Journal of Combinatorial Theory, Series A 118 (2011) 1661–1673

Contents lists available at ScienceDirect

Journal of Combinatorial Theory, Series A

www.elsevier.com/locate/jcta

QSym over Sym has a stable basis

Aaron Lauve [a], Sarah K. Mason [b,1]

[a] *Department of Mathematics & Statistics, Loyola University Chicago, Chicago, IL 60660, United States*
[b] *Department of Mathematics, Wake Forest University, Winston-Salem, NC 27109, United States*

A R T I C L E I N F O

Article history:
Received 18 April 2010
Available online 16 February 2011

Keywords:
Quasisymmetric functions
Symmetric functions
Free modules
Inverting compositions

A B S T R A C T

We prove that the subset of quasisymmetric polynomials conjectured by Bergeron and Reutenauer to be a basis for the coinvariant space of quasisymmetric polynomials is indeed a basis. This provides the first constructive proof of the Garsia–Wallach result stating that quasisymmetric polynomials form a free module over symmetric polynomials and that the dimension of this module is $n!$.

1. Introduction

Quasisymmetric polynomials have held a special place in algebraic combinatorics since their introduction in [7]. They are the natural setting for many enumeration problems [16] as well as the development of Dehn–Somerville relations [1]. In addition, they are related in a natural way to Solomon's descent algebra of the symmetric group [14]. In this paper, we follow [2, Ch. 11] and view them through the lens of invariant theory. Specifically, we consider the relationship between the two subrings $\mathrm{Sym}_n \subseteq \mathrm{QSym}_n \subseteq \mathbb{Q}[\mathbf{x}]$ of symmetric and quasisymmetric polynomials in variables $\mathbf{x} = \mathbf{x}_n := \{x_1, x_2, \dots, x_n\}$. Let $\mathcal{E}_n$ denote the ideal in QSym_n generated by the elementary symmetric polynomials. In 2002, F. Bergeron and C. Reutenauer made a sequence of three successively finer conjectures concerning the quotient ring $\mathrm{QSym}_n/\mathcal{E}_n$. A.M. Garsia and N. Wallach were able to prove the first two in [6], but the third one remained open; we close it here (Corollary 10) with the help of a new basis for QSym_n introduced in [8].

E-mail addresses: lauve@math.luc.edu (A. Lauve), masonsk@wfu.edu (S.K. Mason).
URLs: http://www.math.luc.edu/~lauve (A. Lauve), http://www.wfu.edu/~masonsk (S.K. Mason).

[1] Mason was partially supported by NSF postdoctoral fellowship DMS 0603351.

0097-3165/$ – see front matter © 2011 Published by Elsevier Inc.
doi:10.1016/j.jcta.2011.01.008

1.1. Motivating context

Recall that Sym_n is the ring $\mathbb{Q}[\mathbf{x}]^{\mathfrak{S}_n}$ of invariant polynomials under the permutation action of $\mathfrak{S}_n$ on $\mathbf{x}$ and $\mathbb{Q}[\mathbf{x}]$. One of the crowning results in the invariant theory of $\mathfrak{S}_n$ is that the following statements hold:

(S1) $\mathbb{Q}[\mathbf{x}]^{\mathfrak{S}_n}$ is a polynomial ring, generated, say, by the elementary symmetric polynomials $\mathcal{E}_n = \{e_1(\mathbf{x}), \ldots, e_n(\mathbf{x})\}$;
(S2) the ring $\mathbb{Q}[\mathbf{x}]$ is a free $\mathbb{Q}[\mathbf{x}]^{\mathfrak{S}_n}$-module;
(S3) the *coinvariant space* $\mathbb{Q}[\mathbf{x}]_{\mathfrak{S}_n} = \mathbb{Q}[\mathbf{x}]/(\mathcal{E}_n)$ has dimension $n!$ and is isomorphic to the regular representation of $\mathfrak{S}_n$.

See [11, §§17, 18] for details. Analogous statements hold on replacing $\mathfrak{S}_n$ by any pseudo-reflection group. Since all spaces in question are graded, we may add a fourth item to the list: the *Hilbert series* $H_q(\mathbb{Q}[\mathbf{x}]_{\mathfrak{S}_n}) = \sum_{k\geqslant 0} d_k q^k$, where d_k records the dimension of the kth graded component of $\mathbb{Q}[\mathbf{x}]_{\mathfrak{S}_n}$, satisfies

(S4) $H_q(\mathbb{Q}[\mathbf{x}]_{\mathfrak{S}_n}) = H_q(\mathbb{Q}[\mathbf{x}])/H_q(\mathbb{Q}[\mathbf{x}]^{\mathfrak{S}_n})$.

Before we formulate the conjectures of Bergeron and Reutenauer, we recall another page in the story of Sym_n and the quotient space $\mathbb{Q}[\mathbf{x}]/(\mathcal{E}_n)$. The ring homomorphism ζ from $\mathbb{Q}[\mathbf{x}_{n+1}]$ to $\mathbb{Q}[\mathbf{x}_n]$ induced by the mapping $x_{n+1} \mapsto 0$ respects the rings of invariants (that is, $\zeta : \mathrm{Sym}_{n+1} \twoheadrightarrow \mathrm{Sym}_n$ is a ring homomorphism). Moreover, ζ respects the fundamental bases of monomial (m_λ) and Schur (s_λ) symmetric polynomials of Sym_n, indexed by partitions λ with at most n parts. For example,

$$\zeta\big(m_\lambda(\mathbf{x}_{n+1})\big) = \begin{cases} m_\lambda(\mathbf{x}_n), & \text{if } \lambda \text{ has at most } n \text{ parts,} \\ 0, & \text{otherwise.} \end{cases}$$

The stability of these bases plays a crucial role in representation theory [13]. Likewise, the associated stability of bases for the coinvariant spaces (e.g., of Schubert polynomials [4,12,15]) plays a role in the cohomology theory of flag varieties.

1.2. Bergeron–Reutenauer context

Given that QSym_n is a polynomial ring [14] containing Sym_n, one might ask, by analogy with $\mathbb{Q}[\mathbf{x}]$, how QSym_n looks as a module over Sym_n. This was the question investigated by Bergeron and Reutenauer [3]. (See also [2, §11.2].) They began by computing the quotient $P_n(q) := H_q(\mathrm{QSym}_n)/H_q(\mathrm{Sym}_n)$ by analogy with (S4). Surprisingly, the result was a polynomial in q with non-negative integer coefficients (so it could, conceivably, enumerate the graded space $\mathrm{QSym}_n/(\mathcal{E}_n)$). More astonishingly, sending q to 1 gave $P_n(1) = n!$. This led to the following two conjectures, subsequently proven in [6]:

(Q1) The ring QSym_n is a free module over Sym_n;
(Q2) The dimension of the "coinvariant space" $\mathrm{QSym}_n/(\mathcal{E}_n)$ is $n!$.

In their efforts to prove the conjectures above, Bergeron and Reutenauer introduced the notion of "pure and inverting" compositions B_n with at most n parts. These compositions have the favorable property of being n-stable in that $\mathsf{B}_n \subseteq \mathsf{B}_{n+1}$ and that $\mathsf{B}_{n+1} \setminus \mathsf{B}_n$ are the pure and inverting compositions with exactly $n+1$ parts. They were able to show that the pure and inverting "quasi-monomials" M_β (see Section 2) span $\mathrm{QSym}_n/(\mathcal{E}_n)$ for small n case by case (and that they are $n!$ in number), but the general result remained open. Their final conjecture, which we prove in Corollary 10, is as follows:

(Q3) The set of quasi-monomials $\{M_\beta\colon\ \beta \in \mathsf{B}_n\}$ is a basis for $\mathrm{QSym}_n/\mathcal{E}_n$.

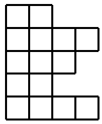

Fig. 1. The diagram associated to the composition $(2,4,3,2,4)$.

The balance of this paper is organized as follows. In Section 2, we recount the details surrounding a new basis $\{\mathcal{S}_\alpha\}$ for QSym_n called the quasisymmetric Schur polynomials. These behave particularly well with respect to the Sym_n action in the Schur basis. In Section 3, we give further details surrounding the "coinvariant space" $\mathrm{QSym}_n/(\mathcal{E}_n)$. These include a bijection between compositions α and pairs (λ,β), with λ a partition and β a pure and inverting composition, that informs our main results. Section 4 contains these results—a proof of (Q3), but with the quasi-monomials M_β replaced by the quasisymmetric Schur polynomials $\mathcal{S}_\beta$. We conclude in Section 5 with some corollaries to the proof. These include (Q3) as originally stated, as well as a version of (Q1) and (Q3) over the integers.

2. Quasisymmetric polynomials

A *composition* of n is a sequence of positive integers summing to n. A polynomial in n variables $\mathbf{x}=\{x_1,x_2,\ldots,x_n\}$ is said to be *quasisymmetric* if and only if for each composition $(\alpha_1,\alpha_2,\ldots,\alpha_k)$, the monomial $x_1^{a_1}x_2^{a_2}\cdots x_k^{a_k}$ has the same coefficient as $x_{i_1}^{\alpha_1}x_{i_2}^{\alpha_2}\cdots x_{i_k}^{\alpha_k}$ for all sequences $1\leqslant i_1<i_2<\cdots<i_k\leqslant n$. For example, $x_1^2x_2+x_1^2x_3+x_2^2x_3$ is a quasisymmetric polynomial in the variables $\{x_1,x_2,x_3\}$. The ring of quasisymmetric polynomials in n variables is denoted QSym_n. (Note that every symmetric polynomial is quasisymmetric.)

It is easy to see that QSym_n has a vector space basis given by the quasi-monomials

$$M_\alpha(\mathbf{x})=\sum_{i_1<\cdots<i_k} x_{i_1}^{\alpha_1}\cdots x_{i_k}^{\alpha_k},$$

for $\alpha=(\alpha_1,\ldots,\alpha_k)$ running over all compositions with at most n parts. It is also evident that QSym_n is a ring. See [10] for a formula for the product of two quasi-monomials. We write $\boldsymbol{l}(\alpha)=k$ for the *length* (number of parts) of α in what follows. We return to the quasi-monomial basis in Section 5, but for the majority of the paper, we focus on the basis of "quasisymmetric Schur polynomials" as its known multiplicative properties assist in our proofs.

2.1. The basis of quasisymmetric Schur polynomials

A *quasisymmetric Schur polynomial* $\mathcal{S}_\alpha$ is defined combinatorially through fillings of composition diagrams. Given a composition $\alpha=(\alpha_1,\alpha_2,\ldots,\alpha_k)$, its associated diagram is constructed by placing α_i boxes, or *cells*, in the ith row from the top. (See Fig. 1.) The cells are labeled using matrix notation; that is, the cell in the jth column of the ith row of the diagram is denoted (i,j). We abuse notation by writing α to refer to the diagram for α.

Given a composition diagram $\alpha=(\alpha_1,\alpha_2,\ldots,\alpha_\ell)$ with largest part m, a *composition tableau* T of *shape* α is a filling of the cells (i,j) of α with positive integers $T(i,j)$ such that

(CT1) entries in the rows of T weakly decrease when read from left to right,
(CT2) entries in the leftmost column of T strictly increase when read from top to bottom,
(CT3) entries satisfy the *triple rule*:
Let (i,k) and (j,k) be two cells in the same column so that $i<j$. If $\alpha_i\geqslant\alpha_j$ then either $T(j,k)<T(i,k)$ or $T(i,k-1)<T(j,k)$. If $\alpha_i<\alpha_j$ then either $T(j,k)<T(i,k)$ or $T(i,k)<T(j,k+1)$.

$$\begin{array}{lll} 1&1&1\\ 2&&\\ 3&3& \end{array}\quad \begin{array}{lll} 1&1&1\\ 2&&\\ 4&3& \end{array}\quad \begin{array}{lll} 1&1&1\\ 2&&\\ 4&4& \end{array}\quad \begin{array}{lll} 1&1&1\\ 3&&\\ 4&4& \end{array}\quad \begin{array}{lll} 2&1&1\\ 3&&\\ 4&4& \end{array}\quad \begin{array}{lll} 2&2&1\\ 3&&\\ 4&4& \end{array}\quad \begin{array}{lll} 2&2&2\\ 3&&\\ 4&4& \end{array}$$

Fig. 2. The composition tableaux encoded in the polynomial $\mathcal{S}_{(3,1,2)}(\mathbf{x}_4) = x_1^3x_2x_3^2 + x_1^3x_2x_3x_4 + x_1^3x_2x_4^2 + x_1^3x_3x_4^2 + x_1^2x_2x_3x_4^2 + x_1x_2^2x_3x_4^2 + x_2^3x_3x_4^2$.

Assign a *weight*, x^T to each composition tableau T by letting a_i be the number of times i appears in T and setting $x^T = \prod x_i^{a_i}$. The quasisymmetric Schur polynomial $\mathcal{S}_\alpha$ corresponding to the composition α is defined by

$$\mathcal{S}_\alpha(\mathbf{x}_n) = \sum_T x^T,$$

the sum being taken over all composition tableaux T of shape α with entries chosen from $[n]$. (See Fig. 2.) Each polynomial $\mathcal{S}_\alpha$ is quasisymmetric and the collection $\{\mathcal{S}_\alpha\colon \boldsymbol{l}(\alpha) \leqslant n\}$ forms a basis for QSym_n [8].

2.2. Sym action in the quasisymmetric Schur polynomial basis

We need several definitions in order to describe the multiplication rule for quasisymmetric Schur polynomials found in [9]. First, given two compositions $\alpha = (\alpha_1, \ldots, \alpha_r)$ and $\beta = (\beta_1, \ldots, \beta_s)$, we say α *contains* β ($\alpha \supseteq \beta$) if $r \geqslant s$ and there is a subsequence $i_1 > \cdots > i_s$ satisfying $\alpha_{i_1} \geqslant \beta_1, \ldots, \alpha_{i_s} \geqslant \beta_s$. The *reverse* of a partition λ is the composition λ^* obtained by reversing the order of its parts. Symbolically, if $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_k)$ then $\lambda^* = (\lambda_k, \ldots, \lambda_2, \lambda_1)$. Let β be a composition, let λ be a partition, and let α be a composition obtained by adding $|\lambda|$ cells to β, possibly between adjacent rows of β. A filling of the cells of α is called a *Littlewood–Richardson composition tableau* of shape $\alpha \supseteq \beta$ if it satisfies the following rules:

(LR1) The ith row from the bottom of β is filled with the entries $k+i$.
(LR2) The content of the appended cells is λ^*.
(LR3) The filling satisfies conditions (CT1) and (CT3) from Section 2.1.
(LR4) The entries in the appended cells, when read from top to bottom, column by column, from right to left, form a *reverse lattice word*. That is, each prefix contains at least as many i's as $(i-1)$'s for each $1 < i \leqslant k$.

The following theorem provides a method for multiplying an arbitrary quasisymmetric Schur polynomial by an arbitrary Schur polynomial.

Proposition 1. *(See [9].) In the expansion*

$$s_\lambda(\mathbf{x}) \cdot \mathcal{S}_\alpha(\mathbf{x}) = \sum_\gamma C_{\lambda\alpha}^{\gamma} \mathcal{S}_\gamma(\mathbf{x}), \tag{1}$$

the coefficient $C_{\lambda\alpha}^{\gamma}$ is the number of Littlewood–Richardson composition tableaux of shape $\gamma \supseteq \alpha$ with appended content λ^*.

3. The coinvariant space for quasisymmetric polynomials

Let $B \subseteq A$ be two $\mathbb{Q}$-algebras with A a free left module over B. This implies the existence of a subset $C \subseteq A$ with $A \simeq B \otimes C$ as vector spaces over $\mathbb{Q}$. In the classical setting of invariant theory (where B is the subring of invariants for some group action on A), this set C is identified as coset representatives for the quotient $A/(B_+)$, where (B_+) is the ideal in A generated by the positive part of the graded algebra $B = \bigoplus_{k \geqslant 0} B_k$.

Now suppose that A and B are graded rings. If A is free over B, then the Hilbert series of C is given as the quotient $H_q(A)/H_q(B)$. Let us try this with the choice $A = \mathrm{QSym}_n$ and $B = \mathrm{Sym}_n$. It is well known that the Hilbert series for QSym_n and Sym_n are given by

$$H_q(\mathrm{QSym}_n) = 1 + \frac{q}{1-q} + \cdots + \frac{q^n}{(1-q)^n} \tag{2}$$

and

$$H_q(\mathrm{Sym}_n) = \prod_{i=1}^{n} \frac{1}{1-q^i}. \tag{3}$$

Let $P_n(q) = \sum_{k \geqslant 0} p_k q^k$ denote the quotient of (2) by (3). It is easy to see that

$$P_n(q) = \prod_{i=1}^{n-1} \left(1 + q + \cdots + q^i\right) \sum_{i=0}^{n} q^i (1-q)^{n-i},$$

and hence $P_n(1) = n!$. It is only slightly more difficult (see (0.13) in [6]) to show that $P_n(q)$ satisfies the recurrence relation

$$P_n(q) = P_{n-1}(q) + q^n\left([n]_q! - P_{n-1}(q)\right), \tag{4}$$

where $[n]_q!$ is the standard q-version of $n!$. Bergeron and Reutenauer use this recurrence to show that p_k is a nonnegative integer for all $k \geqslant 0$ and to produce a set of compositions B_n satisfying $p_k = \#\{\beta \in \mathsf{B}_n \colon |\beta| = k\}$ for all n. In particular, $|\mathsf{B}_n| = n!$.

Let $(\mathcal{E}_n)$ be the ideal in QSym_n generated by all symmetric polynomials with zero constant term and call $R_n := \mathrm{QSym}_n/(\mathcal{E}_n)$ the *coinvariant space for quasisymmetric polynomials*. From the above discussion, R_n has dimension at most $n!$. If the set of quasi-monomials $\{M_\beta \in \mathrm{QSym}_n \colon \beta \in \mathsf{B}_n\}$ are linearly independent over Sym_n, then it has dimension exactly $n!$ and QSym_n becomes a free Sym_n-module of the same dimension.

3.1. Destandardization of permutations

To produce a set B_n of compositions indexing a proposed basis of R_n, first recognize the $[n]_q!$ in (4) as the Hilbert series for the classical coinvariant space $\mathbb{Q}[\mathbf{x}]/(\mathcal{E}_n)$ from (S3). The standard set of compositions indexing this space are the *Artin monomials* $\{x_1^{\alpha_1} \cdots x_n^{\alpha_n} \colon 0 \leqslant \alpha_i \leqslant n-i\}$, but these do not fit into the desired recurrence (4) with n-stability. In [5], Garsia developed an alternative set of monomials indexed by permutations. His "descent monomials" (actually, the "reversed" descent monomials, see [6, §6]) were chosen as the starting point for the recursive construction of the sets B_n. Here we give a description in terms of "destandardized permutations."

In what follows, we view partitions and compositions as words in the alphabet $\mathbb{N} = \{0, 1, 2, \ldots\}$. For example, we write 2543 for the composition $(2, 5, 4, 3)$. The *standardization* $\mathrm{st}(w)$ of a word w of length k is a permutation in $\mathfrak{S}_k$ obtained by first replacing (from left to right) the ℓ_1 1s in w with the numbers $1, \ldots, \ell_1$, then replacing (from left to right) the ℓ_2 2s in w with the numbers $\ell_1 + 1, \ldots, \ell_1 + \ell_2$, and so on. For example, $\mathrm{st}(121) = 132$ and $\mathrm{st}(2543) = 1432$. The *destandardization* $\mathbf{d}(\sigma)$ of a permutation $\sigma \in \mathfrak{S}_k$ is the lexicographically least word $w \in (\mathbb{N}_+)^k$ satisfying $\mathrm{st}(w) = \sigma$. For example, $\mathbf{d}(132) = 121$ and $\mathbf{d}(1432) = 1321$. Let $\mathsf{D}_{(n)}$ denote the compositions $\{\mathbf{d}(\sigma) \colon \sigma \in \mathfrak{S}_n\}$. Finally, given $\mathbf{d}(\sigma) = (\alpha_1, \ldots, \alpha_k)$, let $\mathbf{r}(\sigma)$ denote the vector difference $(\alpha_1, \ldots, \alpha_k) - (1^k)$ (leaving in place any zeros created in the process). For example, $\mathbf{r}(132) = 010$ and $\mathbf{r}(1432) = 0210$. Up to a relabeling, the weak compositions $\mathbf{r}(\sigma)$ are the ones introduced by Garsia in [5]. They are enumerated by $[n]_q!$ as follows. A *descent* in a permutation σ, written in one-line notation, is a position i where $\sigma_i > \sigma_{i+1}$. The *major index* $\mathrm{maj}(\sigma)$ records the sum of the positions i where a descent occurs within σ. It is well known (and readily verified recursively) that the coefficient of q^m in $[n]_q!$ is the number of permutations $\sigma \in \mathfrak{S}_n$ with $\mathrm{maj}(\sigma) = m$. Since the descent positions are preserved by the operators $\mathbf{d}$ and $\mathbf{r}$, the same statistics hold for $\mathsf{D}_{(n)}$ and their weak-composition counterparts.

$$
\begin{array}{ll}
\mathsf{D}_{(1)} = \{\underline{1}\} & \mathsf{B}_0 = \{0\} \\
\mathsf{D}_{(2)} = \{\underline{11}, 21\} & \mathsf{B}_1 = \{0\} \\
\mathsf{D}_{(3)} = \{\underline{111}, 211, 121, 221, 212, \underline{321}\} & \mathsf{B}_2 = \{0, 21\} \\
\mathsf{D}_{(4)} = \{\underline{1111}, 2111, 1211, 1121, 2211, 2121, 1221, 2112, 1212, 2221, 2212, 2122, & \mathsf{B}_3 = \{0, 21, 211, \\
\quad \underline{3211}, 3121, 1321, \underline{3221}, \underline{2321}, 3212, 2312, 2132, \underline{3321}, \underline{3231}, 3213, 4321\} & \quad 121, 221, 212\}
\end{array}
$$

Fig. 3. The sets $\mathsf{D}_{(n)}$ and B_n for small values of n. Compositions $1^n + \mathsf{B}_{n-1}$ are underlined in $\mathsf{D}_{(n)}$.

Bergeron and Reutenauer define their sets B_n recursively in such a way that

- $\mathsf{B}_0 := \{0\}$,
- $1^n + \mathsf{B}_{n-1} \subseteq \mathsf{D}_{(n)}$ and $\mathsf{D}_{(n)}$ is disjoint from B_{n-1}, and
- $\mathsf{B}_n := \mathsf{B}_{n-1} \cup \mathsf{D}_{(n)} \setminus (1^n + \mathsf{B}_{n-1})$.

Here, $1^n + \mathsf{B}_{n-1}$ denotes the vector sums $\{(1^n) + \mathbf{d}\colon\ \mathbf{d} \in \mathsf{B}_{n-1}\}$. Note that the compositions in $\mathsf{D}_{(n)}$ all have length n. Moreover, $1^{n+1} + \mathsf{D}_{(n)} \subseteq \mathsf{D}_{(n+1)}$. Indeed, if $\sigma = \sigma' 1$ is a permutation in $\mathfrak{S}_{n+1}$ with suffix "1" in one-line notation, then $(1^{n+1}) + \mathbf{d}(\mathrm{st}(\sigma')) = \mathbf{d}(\sigma)$. That (4) enumerates B_n is immediate [6, Proposition 6.1]. We give the first few sets B_n and $\mathsf{D}_{(n)}$ in Fig. 3.

3.2. Pure and inverting compositions

We now give an alternative description of the compositions in B_n introduced by Bergeron and Reutenauer which will be easier to work with in what follows. Call a composition α *inverting* if and only if for each $i > 1$ (with i less than or equal to the largest part of α) there exists a pair of indices $s < t$ such that $\alpha_s = i$ and $\alpha_t = i - 1$. For example, 13112312 is inverting while 21123113 is not. Any composition α admits a unique factorization

$$\alpha = \gamma k^{i_k} \cdots 2^{i_2} 1^{i_1} \quad (i_j \geqslant 1), \tag{5}$$

such that γ is a composition that does not contain any of the values from 1 to k, and k is maximal (but possibly zero). We say α is *pure* if and only if this maximal k is even. (Note that if the last part of a composition is not 1, then $k = 0$ and the composition is pure.) For example, 5435211 is pure with $k = 2$ while 3231 is impure since $k = 1$.

Proposition 2. *(See [3].) The set of inverting compositions of length n is precisely* $\mathsf{D}_{(n)}$. *The set of pure and inverting compositions of length at most n is precisely* B_n.

We reprise the proof of Bergeron and Reutenauer, for the sake of completeness.

Proof. Let $\mathcal{D}_{(n)}$ denote the set of inverting compositions of length n. The destandardization procedure makes it clear that $\mathsf{D}_{(n)} \subseteq \mathcal{D}_{(n)}$. For the reverse containment, we use induction on n to show that $|\mathcal{D}_{(n)}| = n!$. (The base case $n = 1$ is trivially satisfied.) Let $\alpha = (a_1, \ldots, a_{n-1})$ be one of the $(n-1)!$ compositions in $\mathcal{D}_{(n-1)}$. We construct n distinct compositions by inserting a new part between positions k and $k+1$ in α (for all $0 \leqslant k < n-1$). Define this part $m_k(\alpha)$ by

$$m_k(\alpha) = \max\bigl(\{a_i\colon\ i \leqslant k\} \cup \{1 + a_j\colon\ j > k\}\bigr).$$

To reverse the procedure, simply remove the rightmost maximal value appearing in the inverting composition of length n. Conclude that applying the procedure to $\mathcal{D}_{(n-1)}$ results in $n!$ distinct elements in $\mathcal{D}_{(n)}$. Finally, since the reverse map from $\mathcal{D}_{(n)}$ to $\mathcal{D}_{(n-1)}$ is an n to 1 map, we get that $|\mathcal{D}_{(n)}| = n!$.

Turning to B_n, we argue that $\mathsf{B}_n \cap \mathsf{D}_{(n)}$ are the pure compositions in $\mathsf{D}_{(n)}$ of length $n \geqslant 0$. This will complete the proof, since by construction and the previous paragraph, the compositions B_n are inverting. (Indeed, $\mathsf{B}_n \subseteq \bigcup_{0 \leqslant i \leqslant n} \mathsf{D}_{(i)}$, setting $\mathsf{D}_{(0)} = \{0\}$.) We argue by induction on n. (The base case $n = 0$ is trivially satisfied.) Note that if $\alpha \in \mathsf{D}_{(n)}$ is impure, then k is odd in the factorization (5), and $\alpha' := \alpha - (1^n)$ is pure. That is, $\alpha' \in \mathsf{B}_{n-1} \subseteq \mathsf{B}_n$. These are precisely the compositions eliminated from $\mathsf{D}_{(n)}$ in constructing B_n, for $\mathsf{B}_n := \mathsf{B}_{n-1} \cup \mathsf{D}_{(n)} \setminus (1^n + \mathsf{B}_{n-1})$. In other words, if $\alpha \in \mathsf{D}_{(n)}$ is pure, then $\alpha \in \mathsf{B}_n$. $\square$

$$\begin{array}{rccccccccccccc} \lambda = & 1 & 4 & 2 & 1 & 1 & 4 & 5 & 2 & 4 & 1 & 1 \\ \beta = & 2 & 4 & 3 & 1 & 1 & 3 & 4 & 2 & 3 & & \\ \hline \phi(\lambda,\beta) = & 3 & 8 & 5 & 2 & 2 & 7 & 9 & 4 & 7 & 1 & 1 \end{array}$$

Fig. 4. An example of the map $\phi : \mathsf{PB}_{13,49} \to \mathsf{C}_{13,49}$.

3.3. A bijection

Let $\mathsf{C}_{n,d}$ be the set of all compositions of d into at most n parts and set $\mathsf{PB}_{n,d} := \{(\lambda,\beta)$: λ a partition, $\beta \in \mathsf{B}_n$, $|\lambda| + |\beta| = d$, and $\boldsymbol{l}(\lambda) \leqslant n$, $\boldsymbol{l}(\beta) \leqslant n\}$. We define a map $\phi : \mathsf{PB}_{n,d} \to \mathsf{C}_{n,d}$ as follows.

Let (λ,β) be an arbitrary element of $\mathsf{PB}_{n,d}$. Then $\phi((\lambda,\beta))$ is the composition obtained by adding λ_i to the ith largest part of β for each $1 \leqslant i \leqslant \boldsymbol{l}(\lambda)$, where if $\beta_j = \beta_k$ and $j < k$, then β_j is considered smaller than β_k. If $\boldsymbol{l}(\lambda) > \boldsymbol{l}(\beta)$, append zeros after the last part to lengthen β before applying ϕ. (See Fig. 4.)

Proposition 3. *The map ϕ is a bijection between $\mathsf{PB}_{n,d}$ and $\mathsf{C}_{n,d}$.*

Proof. We prove this by describing the inverse ϕ^{-1} algorithmically. Let α be an arbitrary composition in $\mathsf{C}_{n,d}$ and set $(\lambda,\beta) := (\emptyset,\alpha)$.

(1) If β is pure and inverting, then $\phi^{-1}(\alpha) := (\lambda,\beta)$.
(2) If β is impure and inverting, then set $\phi^{-1}(\alpha) := (\lambda + (1^n), \beta - (1^n))$.
(3) If β is not inverting, then let j be the smallest part of β such that there does not exist a pair of indices $s < t$ such that $\beta_s = j$ and $\beta_t = j - 1$. Let m be the number of parts of β which are greater than or equal to j. Replace β with the composition obtained by subtracting 1 from each part greater than or equal to j and replace λ with the partition obtained by adding 1 to each of the first m parts.
(4) Repeat Steps (1)–(4) until ϕ^{-1} is obtained, that is, until Step (1) or (2) above is followed.

Notice that the composition $\beta - (1^n)$ in Step (2) is pure and inverting, since subtracting one from each part will change k from an even number to an odd number without affecting the inversions.

To see that $\phi\phi^{-1} = \mathbb{1}$, consider an arbitrary composition α. If α is pure and inverting, then $\phi\phi^{-1}(\alpha) = \phi(\emptyset,\alpha) = \alpha$. If α is impure and inverting, then $\phi(\phi^{-1}(\alpha)) = \phi(((1^{\boldsymbol{l}(\alpha)}), \alpha - (1^{\boldsymbol{l}(\alpha)}))) = \alpha$. Finally, consider a composition α which is not inverting. Note that the largest entry in α is decreased at each iteration of Step (3). Therefore the largest entry in the partition records the number of times the largest entry in α is decreased. Similarly, for each $i \leqslant \boldsymbol{l}(\lambda)$, the ith largest entry in α is decreased by one λ_i times. This means that the ith largest part of α is obtained by adding λ_i to the ith largest part of β and therefore $\phi\phi^{-1} = \mathbb{1}$.

To see that $\phi^{-1}\phi = \mathbb{1}$, consider an arbitrary pair (λ,β) such that β is a pure and inverting composition of length less than or equal to n, λ is a partition of length less than or equal to n, and $|\lambda| + |\beta| = d$. Apply the map ϕ to obtain a composition $\alpha = \phi(\lambda,\beta)$ of d of length less than or equal to n. Let ℓ be the length of λ and let m be the size of the least part of α which was modified during the procedure mapping (λ,β) to α. (Recall that if two parts are equal, the part to the right is considered to be larger.) Let k be the index of this part, so that $\alpha_k = m$. No part α_i with $i < k$ is equal to m (by construction) and no part α_j with $j > k$ is equal to $m - 1$ for otherwise $\beta_j \geqslant \beta_k$ and hence β_j would have been modified before β_k, a contradiction on the assumption that α_k is the smallest part of α which was modified during the map ϕ. Thus α violates the inverting condition at level m. The parts of α smaller than m do not violate the inverting condition since they appear as in β. Therefore the map ϕ^{-1} begins by subtracting one from each of the parts of α which are greater than or equal to m. Note that these are precisely the ℓ largest parts of α, since ℓ of the parts were modified and the smallest of the modified parts is α_k. In particular, any parts of α obtained from β by adding 1 during the map ϕ are returned to their initial values during this step.

The next step in ϕ^{-1} repeats the procedure described in the above paragraph replacing λ with $\lambda - (1^\ell)$. Therefore, the next step subtracts one from each of the parts of α which were modified by

$$\alpha \mapsto \begin{pmatrix} \lambda \\ \beta \end{pmatrix} : \begin{matrix} \emptyset \\ 3\,\underline{8}\,\underline{5}\,2\,2\,\underline{7}\,\underline{9}\,\underline{\underline{4}}\,\underline{7}\,1\,1 \end{matrix} \rightarrow \begin{matrix} 1\,1 \qquad 1\,1\,1\,1 \\ 3\,\underline{7}\,4\,2\,2\,\underline{\underline{6}}\,\underline{8}\,3\,\underline{\underline{6}}\,1\,1 \end{matrix}$$

$$\downarrow$$

$$\begin{matrix} 3\,1 \qquad 3\,3\,1\,3 \\ 3\,5\,4\,2\,2\,4\,\underline{\underline{6}}\,3\,4\,1\,1 \end{matrix} \leftarrow \begin{matrix} 2\,1 \qquad 2\,2\,1\,2 \\ 3\,\underline{6}\,4\,2\,2\,\underline{\underline{5}}\,\underline{7}\,3\,\underline{\underline{5}}\,1\,1 \end{matrix}$$

$$\downarrow$$

$$\begin{matrix} 3\,1 \qquad 3\,4\,1\,3 \\ 3\,5\,4\,2\,2\,4\,5\,3\,4\,1\,1 \end{matrix} \rightarrow \begin{matrix} 1\,4\,2\,1\,1\,4\,5\,2\,4\,1\,1 \\ 2\,4\,3\,1\,1\,3\,4\,2\,3 \end{matrix} \rightarrow \begin{pmatrix} 54442211111 \\ 243113423 \end{pmatrix}.$$

Fig. 5. The map $\phi^{-1} : \mathsf{C}_{13,49} \to \mathsf{PB}_{13,49}$ applied to $\alpha = 38522794711$. Parts j from Step (3) of the algorithm are marked with a double underscore.

the addition of a part λ_i of λ such that $\lambda_i > 1$. Since 1 was subtracted from each of these parts during the first step and 1 was subtracted from each of these parts during the second step, the end result after two steps is that 2 is subtracted from each part of α modified by the addition of a part of λ greater than or equal to 2. Continuing in this manner, each of the parts of λ are removed from the composition α until the original composition β is produced, together with the original partition λ. Therefore $\phi^{-1}\phi = \mathbb{1}$ so that the map ϕ is a bijection. $\square$

Fig. 5 illustrates the algorithmic description of ϕ^{-1} as introduced in the proof of Proposition 3 on $\alpha = 38522794711$.

4. Main theorem

Let B_n be as in Section 3 and set $\mathcal{B}_n := \{\mathcal{S}_\beta \colon \beta \in \mathsf{B}_n\}$. We prove the following.

Theorem 4. *The set $\mathcal{B}_n$ is a basis for the* Sym_n*-module R_n.*

To prove this, we analyze the quasisymmetric polynomials $\mathrm{QSym}_{n,d}$ in n variables of homogeneous degree d. Note that $\mathrm{QSym}_n = \bigoplus_{d \geqslant 0} \mathrm{QSym}_{n,d}$. Therefore, if $\mathfrak{C}_{n,d}$ is a basis for $\mathrm{QSym}_{n,d}$, then the collection $\bigcup_{d \geqslant 0} \mathfrak{C}_{n,d}$ is a basis for QSym_n. First, we introduce a useful term order.

4.1. The lexrev order

Each composition α can be rearranged to form a partition $\boldsymbol{\lambda}(\alpha)$ by arranging the parts in weakly decreasing order. Recall the *lexicographic* order $\geqslant_{\mathrm{lex}}$ on partitions of n, which states that $\lambda \geqslant_{\mathrm{lex}} \mu$ if and only if the first nonzero entry in $\lambda - \mu$ is positive. For two compositions α and γ of n, we say that α is larger then γ in *lexrev* order (written $\alpha \succcurlyeq \gamma$) if and only if either

- $\boldsymbol{\lambda}(\alpha) \geqslant_{\mathrm{lex}} \boldsymbol{\lambda}(\gamma)$, or
- $\boldsymbol{\lambda}(\alpha) = \boldsymbol{\lambda}(\gamma)$ and α is lexicographically larger than γ when reading right to left.

For instance, we have

$$4 \succcurlyeq 13 \succcurlyeq 31 \succcurlyeq 22 \succcurlyeq 112 \succcurlyeq 121 \succcurlyeq 211 \succcurlyeq 1111.$$

Remark. Extend lexrev to weak compositions of n of length at most n by padding the beginning of α or γ with zeros as necessary, so $\boldsymbol{l}(\alpha) = \boldsymbol{l}(\gamma) = n$. Viewing these as exponent vectors for monomials in $\mathbf{x}$ provides a term ordering on $\mathbb{Q}[\mathbf{x}]$. However, it is not good term ordering in the sense that it is not multiplicative: given exponent vectors α, β, and γ with $\alpha \succcurlyeq \gamma$, it is not necessarily the case that $\alpha + \beta \succcurlyeq \gamma + \beta$. This is likely the trouble encountered in [3] and [6] when trying to prove the Bergeron–Reutenauer conjecture (Q3). We circumvent this difficulty by working with the Schur polynomials s_λ and the quasisymmetric Schur polynomials $\mathcal{S}_\alpha$. We consider leading polynomials $\mathcal{S}_\gamma$ instead of leading monomials x^γ. The leading term $\mathcal{S}_\gamma$ in a product $s_\lambda \cdot \mathcal{S}_\alpha$ is readily found.

4.2. Proof of main theorem

We claim that the collection $\mathfrak{C}_{n,d} = \{s_\lambda \mathcal{S}_\beta\colon\ |\lambda| + |\beta| = d,\ \boldsymbol{l}(\lambda) \leqslant n,\ \boldsymbol{l}(\beta) \leqslant n,\ \text{and}\ \beta \in \mathsf{B}_n\}$ is a basis for $\mathrm{QSym}_{n,d}$, which in turn implies that $\mathcal{B}_n$ is a basis for R_n. To prove this, we make use of a special Littlewood–Richardson composition tableau called the *super filling*. Consider a composition β and a partition $\lambda = (\lambda_1, \ldots, \lambda_k)$. If $\boldsymbol{l}(\lambda) > \boldsymbol{l}(\beta)$ then append $\boldsymbol{l}(\lambda) - \boldsymbol{l}(\beta)$ zeros to the end of β. Fill the cells in the ith row from the bottom of β with the entries $k + i$. Append λ_i cells to the ith longest row of β. (If two rows of β have equal length, the lower of the rows is considered longer.) These new cells are then filled so that their entries have content λ^* as follows. Fill the new cells in the jth longest row with the entries λ_{k-j+1} unless two rows are of the same length. If two rows are the same length, fill the lower row with the lesser entries. The resulting filling is called the *super filling* $S(\lambda, \beta)$.

Proposition 5. *The super filling $S(\lambda, \beta)$ obtained from composition β and partition λ is a filling satisfying* (LR1)–(LR4)*.*

Proof. The super filling $S(\lambda, \beta)$ satisfies (LR1) and (LR2) by construction. We must prove that the filling also satisfies (LR3) and (LR4). Note that since $S(\lambda, \beta)$ satisfies (CT1) by construction, we need only prove that the entries in the filling satisfy the triple condition (CT3) and the lattice condition (LR4). In the following, let α be the shape of $S(\lambda, \beta)$.

To prove that the filling $S(\lambda, \beta)$ satisfies (CT3), consider an arbitrary pair of cells (i, k) and (j, k) in the same column. If $\alpha_i \geqslant \alpha_j$ then $\beta_i \geqslant \beta_j$, since the entries from λ are appended to the rows of β from largest row to smallest row. Therefore if (i, k) is a cell in the diagram of β then $T(j, k) < T(i, k) = T(i, k-1)$ regardless of whether or not (j, k) is in the diagram of β. If (i, k) is not in the diagram of β then (j, k) cannot be in the diagram of β since $\beta_i \geqslant \beta_j$. Therefore $T(j, k) < T(i, k)$ since the smaller entry is placed into the shorter row, or the lower row if the rows have equal length.

If $\alpha_i < \alpha_j$ then $\beta_i \leqslant \beta_j$. If $T(i, k) \leqslant T(j, k)$ then (i, k) is not in the diagram of β. If $(j, k+1)$ is in the diagram of β then $T(i, k) < T(j, k+1)$ since the entries in the diagram of β are larger than the appended entries. Otherwise the cell $(j, k+1)$ is filled with a larger entry than (i, k) since the longer rows are filled with larger entries and $\alpha_j > \alpha_i$. Therefore the entries in $S(\lambda, \beta)$ satisfy (CT3).

To see that the entries in $S(\lambda, \beta)$ satisfy (LR4), consider an entry i. We must show that an arbitrary prefix of the reading word contains at least as many i's as $(i-1)$'s. (Note that this is true when the prefix chosen is the entire reading word since $\lambda_i^* \geqslant \lambda_{i-1}^*$.) Let c_i be the rightmost column of $S(\lambda, \beta)$ containing the letter i and let c_{i-1} be the rightmost column of $S(\lambda, \beta)$ containing the letter $i-1$. Note that all entries not in the diagram of β in a given row are equal. If $c_i > c_{i-1}$ then every prefix will contain at least as many i's as $(i-1)$'s since there will always be at least one i appearing before any pairs i, $i-1$ in reading order. If $c_i = c_{i-1}$, then the entry i will appear in a higher row than the entry $i-1$ and hence will be read first for each column containing both an i and an $i-1$. Therefore the reading word is a reverse lattice word and hence the filling satisfies (LR4). $\square$

Proof of Theorem 4. Order the compositions of d into at most n parts by the lexrev order. To define the ordering on the elements of $\mathfrak{C}_{n,d}$, note that their indices are pairs of the form (λ, β), where λ is a partition of some $k \leqslant d$ and β is a composition of $d - k$ which lies in B_n. We claim that the leading term in the quasisymmetric Schur polynomial expansion of $s_\lambda \mathcal{S}_\beta$ is the polynomial $\mathcal{S}_{\phi(\lambda,\beta)}$. To see this, recall from Proposition 1 that the terms of $s_\lambda \mathcal{S}_\beta$ are given by Littlewood–Richardson composition tableaux of shape $\alpha \supseteq \beta$ and appended content λ^*, where α is an arbitrary composition shape obtained by appending $|\lambda|$ cells to the diagram of β so that conditions (CT1) and (CT3) are satisfied.

To form the largest possible composition (in lexrev order), one must first append as many cells as possible to the longest row of β, where again the lower of two equal rows is considered longer. The filling of this new longest row must end in an $L := \boldsymbol{l}(\lambda)$, since the reading word of the Littlewood–Richardson composition tableau must satisfy (LR4). No entry smaller than L can appear to the left of L in this row, since the row entries are weakly decreasing from left to right. This implies that the maximum possible number of entries that could be added to the longest row of β is λ_1. Similarly,

$$\begin{array}{r} \\ s_4 \\ s_{31} \\ s_1\cdot\mathcal{S}_{21} \\ s_{22} \\ s_{211} \\ \mathcal{S}_{121} \\ \mathcal{S}_{211} \end{array}
\begin{array}{c} \begin{array}{ccccccc} 4 & 13 & 31 & 22 & 112 & 121 & 211 \end{array} \\ \left(\begin{array}{ccccccc}
1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\
\cdot & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\
\cdot & \cdot & 1 & 1 & \cdot & \cdot & 1 \\
\cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot \\
\cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot \\
\cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot \\
\cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1
\end{array}\right)\end{array}$$

Fig. 6. The transition matrix for $n=3$, $d=4$.

the maximum possible number of entries that can be added to the second longest row of β is λ_2 and so on. If $\boldsymbol{l}(\lambda) > \boldsymbol{l}(\beta)$, append the extra parts of λ (from least to greatest, top to bottom) after the bottom row of β. The resulting shape is precisely the shape of $S(\lambda,\beta)$ which is equal to $\phi(\lambda,\beta)$ since β is a pure and inverting composition. Therefore there is at least one Littlewood–Richardson composition tableau of shape $\phi(\lambda,\beta)$ since $S(\lambda,\beta)$ is a Littlewood–Richardson composition tableau by Proposition 5.

The shape of the Littlewood–Richardson composition tableau $S(\lambda,\beta)$ corresponds to the largest composition appearing as an index of a quasisymmetric Schur polynomial in the expansion of $s_\lambda\mathcal{S}_\beta$, implying that $\mathcal{S}_{\phi(\lambda,\beta)}$ is indeed the leading term in this expansion. Since ϕ is a bijection, the entries in $\mathfrak{C}_{n,d}$ span $\mathrm{QSym}_{n,d}$ and are linearly independent. Therefore $\mathfrak{C}_{n,d}$ is a basis for $\mathrm{QSym}_{n,d}$ and hence $\mathcal{B}_n$ is a basis for the Sym_n-module R_n. □

Remark 6. Note that in the proof of Theorem 4, the entries appearing in the filling of shape $\phi(\lambda,\alpha)$ are uniquely determined by the lattice condition (LR4). This implies that $C_{\lambda,\alpha}^{\phi(\lambda,\alpha)}=1$. This fact allows us to work over $\mathbb{Z}$, a slightly more general setting than working over $\mathbb{Q}$. (See Section 5.3 for details.)

The transition matrix between the basis $\mathfrak{C}_{3,4}$ and the quasisymmetric Schur polynomial basis for $\mathrm{QSym}_{3,4}$ is given in Fig. 6.

5. Corollaries and applications

5.1. Closing the Bergeron–Reutenauer conjecture

The relationship between the monomial basis and quasisymmetric Schur basis was investigated in [8, Thm. 6.1 & Prop. 6.7]. We recall the pertinent facts.

Proposition 7. *(See [8].) The polynomials M_γ are related to the polynomials $\mathcal{S}_\alpha$ as follows:*

$$\mathcal{S}_\alpha=\sum_\gamma K_{\alpha,\gamma}M_\gamma, \tag{6}$$

where $K_{\alpha,\gamma}$ counts the number of composition tableaux T of shape α and content γ. Moreover, $K_{\alpha,\alpha}=1$ and $K_{\alpha,\gamma}=0$ whenever $\boldsymbol{\lambda}(\alpha) <_{\mathrm{lex}} \boldsymbol{\lambda}(\gamma)$.

We need a bit more to prove Conjecture (Q3).

Lemma 8. *In the notation of Proposition 7, $K_{\alpha,\gamma}=0$ whenever $\boldsymbol{\lambda}(\alpha)=\boldsymbol{\lambda}(\gamma)$ and $\alpha\neq\gamma$.*

Proof. We argue by induction on the largest part of α such that if $\boldsymbol{\lambda}(\alpha)=\boldsymbol{\lambda}(\gamma)$, and T is a composition tableau with shape α and content γ, then $\alpha=\gamma$.

The base case is trivial, for if the largest part of α is 1, then $\alpha=\gamma=(1^d)$ for some d. Now suppose α has largest part l. We claim that all rows i in T of length l must be filled only with i's. This claim finishes the proof. Indeed, we learn that $\alpha_i=\gamma_i$ for all such i. Thus we may apply the induction hypothesis to the new compositions α' and γ' obtained by deleting the largest parts from each.

To prove the claim, suppose row i of T has length l and is not filled with all i's. Let (i,k) be the rightmost cell in row i containing the entry i. The i in column $k+1$ must appear in a lower row, say row j, by condition (CT1) since the entries above row i in the first column must be less than i. This implies that $T(i,k)=T(j,k+1)$. But $T(j,k)\geqslant T(j,k+1)$ and hence $T(j,k)\geqslant T(i,k)$, so (CT3) is violated regardless of which row is longer. Therefore row i must be filled only with i's and the claim follows by induction. $\square$

Theorem 9. *In the expansion $M_\alpha=\sum_\gamma \tilde{K}_{\alpha,\gamma}\mathcal{S}_\gamma$, $\tilde{K}_{\alpha,\alpha}=1$ and $\tilde{K}_{\alpha,\gamma}=0$ whenever $\alpha\prec\gamma$.*

Proof. From Proposition 7 and Lemma 8, we learn that $K_{\alpha,\gamma}=0$ whenever $\alpha\prec\gamma$. (The proposition handles the first condition in the definition of the lexrev order and the lemma handles the second condition.) Now arrange the integers $K_{\alpha,\gamma}$ in a matrix K, ordering the rows and columns by $\succcurlyeq$. The previous observation shows that this change of basis matrix is upper-unitriangular. Consequently, the same holds true for $\tilde{K}=K^{-1}$. $\square$

We are ready to prove Conjecture (Q3). Let B_n and R_n be as in Section 4.

Corollary 10. *The set $\{M_\beta\colon\ \beta\in\mathsf{B}_n\}$ is a basis for the* Sym_n*-module R_n.*

Proof. We show that the collection $\mathfrak{M}_{n,d}=\{s_\lambda M_\beta\colon\ |\lambda|+|\beta|=d,\ \boldsymbol{l}(\lambda)\leqslant n,\ \boldsymbol{l}(\beta)\leqslant n,\ \text{and}\ \beta\in\mathsf{B}_n\}$ is a basis for $\mathrm{QSym}_{n,d}$, which in turn implies that $\{M_\beta\colon\ \beta\in\mathsf{B}_n\}$ is a basis for R_n. We first claim that the leading term in the quasisymmetric Schur polynomial expansion of $s_\lambda M_\beta$ is indexed by the composition $\phi(\lambda,\beta)$. The corollary will easily follow.

Applying Theorem 9, we may write $s_\lambda M_\beta$ as

$$s_\lambda M_\beta=s_\lambda\mathcal{S}_\beta+\sum_{\beta\succ\gamma}\tilde{K}_{\beta,\gamma}s_\lambda\mathcal{S}_\gamma.$$

Note that for any composition γ, the leading term of $s_\lambda S_\gamma$ is indexed by $\phi(\lambda,\gamma)$. This follows by the same reasoning used in the proof of Theorem 4. To prove the claim, it suffices to show that $\beta\succ\gamma\implies\phi(\lambda,\beta)\succ\phi(\lambda,\gamma)$.

Assume first that $\boldsymbol{\lambda}(\beta)=\boldsymbol{\lambda}(\gamma)$. Let i be the greatest integer such that $\beta_i>\gamma_i$. The map ϕ adds λ_j cells to β_i and λ_k cells to γ_i, where $\lambda_j\geqslant\lambda_k$. Therefore $\beta_i+\lambda_j>\gamma_i+\lambda_k$. Since the parts of $\phi(\lambda,\beta)$ and $\phi(\lambda,\gamma)$ are equal after part i, we have $\phi(\lambda,\beta)\succcurlyeq\phi(\lambda,\gamma)$.

Next assume that $\boldsymbol{\lambda}(\beta)\succ\boldsymbol{\lambda}(\gamma)$. Consider the smallest i such that the ith largest part β_j of β is not equal to the ith largest part γ_k of γ. The map ϕ adds λ_i cells to β_j and to γ_k, so that $\beta_j+\lambda_i>\gamma_k+\lambda_i$. Since the largest $i-1$ parts of $\phi(\lambda,\beta)$ and $\phi(\lambda,\gamma)$ are equal, we have $\boldsymbol{\lambda}(\phi(\lambda,\beta))\succ\boldsymbol{\lambda}(\phi(\lambda,\gamma))$.

We now use the claim to complete the proof. Following the proof of Theorem 4, we arrange the products $s_\lambda M_\beta$ as row vectors written in the basis of quasisymmetric Schur polynomials. The claim shows that the corresponding matrix is upper-unitriangular. Thus $\mathfrak{M}_{n,d}$ forms a basis for $\mathrm{QSym}_{n,d}$, as desired. $\square$

5.2. Triangularity

It was shown in Section 4 that the transition matrix between the bases $\mathfrak{C}$ and $\{\mathcal{S}_\alpha\}$ is triangular with respect to the lexrev ordering. Here, we show that a stronger condition holds: it is triangular with respect to a natural partial ordering on compositions. Every composition α has a corresponding partition $\boldsymbol{\lambda}(\alpha)$ obtained by arranging the parts of α in weakly decreasing order. A partition λ is said to *dominate* a partition μ iff $\sum_{i=1}^k\lambda_i\geqslant\sum_{i=1}^k\mu_i$ for all k. Let $C^\alpha_{\lambda,\beta}$ be the coefficient of $\mathcal{S}_\alpha$ in the expansion of the product $s_\lambda\mathcal{S}_\beta$.

Theorem 11. *If $\boldsymbol{\lambda}(\alpha)$ is not dominated by $\boldsymbol{\lambda}(\phi(\lambda,\beta))$, then $C^\alpha_{\lambda,\beta}=0$.*

Proof. Let (λ,β) be an arbitrary element of $\mathsf{PB}_{n,d}$ and let α be an arbitrary element of $\mathsf{C}_{n,d}$. Set $\gamma := \phi(\lambda,\beta)$. If $\gamma \preccurlyeq \alpha$ then $C^{\alpha}_{\lambda,\beta} = 0$ (by the proof of Theorem 4) and we are done.

Hence, assume that $\alpha \succ \phi(\lambda,\beta) = \gamma$ and that $\boldsymbol{\lambda}(\alpha)$ is not dominated by $\boldsymbol{\lambda}(\gamma)$. Let k be the smallest positive integer such that $\sum_{i=1}^{k} \boldsymbol{\lambda}(\alpha)_i > \sum_{i=1}^{k} \boldsymbol{\lambda}(\gamma)_i$. (Such an integer exists since $\boldsymbol{\lambda}(\alpha)$ is not dominated by $\boldsymbol{\lambda}(\gamma)$.) Therefore $\sum_{i=1}^{k} \boldsymbol{\lambda}(\alpha)_i - \sum_{i=1}^{k} \boldsymbol{\lambda}(\beta)_i > \sum_{i=1}^{k} \boldsymbol{\lambda}(\gamma)_i - \sum_{i=1}^{k} \boldsymbol{\lambda}(\beta)_i$ and there are more entries in the longest k rows of $\alpha \supseteq \beta$ then there are in the longest k rows of $\gamma \supseteq \beta$. This implies that there are more than $\sum_{i=1}^{k} \lambda_i$ entries from $\alpha \supseteq \beta$ contained in the longest k rows of α, since there are $\sum_{i=1}^{k} \lambda_i$ entries in the longest k rows of $\gamma \supseteq \beta$. This implies that in a Littlewood–Richardson composition tableau of shape $\alpha \supseteq \beta$, the longest k rows must contain an entry less than $L-k+1$ where $L = \boldsymbol{l}(\lambda)$.

The rightmost entry in the ith longest row of $\alpha \supseteq \beta$ must be $L-i+1$ for otherwise the filling would not satisfy the reverse lattice condition. This means that the longest k rows of α must contain only entries greater than or equal to $L-i+1$, which contradicts the assertion that an entry less than $L-k+1$ appears among the k longest rows of α. Therefore there is no such Littlewood–Richardson composition tableau of shape α and so $C^{\alpha}_{\lambda,\beta} = 0$. $\square$

5.3. *Integrality*

Up to this point, we have been working with the symmetric and quasisymmetric polynomials over the rational numbers, but their defining properties are equally valid over the integers. Briefly, bases for $\mathrm{Sym}_n(\mathbb{Z})$ and $\mathrm{QSym}_n(\mathbb{Z})$ are the Schur polynomials s_λ and the monomial quasisymmetric polynomials M_α, respectively. See [13] and [10] for details.

Lemma 12. *The polynomials $\{\mathcal{S}_\alpha \colon \boldsymbol{l}(\alpha) \leqslant n\}$ form a basis of $\mathrm{QSym}_n(\mathbb{Z})$.*

Proof. Lemma 8 states that the change of basis matrix K from $\{\mathcal{S}_\alpha\}$ to $\{M_\alpha\}$ is upper-unitriangular and integral. In particular, K is invertible over $\mathbb{Z}$, meaning that $\{\mathcal{S}_\alpha\}$ is a basis for $\mathrm{QSym}_n(\mathbb{Z})$. $\square$

One consequence of the proof of Theorem 4 is that $C^{\phi(\lambda,\beta)}_{\lambda,\beta} = 1$. (See Remark 6.) We exploit this fact below to prove stronger versions of Conjectures (Q1) and (Q3).

Corollary 13. *The algebra $\mathrm{QSym}_n(\mathbb{Z})$ is a free module over $\mathrm{Sym}_n(\mathbb{Z})$. A basis is given by $\{s_\lambda \mathcal{S}_\beta \colon \beta \in \Pi_n,\ \boldsymbol{l}(\lambda) \leqslant n$, and $\boldsymbol{l}(\beta) \leqslant n\}$. Replacing $\mathcal{S}_\beta$ by M_β results in an alternative basis.*

Proof. Theorem 4 combines with Proposition 1 (and the fact that $C^{\phi(\lambda,\beta)}_{\lambda,\beta} = 1$) to establish an upper-unitriangular, integral change of basis matrix C between $\{\mathcal{S}_\alpha \colon \boldsymbol{l}(\alpha) \leqslant n\}$ and $\{s_\lambda \mathcal{S}_\beta \colon \beta \in \Pi_n,\ \boldsymbol{l}(\lambda) \leqslant n$, and $\boldsymbol{l}(\beta) \leqslant n\}$. Since the former is an integral basis for $\mathrm{QSym}_n(\mathbb{Z})$, so is the latter. Composition of K, C and K^{-1} establishes the result for the monomial quasisymmetric polynomials. $\square$

Acknowledgments

We thank François Bergeron for sharing the story in Section 1.2 with us and encouraging us to write this paper. Our approach follows an idea that he proposed during CanaDAM 2009. We also gratefully acknowledge several beneficial discussions we had with Adriano Garsia, Christophe Reutenauer, and Frank Sottile.

References

[1] M. Aguiar, N. Bergeron, F. Sottile, Combinatorial Hopf algebras and generalized Dehn–Sommerville relations, Compos. Math. 142 (1) (2006) 1–30.
[2] F. Bergeron, Algebraic Combinatorics and Coinvariant Spaces, CMS Treatises Math., Canadian Mathematical Society, Ottawa, ON, 2009.
[3] F. Bergeron, C. Reutenauer, The coinvariant space for quasisymmetric polynomials, unpublished manuscript.

[4] S. Fomin, A.N. Kirillov, Combinatorial B_n-analogues of Schubert polynomials, Trans. Amer. Math. Soc. 348 (9) (1996) 3591–3620.
[5] A.M. Garsia, Combinatorial methods in the theory of Cohen-Macaulay rings, Adv. Math. 38 (3) (1980) 229–266.
[6] A.M. Garsia, N. Wallach, Qsym over Sym is free, J. Combin. Theory Ser. A 104 (2) (2003) 217–263.
[7] I.M. Gessel, Multipartite P-partitions and inner products of skew Schur functions, in: Combinatorics and algebra, Boulder, CO, 1983, in: Contemp. Math., vol. 34, Amer. Math. Soc., Providence, RI, 1984, pp. 289–317.
[8] J. Haglund, K. Luoto, S.K. Mason, S. van Willigenburg, Quasisymmetric Schur functions, J. Combin. Theory Ser. A 118 (2) (2011) 463–490.
[9] J. Haglund, K. Luoto, S.K. Mason, S. van Willigenburg, Refinements of the Littlewood–Richardson rule, Trans. Amer. Math. Soc. 363 (2011) 1665–1686.
[10] M. Hazewinkel, Symmetric functions, noncommutative symmetric functions, and quasisymmetric functions, Acta Appl. Math. 75 (1) (2003) 55–83.
[11] R. Kane, Reflection Groups and Invariant Theory, CMS Books Math./Ouvrages Math. SMC, vol. 5, Springer-Verlag, New York, 2001.
[12] A. Lascoux, M.-P. Schützenberger, Polynômes de Schubert, C. R. Acad. Sci. Paris Sér. I Math. 294 (13) (1982) 447–450.
[13] I.G. Macdonald, Symmetric Functions and Hall Polynomials, second ed., Oxford Math. Monogr., The Clarendon Press Oxford University Press, New York, 1995, with contributions by A. Zelevinsky, Oxford Science Publications.
[14] C. Malvenuto, C. Reutenauer, Duality between quasi-symmetric functions and the Solomon descent algebra, J. Algebra 177 (3) (1995) 967–982.
[15] L. Manivel, Symmetric Functions, Schubert Polynomials and Degeneracy Loci, SMF/AMS Texts Monogr., vol. 6, Amer. Math. Soc., Providence, RI, 2001, translated from the 1998 French original by John R. Swallow, Cours Spécialisés (Specialized Courses), 3.
[16] R.P. Stanley, Enumerative Combinatorics, vol. 2, Cambridge Stud. Adv. Math., vol. 62, Cambridge University Press, Cambridge, 1999.