modules using the author’s GAP programs, implementing these simpler expressions has been observed to have some benefits in terms of computational efficiency.

We now define the cyclotomic quotients of the quiver Hecke algebras and their Specht modules. Choose a positive integer l , and an l -tuple $(\kappa_1, \dots, \kappa_l) \in (\mathbb{Z}/e\mathbb{Z})^l$. For each $a \in \mathbb{Z}/e\mathbb{Z}$ define $\kappa(a)$ to be the number of values of j for which $\kappa_j = a$, and define $\mathcal{R}_n^\kappa$ to be the quotient of $\mathcal{R}_n$ by the ideal generated by the elements $y_1^{\kappa(i_1)} e(\mathbf{i})$ for all $\mathbf{i} \in (\mathbb{Z}/e\mathbb{Z})^n$. We use the same notation for the standard generators of $\mathcal{R}_n$ and their images in $\mathcal{R}_n^\kappa$.

6.2 Row permutation modules and Specht modules

From now on, we stick to the case $l = 1$, which corresponds to the Iwahori–Hecke algebra of type A ; there is no essential difference in the homogeneous Garnir relations for arbitrary l , and we save on notation and terminology by restricting to this special case. We work only with ungraded modules, since the grading plays no part in our results.

In the case $l = 1$, the Specht modules (by which we mean the *row Specht modules* of [KMR, §5.4]) are labelled by partitions. For the sake of alignment with the rest of the literature, we now change to using the *English* convention for Young diagrams in this section, where the first coordinate increases down the page and the second increases from left to right.

Suppose π is a partition of n . If (a, b) is a node, its *residue* is defined to be $\text{res}((a, b)) = b - a + \kappa_1 \pmod{e}$. A π -*tableau* is a bijection from the Young diagram of π to the set $\{1, \dots, n\}$. We specify a particular π -tableau T^π by assigning the numbers $1, \dots, n$ to the nodes

$$(1, 1), (1, 2), \dots, (1, \pi_1), (2, 1), (2, 2), \dots, (2, \pi_2), (3, 1), (3, 2), \dots$$

in order. For example, when $\pi = (6, 4, 1^2)$ we have

$$T^\pi = \begin{array}{|c|c|c|c|c|c|} \hline 1 & 2 & 3 & 4 & 5 & 6 \\ \hline 7 & 8 & 9 & 10 & & \\ \hline 11 & & & & & \\ \hline 12 & & & & & \\ \hline \end{array}.$$

We define an element $\mathbf{i}^\pi$ of $(\mathbb{Z}/e\mathbb{Z})^n$ by setting $\mathbf{i}_j^\pi = \text{res}((T^\pi)^{-1}(j))$. For example, with $\pi = (6, 4, 1^2)$, $e = 4$ and $\kappa_1 = 0$, $\mathbf{i}^\pi$ is the sequence $(0, 1, 2, 3, 0, 1, 3, 0, 1, 2, 2, 1)$.

Now we can define the *row permutation module* M^π from [KMR, §5.3]. This has a single generator m^π , and relations as follows:

- $e(\mathbf{i}^\pi)m^\pi = m^\pi$;
- $y_j m^\pi = 0$ for all j ;
- $\psi_j m^\pi = 0$ whenever j and $j + 1$ lie in the same row of T^π .

It is easy to write down a basis for M^π . Say that a π -tableau T is *row-strict* if the entries increase from left to right along the rows. If T is row-strict, let w^T be the permutation taking T^μ to T , and define $\psi^T := \psi_{w^T}$, and $m^T = \psi^T m^\pi$. Then by [KMR, Theorem 5.6] the set $\{m^T \mid T \text{ a row-strict } \pi\text{-tableau}\}$ is a basis for M^π .

The Specht module S^π is the quotient of M^π by the homogeneous Garnir relations, which we now define. Say that a node $\mathbf{n} = (a, b)$ of π is a *Garnir node* if $(a + 1, b)$ is also a node of π . If $\mathbf{n}$ is a Garnir node, define the *Garnir belt* $\mathbf{B}^{\mathbf{n}}$ to be the set of nodes

$$\{(a, b), (a, b + 1), \dots, (a, \pi_a), (a + 1, 1), (a + 1, 2), \dots, (a + 1, b)\}.$$

The *Garnir tableau* $G^{\mathbf{n}}$ is the tableau defined by taking T^π and rearranging the entries within $\mathbf{B}^{\mathbf{n}}$ so that they increase from bottom left to top right. For example, with $\pi = (6, 4, 1^2)$ and $\mathbf{n} = (1, 3)$, the tableau $G^{\mathbf{n}}$ (with the Garnir belt $\mathbf{B}^{\mathbf{n}}$ shaded) is as follows.

$$G^{\mathbf{n}} = \begin{array}{|c|c|c|c|c|c|} \hline 1 & 2 & 6 & 7 & 8 & 9 \\ \hline 3 & 4 & 5 & 10 & & \\ \hline 11 & & & & & \\ \hline 12 & & & & & \\ \hline \end{array}$$

Now we define *bricks*: an n -brick is defined to be a set of e consecutive nodes in the same row of $\mathbf{B}^n$ of which the leftmost has the same residue as n . For example, take $\pi = (6, 4, 1^2)$ and $n = (1, 3)$ as above. If $e = 2$, then the bricks are

$$\{(1, 3), (1, 4)\}, \{(1, 5), (1, 6)\}, \{(2, 2), (2, 3)\},$$

while if $e = 3$ the bricks are

$$\{(1, 3), (1, 4), (1, 5)\}, \{(2, 1), (2, 2), (2, 3)\}.$$

Let f denote the number of bricks in row a and g the number of bricks in row $a + 1$, and set $k = f + g$. If $k > 0$, let $d = d^n$ be the smallest number which is contained in a brick in G^n , and for $1 \leq i < k$ define w_i to be the permutation

$$(d + ie - e, d + ie)(d + ie - e + 1, d + ie + 1) \dots (d + ie - 1, d + ie + e - 1),$$

Recalling the notation ψ_w for $w \in \mathfrak{S}_n$ from above, set $\sigma_i^n = \psi_{w_i}$; unlike the situation for arbitrary w , ψ_{w_i} does not depend on the choice of reduced expression for w_i . Also set $\tau_i^n = \sigma_i^n + 1$.

Remark. In the definition [KMR, (5.7)], the elements τ_i^n, σ_i^n include an idempotent $e(\mathbf{i}^n)$ as a factor in order to make them homogeneous; but this factor is unnecessary for our purposes, so we prefer the simpler version.

Now let T^n be the tableau obtained from G^n by re-ordering the bricks so that their entries increase along row a and then row $a + 1$.

Continuing the example above, we illustrate the tableau T^n for $e = 2, 3$, shading the Garnir belt and outlining the bricks:

$$T^n = \begin{array}{|c|c|c|c|c|c|} \hline 1 & 2 & 4 & 5 & 6 & 7 \\ \hline 3 & 8 & 9 & 10 & & \\ \hline 11 & & & & & \\ \hline 12 & & & & & \\ \hline \end{array} \quad \text{if } e = 2, \quad T^n = \begin{array}{|c|c|c|c|c|c|} \hline 1 & 2 & 3 & 4 & 5 & 9 \\ \hline 6 & 7 & 8 & 10 & & \\ \hline 11 & & & & & \\ \hline 12 & & & & & \\ \hline \end{array} \quad \text{if } e = 3.$$

Define ψ^{T^n} as above; as noted in [KMR, §5.4], ψ^{T^n} is independent of the choice of reduced expression for w_{T^n} .

Now we define elements of $\mathcal{R}_n^\kappa$ corresponding to the elements t_λ and s_λ of $\mathbb{F}\mathfrak{S}_k$ defined in Section 5. For $\lambda \in \mathcal{P}_{f,g}$, define

$$\tau_\lambda^n = \left(\tau_{\lambda_1+g}^n \tau_{\lambda_1+g+1}^n \dots \tau_{f+g-1}^n \right) \left(\tau_{\lambda_2+g-1}^n \tau_{\lambda_2+g}^n \dots \tau_{f+g-2}^n \right) \dots \left(\tau_{\lambda_g+1}^n \tau_{\lambda_g+2}^n \dots \tau_f^n \right),$$

and define σ_λ^n similarly. Then the elements $\sigma_\lambda^n \psi^{T^n}$ are precisely the elements ψ^T for row-strict tableaux T obtained from G^n by permuting the bricks.

Now we can define the homogeneous Garnir relations and the Specht module. Given a Garnir node n , the corresponding *Garnir element* is

$$g^n = \sum_{\lambda \in \mathcal{P}_{f,g}} \tau_\lambda^n \psi^{T^n}.$$

The Specht module S^π is defined to be the quotient of M^π by the relations $g^n m^\pi = 0$ for all Garnir nodes n of π .

6.3 Re-writing the Garnir relations

The difficulty with the definition of the Specht module given above, especially from a computational point of view, is that the individual terms $\tau_\lambda^n \psi^{T^n} m^\pi$ appearing in the Garnir relation are not readily expressed in terms of the standard basis $\{m^\pi \mid T \text{ row-strict}\}$ for M^π ; to express it in these terms, some quite involved reduction is required using the defining relations in $\mathcal{R}_n$.

Example. Take $\pi = (11, 5, 3, 1)$, $n = (1, 5)$ and $e = 3$. Then we have $f = 2$, $g = 1$, and

$$\sigma_1^n = \psi_9 \psi_8 \psi_7 \psi_{10} \psi_9 \psi_8 \psi_{11} \psi_{10} \psi_9, \quad \sigma_2^n = \psi_{12} \psi_{11} \psi_{10} \psi_{13} \psi_{12} \psi_{11} \psi_{14} \psi_{13} \psi_{12},$$

while $\psi^{T^n} = \psi_6 \psi_5 \psi_7 \psi_6 \psi_8 \psi_7 \psi_9 \psi_8 \psi_{10} \psi_9 \psi_{11} \psi_{10} \psi_{15} \psi_{14} \psi_{13} \psi_{12} \psi_{11}$. Hence the Garnir element g^n is

$$\begin{aligned} \left(\tau_\emptyset^n + \tau_{(1)}^n + \tau_{(2)}^n \right) \psi^{T^n} &= ((\sigma_1^n + 1)(\sigma_2^n + 1) + (\sigma_2^n + 1) + 1) \psi^{T^n} \\ &= (\sigma_1^n \sigma_2^n + \sigma_1^n + 2\sigma_2^n + 3) \psi^{T^n}. \end{aligned}$$

A non-trivial calculation using the defining relations for $\mathcal{R}_n$ and the fact that m^π is killed by ψ_5 , ψ_6 and ψ_7 shows that $\sigma_1^n \psi^{T^n} m^\pi = 0$. Hence

$$g^n m^\pi = \left(\sigma_\emptyset^n + 2\sigma_{(1)}^n + 3\sigma_{(2)}^n \right) \psi^{T^n} m^\pi = m^{G^n} + 2m^U + 3m^{T^n},$$

where

$$\begin{aligned} G^n &= \begin{array}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|} \hline 1 & 2 & 3 & 4 & 10 & 11 & 12 & 13 & 14 & 15 & 16 & & & & & \\ \hline 5 & 6 & 7 & 8 & 9 & & & & & & & & & & & \\ \hline 17 & 18 & 19 & & & & & & & & & & & & & \\ \hline 20 & & & & & & & & & & & & & & & \\ \hline \end{array}, \\ U &= \begin{array}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|} \hline 1 & 2 & 3 & 4 & 7 & 8 & 9 & 13 & 14 & 15 & 16 & & & & & \\ \hline 5 & 6 & 10 & 11 & 12 & & & & & & & & & & & \\ \hline 17 & 18 & 19 & & & & & & & & & & & & & \\ \hline 20 & & & & & & & & & & & & & & & \\ \hline \end{array}, \\ T^n &= \begin{array}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|} \hline 1 & 2 & 3 & 4 & 7 & 8 & 9 & 10 & 11 & 12 & 16 & & & & & \\ \hline 5 & 6 & 13 & 14 & 15 & & & & & & & & & & & \\ \hline 17 & 18 & 19 & & & & & & & & & & & & & \\ \hline 20 & & & & & & & & & & & & & & & \\ \hline \end{array}. \end{aligned}$$

We wish to generalise the above example, to re-write $g^n m^\pi$ for an arbitrary Garnir node n as a linear combination of elements $\sigma_\lambda^n \psi^{T^n} m^\pi$. Fortunately, all the computations that we need with the defining relations for $\mathcal{R}_n$ have already been done in [KMR].

Given a Garnir node n of π , define the *brick permutation space* $T^{\pi, n}$ to be the $\mathbb{F}$ -subspace of M^π spanned by all elements of the form $\sigma_{i_1}^n \dots \sigma_{i_s}^n \psi^{T^n} m^\pi$. Then the σ_i^n , and hence the τ_i^n , act on $T^{\pi, n}$, and we have the following.

Theorem 6.1 [KMR, Theorem 5.11]. *As operators on $T^{\pi, n}$, the elements $\tau_1^n, \dots, \tau_{k-1}^n$ satisfy the Coxeter relations for the symmetric group $\mathfrak{S}_k$, and hence $T^{\pi, n}$ can be considered as an $\mathbb{F}\mathfrak{S}_k$ -module. In fact, as an $\mathbb{F}\mathfrak{S}_n$ -module $T^{\pi, n}$ is isomorphic to the Young permutation module $\mathcal{M}^{(f, g)}$, with an isomorphism given by mapping $\psi^{T^n} m^\pi$ to the standard generator m .*

By Theorem 6.1 and the discussion at the start of Section 5, the sets

$$\left\{ \tau_\lambda^n \psi^{T^n} m^\pi \mid \lambda \in \mathcal{P}_{f,g} \right\} \quad \text{and} \quad \left\{ \sigma_\lambda^n \psi^{T^n} m^\pi \mid \lambda \in \mathcal{P}_{f,g} \right\}$$

are $\mathbb{F}$ -bases of $T^{\pi,n}$. Theorems 5.2 and 5.4 give the transition coefficients between these bases, while Theorem 5.10 gives the sum of the elements of the first basis in terms of the second basis. So we can immediately deduce the following, which is the main result of this section.

Theorem 6.2. *Suppose π is a partition and $\mathbf{n} = (a, b)$ is a Garnir node of π , and let f, g denote the numbers of bricks in rows a and $a + 1$ of $\mathbf{B}^n$ respectively. Define σ_λ^n and τ_λ^n for $\lambda \in \mathcal{P}_{f,g}$ as above, and define the modified Garnir element*

$$\hat{g}^n = \sum_{\lambda \in \mathcal{P}_{f,g}} F(\lambda) \sigma_\lambda^n \psi^{T^n}.$$

Then:

1. $\sigma_\mu^n \psi^{T^n} m^\pi = \sum_{\lambda \in \mathcal{P}_{f,g}} (-1)^{|\lambda|+|\mu|} e_{\lambda\mu} \tau_\lambda^n \psi^{T^n} m^\pi;$
2. $\tau_\mu^n \psi^{T^n} m^\pi = \sum_{\lambda \in \mathcal{P}_{f,g}} i_{\lambda\mu} \sigma_\lambda^n \psi^{T^n} m^\pi;$
3. $\hat{g}^n m^\pi = g^n m^\pi.$

Hence the Specht module S^π is the quotient of M^π by the relations $\hat{g}^n m^\pi = 0$ for all Garnir nodes $\mathbf{n}$ of π .

Remarks.

1. Statements (1) and (2) in the theorem are not necessary for the Garnir relations, but may be of interest for computation in the row permutation modules; indeed, the coefficients in (1) are essentially the coefficients c_w appearing in [KMR, Corollary 5.12].
2. We emphasise that the identities in Theorem 6.2 are not true if the terms m^π are omitted; in general, τ_μ^n is not a linear combination of σ_λ^n s (even if the idempotent $e(\mathbf{i}^n)$ from [KMR, (5.7)] is reinstated in the definition of σ_i^n and τ_i^n).

We end with an example which illustrates how the Garnir relations are used in Specht module computations. Say that a π -tableau is *standard* if the entries are increasing both along the rows and down the columns. If we let v^T denote the image of m^T in the Specht module, then by [KMR, Corollary 6.24] the set $\{v^T \mid T \text{ a standard } \pi\text{-tableau}\}$ is a basis for S^π . The Garnir relations allow one to write v^T as a linear combination of this basis when T is a row-strict, but not standard, tableau. The foundation for this is [KMR, Lemma 5.13] which does this in the case where $T = G^n$ for a Garnir node $\mathbf{n}$.

For our example, take $\pi = (8, 4)$ and $e = 2$. Let $\mathbf{n} = (1, 4)$. Then $f = g = 2$, and

$$G^n = \begin{array}{|c|c|c|c|c|c|c|c|} \hline 1 & 2 & 3 & 8 & 9 & 10 & 11 & 12 \\ \hline 4 & 5 & 6 & 7 & & & & \\ \hline \end{array}.$$

As with any Garnir node, G^n is row-strict but not standard. The modified Garnir element is

$$\hat{g}^n = \left(\sigma_\emptyset^n + 2\sigma_{(1)}^n + 3\sigma_{(2)}^n + 3\sigma_{(1^2)}^n + 6\sigma_{(2,1)}^n + 6\sigma_{(2^2)}^n \right) \psi^{T^n},$$

and the corresponding Garnir relation yields

$$v^{G^n} = -2v^U - 3v^V - 3v^W - 6v^X - 6v^Y,$$

where U, V, W, X, Y are standard tableaux that we describe as follows. We represent the elements $\psi^U, \psi^V, \psi^W, \psi^X, \psi^Y$ using the braid diagrams of Khovanov and Lauda [KL]; by using shaded squares, we illustrate (reverting to the Russian convention!) how the labelling partitions $\lambda \in \mathcal{P}_{2,2}$ arise.

$G^n = \begin{array}{ c c c c c c c c } \hline 1 & 2 & 3 & 8 & 9 & 10 & 11 & 12 \\ \hline 4 & 5 & 6 & 7 & & & & \\ \hline \end{array},$	$\psi^{G^n} = \sigma_{\emptyset}^n \psi^{T^n} =$	
$U = \begin{array}{ c c c c c c c c } \hline 1 & 2 & 3 & 6 & 7 & 10 & 11 & 12 \\ \hline 4 & 5 & 8 & 9 & & & & \\ \hline \end{array},$	$\psi^U = \sigma_{(1)}^n \psi^{T^n} =$	
$V = \begin{array}{ c c c c c c c c } \hline 1 & 2 & 3 & 6 & 7 & 8 & 9 & 12 \\ \hline 4 & 5 & 10 & 11 & & & & \\ \hline \end{array},$	$\psi^V = \sigma_{(2)}^n \psi^{T^n} =$	
$W = \begin{array}{ c c c c c c c c } \hline 1 & 2 & 3 & 4 & 5 & 10 & 11 & 12 \\ \hline 6 & 7 & 8 & 9 & & & & \\ \hline \end{array},$	$\psi^W = \sigma_{(1^2)}^n \psi^{T^n} =$	
$X = \begin{array}{ c c c c c c c c } \hline 1 & 2 & 3 & 4 & 5 & 8 & 9 & 12 \\ \hline 6 & 7 & 10 & 11 & & & & \\ \hline \end{array},$	$\psi^X = \sigma_{(2,1)}^n \psi^{T^n} =$	
$Y = \begin{array}{ c c c c c c c c } \hline 1 & 2 & 3 & 4 & 5 & 6 & 7 & 12 \\ \hline 8 & 9 & 10 & 11 & & & & \\ \hline \end{array},$	$\psi^Y = \sigma_{(2^2)}^n \psi^{T^n} =$	

References

- [AK] S. Ariki & K. Koike, ‘A Hecke algebra of $(\mathbb{Z}/r\mathbb{Z}) \wr \mathfrak{S}_n$ and construction of its irreducible representations’, *Adv. Math.* **106** (1994), 216–43. [cited on p. 32]
- [BM] M. Broué & G. Malle, ‘Zyklotomische Heckealgebren’, *Astérisque* **212** (1993), 119–89. [32]
- [BK] J. Brundan & A. Kleshchev, ‘Blocks of cyclotomic Hecke algebras and Khovanov–Lauda algebras’, *Invent. Math.* **178** (2009), 451–84. [32]
- [BKW] J. Brundan, A. Kleshchev & W. Wang, ‘Graded Specht modules’, *J. Reine Angew. Math.* **655** (2011), 61–87. [32]
- [G] H. Garnir, ‘Théorie de la représentation lineaire des groupes symétriques’, *Mém. Soc. Roy. Sci. Liège* **10** (1950), 100pp. [32]
- [KW] R. Kenyon & D. Wilson, ‘Double-dimer pairings and skew Young diagrams’, *Electron. J. Combin.* **18** (2011), Paper 130, 22pp. [4, 28, 29, 30, 31]
- [KL] M. Khovanov & A. Lauda, ‘A diagrammatic approach to categorification of quantum groups I’, *Represent. Theory* **13** (2009), 309–47. [37]
- [Ki] J. S. Kim, ‘Proofs of two conjectures of Kenyon and Wilson on Dyck tilings’, *J. Combin. Theory Ser. A* **119** (2012), 1692–1710. [30, 31]
- [KMPW] J. S. Kim, K. Mészáros, G. Panova & D. Wilson, ‘Dyck tilings, linear extensions, descents, and inversions’, arXiv:1205.6578. [31]
- [KMR] A. Kleshchev, A. Mathas & A. Ram, ‘Universal graded Specht modules for cyclotomic Hecke algebras’, *Proc. Lond. Math. Soc.* **105** (2012), 1245–89. [1, 32, 33, 34, 35, 36]
- [Kn] D. Knuth, *The art of computer programming. Volume 3: sorting and searching.*, Addison–Wesley Series in Computer Science and Information Processing. Addison–Wesley Publishing Co., Reading, Mass.–London–Don Mills, Ont., 1973. [31]
- [SZ] K. Shigechi & P. Zinn-Justin, ‘Path representation of maximal parabolic Kazhdan–Lusztig polynomials’, *J. Pure Appl. Algebra* **216** (2012), 2533–2548. [2]
- [S] D. Speyer, ‘A function from partitions to natural numbers – is it familiar?’, mathoverflow.net/q/132548. [31]