

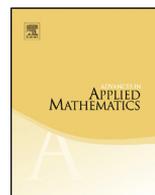


ELSEVIER

Contents lists available at ScienceDirect

Advances in Applied Mathematics

www.elsevier.com/locate/yaama



Chromatic symmetric functions in noncommuting variables revisited

Samantha Dahlberg^a, Stephanie van Willigenburg^{b,*}

^a School of Mathematical and Statistical Sciences, Arizona State University, Tempe AZ 85287-1804, USA

^b Department of Mathematics, University of British Columbia, Vancouver BC V6T 1Z2, Canada

ARTICLE INFO

Article history:

Received 13 May 2019

Received in revised form 13 August 2019

Accepted 5 September 2019

Available online xxxx

MSC:

05A18

05C15

05C25

05E05

16T30

Keywords:

Chromatic symmetric function

 e -positive

Schur-positive

Symmetric function in

noncommuting variables

ABSTRACT

In 1995 Stanley introduced a generalization of the chromatic polynomial of a graph G , called the chromatic symmetric function, X_G , which was generalized to noncommuting variables, Y_G , by Gebhard-Sagan in 2001. Recently there has been a renaissance in the study of X_G , in particular in classifying when X_G is a positive linear combination of elementary symmetric or Schur functions.

We extend this study from X_G to Y_G , including establishing the multiplicativity of Y_G , and showing Y_G satisfies the k -deletion property. Moreover, we completely classify when Y_G is a positive linear combination of elementary symmetric functions in noncommuting variables, and similarly for Schur functions in noncommuting variables, in the sense of Bergeron-Hohlweg-Rosas-Zabrocki. We further establish the natural multiplicative generalization of the fundamental theorem of symmetric functions, now in noncommuting variables, and obtain numerous new bases for this algebra whose generators are chromatic symmetric functions in noncommuting variables. Finally, we show that of all known symmetric func-

* Corresponding author.

E-mail addresses: sdahlber@asu.edu (S. Dahlberg), steph@math.ubc.ca (S. van Willigenburg).

tions in noncommuting variables, only all elementary and specified Schur ones can be realized as Y_G for some G .

© 2019 Elsevier Inc. All rights reserved.

1. Introduction

In 1912, Birkhoff introduced the chromatic polynomial of a graph G as a means to solve the four colour problem [5]. In 1995 Stanley generalized this to the chromatic symmetric function of G , X_G , [26], which not only generalized theorems of the chromatic polynomial such as the Broken Circuit Theorem [26], but also led to infinitely many new bases of the algebra of symmetric functions, Sym [8,11], a triple-deletion rule [22], and connections to representation theory and algebraic geometry [19]. However, much research on X_G has been devoted to proving the Stanley-Stembridge conjecture [28, Conjecture 5.5], which when studied in terms of X_G [26, Conjecture 5.1] says that if a poset is $(3 + 1)$ -free, then the chromatic symmetric function of its incomparability graph is a positive linear combination of elementary symmetric functions. The elementary symmetric functions arise in a variety of contexts, and one of the best known is the fundamental theorem of symmetric functions, which roughly states that Sym is generated by the elementary symmetric functions indexed by positive integers. While the $(3 + 1)$ -conjecture still stands, there has been much progress made towards it, for example, [7,10–13,15–18,20,30,32], and the related question of when X_G is a positive linear combination of Schur functions [14,23,25,27]. Most of these results have been achieved by working directly with X_G , however, there has been notable success in employing its generalization to quasisymmetric functions [7,25] and symmetric functions in noncommuting variables Y_G [10,16].

The algebra of symmetric functions in noncommuting variables, NCSym , was first studied by Wolf [31], who aimed to provide an analogue of the fundamental theorem of symmetric functions in this setting. However, little more was done in this area until Rosas and Sagan gave a systematic study of NCSym in 2004 [24], and also give a nice survey of this intervening work. They gave analogues in NCSym to well-known constructs in Sym such as Jacobi-Trudi determinants, and the RSK algorithm, and bases analogous to those in Sym with the exception of the basis of Schur functions. Such an analogue was found by Bergeron, Hohlweg, Rosas and Zabrocki who connected it to the Grothendieck bialgebra of the semi-tower of partition lattice algebras [2]. Bergeron, Reutenauer, Rosas and Zabrocki furthermore introduced a natural Hopf algebra structure on NCSym [3] and the antipode was subsequently computed by Lauve and Mastnak [21]. Bergeron and Zabrocki uncovered further algebraic structure by proving NCSym was free and cofree [4], and NCSym was also shown to be isomorphic to the algebra of rook placements by Can and Sagan [6]. Moreover, NCSym is connected with the supercharacter theory of all unipotent upper-triangular matrices over a finite field [1,29].

With the role of NCSym becoming ever more prominent, it is therefore befitting that the many new results regarding X_G be generalized to Y_G and that the question of Y_G being a positive linear combination of elementary symmetric functions or Schur functions in NCSym be answered. This paper achieves both of these goals, and uses Y_G to establish a number of new results regarding the algebraic structure of NCSym.

More precisely, this paper is structured as follows. In the next section we recall the relevant definitions. Then in Section 3 we prove the multiplicativity of Y_G in Proposition 3.5, generalizing Stanley’s [26, Proposition 2.3], and use it to establish a multiplicative version in NCSym of the fundamental theorem of symmetric functions in Theorem 4.12. We also give a formula for Y_G in terms of a Möbius function in Theorem 3.1, generalizing Stanley’s [26, Theorem 2.6]. We also show in Proposition 3.6 that Y_G exhibits the triple-deletion property of X_G proved by Orellana and Scott [22, Theorem 3.1] and its generalization, the k -deletion property, proved by the authors [11, Proposition 5]. In Section 4 we classify when Y_G is a positive (or negative) linear combination of elementary symmetric functions or Schur functions in NCSym in Theorems 4.14 and 4.8, respectively. In particular we show that Y_G is *always* a positive or negative linear combination in the latter case, and use this to show the same is true for all elementary symmetric functions in NCSym in Corollary 4.10. Lastly, in Section 5 we show that the Y_G generate new bases for NCSym in Theorem 5.2, generalizing the result of Cho and the second author [8, Theorem 5]. We conclude by showing that, with the exception of all elementary symmetric functions in NCSym and some specified Schur functions in NCSym, Y_G is never a currently known function in NCSym in Proposition 5.3, Theorem 5.4 and Theorem 5.6.

2. Background

We begin by giving necessary definitions and results that will be used throughout our paper.

Let n be a positive integer. Then we say an *integer partition* $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_{\ell(\lambda)})$ of n is an unordered list of positive integers whose sum is n , and denote this by $\lambda \vdash n$. We call the λ_i for $1 \leq i \leq \ell(\lambda)$ the *parts* of λ , call $\ell(\lambda)$ the *length* of λ and list the parts in weakly decreasing order. For example, $\lambda = (3, 2, 2, 1) \vdash 8$ and $\ell(\lambda) = 4$. We also write $\lambda = (1^{m_1}, 2^{m_2}, \dots, n^{m_n})$ to indicate that i appears in λ m_i times for $1 \leq i \leq n$. For example, our previous λ can be written as $\lambda = (1^1, 2^2, 3^1, 4^0, 5^0, 6^0, 7^0, 8^0)$. With this in mind, we define $\lambda! = \lambda_1! \lambda_2! \cdots \lambda_{\ell(\lambda)}!$ and $\lambda^! = m_1! m_2! \cdots m_n!$ similarly. For example, $(3, 2, 2, 1)! = 3!2!2!1! = 24$ and $(3, 2, 2, 1)^! = 1!2!1!0!0!0!0! = 2$.

Let $[n] = \{1, 2, \dots, n\}$. Then we say a *set partition* π of $[n]$ is a family of disjoint non-empty sets $B_1, B_2, \dots, B_{\ell(\pi)}$ whose union is $[n]$, and denote this by

$$\pi = B_1/B_2/\cdots/B_{\ell(\pi)} \vdash [n].$$

We call the B_i for $1 \leq i \leq \ell(\pi)$ the *blocks* of π , call $\ell(\pi)$ the *length* of π and list the blocks by increasing least element. For ease of notation we usually omit the set parentheses and commas of the blocks. For example, if π is the family of disjoint sets $\{1, 3, 4\}, \{2, 5\}, \{6\}, \{7, 8\}$ then we write

$$\pi = 134/25/6/78 \vdash [8]$$

and $\ell(\pi) = 4$. Note that every set partition $\pi \vdash [n]$ determines an integer partition $\lambda \vdash n$ by

$$\lambda(\pi) = \lambda(B_1/B_2/\dots/B_{\ell(\pi)}) = (|B_1|, |B_2|, \dots, |B_{\ell(\pi)}|)$$

listed in weakly decreasing order. For example, $\lambda(134/25/6/78) = (3, 2, 2, 1)$. For a finite set of integers, S , define $S + n = \{s + n : s \in S\}$. Then for two set partitions $\pi \vdash [n]$ and $\sigma = B_1/B_2/\dots/B_{\ell(\sigma)} \vdash [m]$ we define their *slash product* to be

$$\pi \mid \sigma = \pi/(B_1 + n)/(B_2 + n)/\dots/(B_{\ell(\sigma)} + n) \vdash [n + m].$$

Example 2.1. If $\pi = 134/25 \vdash [5]$ and $\sigma = 1/23 \vdash [3]$ then $\pi \mid \sigma = 134/25/6/78 \vdash [8]$.

We say a set partition π is *atomic* if there do not exist two non-empty set partitions σ_1, σ_2 such that $\pi = \sigma_1 \mid \sigma_2$. It is not hard to see that given any set partition π , it can be written uniquely as

$$\pi = \alpha_1 \mid \alpha_2 \mid \dots \mid \alpha_k$$

such that each α_i is non-empty and atomic, which we will call the *atomic decomposition* of π .

Example 2.2. The atomic decomposition of $\pi = 134/25/6/78 = 134/25 \mid 1 \mid 12$. Note that for ease of readability larger spaces have been intentionally inserted around the slash product symbols.

The set partitions of $[n]$ are the elements of the *partition lattice* Π_n ordered by *refinement*, namely for set partitions $\pi, \sigma \vdash [n]$ we have $\pi \leq \sigma$ if and only if every block of π is contained in some block of σ . For example, $134/25/6/78 < 1346/25/78$. The partition lattice Π_n has rank function $r(\pi) = n - \ell(\pi)$ and respective minimal and maximal elements

$$\begin{aligned} \hat{0}_n &= 1/2/\dots/n \\ \hat{1}_n &= 12 \dots n. \end{aligned}$$

Its Möbius function is known and satisfies

$$\mu_{\Pi}(\hat{0}_n, \hat{1}_n) = (-1)^{n-1}(n-1)! \tag{2.1}$$

and

$$\mu_{\Pi}(\hat{0}_n, \pi) = \prod_i (-1)^{\lambda_i-1}(\lambda_i-1)! \tag{2.2}$$

where $\lambda = \lambda(\pi)$.

We will now use set partitions to define the *algebra of symmetric functions in non-commuting variables* x_1, x_2, \dots

$$\text{NCSym} \subset \mathbb{Q} \ll x_1, x_2, \dots \gg$$

that is a graded algebra

$$\text{NCSym} = \text{NCSym}^0 \oplus \text{NCSym}^1 \oplus \dots$$

where $\text{NCSym}^0 = \text{span}\{1\}$. The n -th graded piece for $n \geq 1$ has the following bases

$$\begin{aligned} \text{NCSym}^n &= \text{span}\{m_{\pi} : \pi \vdash [n]\} = \text{span}\{p_{\pi} : \pi \vdash [n]\} \\ &= \text{span}\{e_{\pi} : \pi \vdash [n]\} = \text{span}\{\mathbf{x}_{\pi} : \pi \vdash [n]\} \end{aligned}$$

where these functions are defined as follows.

The *monomial symmetric function* in NCSym , m_{π} where $\pi \vdash [n]$, is given by

$$m_{\pi} = \sum_{(i_1, i_2, \dots, i_n)} x_{i_1} x_{i_2} \cdots x_{i_n}$$

summed over all tuples (i_1, i_2, \dots, i_n) with $i_j = i_k$ if and only if j and k are in the same block of π .

Example 2.3. $m_{13/2} = x_1 x_2 x_1 + x_2 x_1 x_2 + x_1 x_3 x_1 + x_3 x_1 x_3 + x_2 x_3 x_2 + x_3 x_2 x_3 + \dots$

Meanwhile, the *power sum symmetric function* in NCSym , p_{π} where $\pi \vdash [n]$, is given by

$$p_{\pi} = \sum_{(i_1, i_2, \dots, i_n)} x_{i_1} x_{i_2} \cdots x_{i_n}$$

summed over all tuples (i_1, i_2, \dots, i_n) with $i_j = i_k$ if j and k are in the same block of π .

Example 2.4. $p_{13/2} = x_1 x_2 x_1 + x_2 x_1 x_2 + \dots + x_1^3 + x_2^3 + \dots$

Similarly, the elementary symmetric function in NCSym, e_π where $\pi \vdash [n]$, is given by

$$e_\pi = \sum_{(i_1, i_2, \dots, i_n)} x_{i_1} x_{i_2} \cdots x_{i_n}$$

summed over all tuples (i_1, i_2, \dots, i_n) with $i_j \neq i_k$ if j and k are in the same block of π .

Example 2.5. $e_{13/2} = x_1 x_1 x_2 + x_1 x_2 x_2 + x_2 x_2 x_1 + x_2 x_1 x_1 + \cdots + x_1 x_2 x_3 + x_2 x_3 x_4 + \cdots$

Our final basis, for now, was defined by Bergeron, Hohlweg, Rosas and Zabrocki [2] to answer Rosas and Sagan’s question [24, Section 9] of whether there exists a basis for NCSym that reflects properties of Schur functions. Motivated by representation theory their Schur function in NCSym, \mathbf{x}_π where $\pi \vdash [n]$, is given by

$$\mathbf{x}_\pi = \sum_{\sigma \leq \pi} \mu_\Pi(\sigma, \pi) p_\sigma \tag{2.3}$$

with

$$p_\pi = \sum_{\sigma \leq \pi} \mathbf{x}_\sigma. \tag{2.4}$$

Example 2.6.

$$p_{13/2} = \mathbf{x}_{13/2} + \mathbf{x}_{1/2/3}$$

If we have a basis $\{b_\pi\}_{\pi \vdash [n], n \geq 1}$ of NCSym we say that $f = \sum_\pi c_\pi b_\pi$ is *b-positive* if and only if $c_\pi \geq 0$ for all π , and is *b-negative* if and only if $c_\pi \leq 0$ for all π . We also note that the names for our bases were chosen as the literature developed because of the projection map

$$\rho : \mathbb{Q} \ll x_1, x_2, \dots \gg \longrightarrow \mathbb{Q}[[x_1, x_2, \dots]]$$

that lets the variables commute, and in particular takes NCSym to Sym, with its analogous bases. In particular it takes the monomial symmetric functions, power sum symmetric functions and elementary symmetric functions in NCSym to their respective scalar multiples in Sym [24, Theorem 2.1].

The final concepts that we will need in order to define our objects of study come from graph theory. Let G be a graph with vertices $V(G)$ and edges $E(G)$ and for our purposes we will always assume that G is finite and simple. Furthermore, we will assume that G is labelled, namely if G has n vertices then the vertices are labelled distinctly with $1, 2, \dots, n$, which we term G with distinct vertex labels in $[n]$. Two types of graph that will be particularly useful to us are trees, which are connected graphs containing

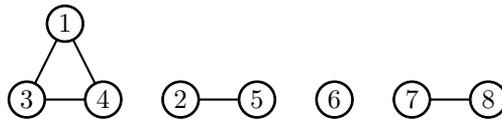
no cycles whose degree 1 vertices are each called a *leaf*, and the complete graphs K_n for $n \geq 1$ that consist of n vertices, every pair of which are adjacent.

Furthermore, given $\pi = B_1/B_2/\dots/B_{\ell(\pi)} \vdash [n]$ we define

$$K_\pi = K_{|B_1|} \cup K_{|B_2|} \cup \dots \cup K_{|B_{\ell(\pi)}|}$$

where \cup denotes disjoint union of graphs, and the labels on $K_{|B_i|}$ are the elements of the block B_i .

Example 2.7. $K_{134/25/6/78}$ is the following graph.



Two other tools that will be useful will be those of deletion and contraction. For the first of these, if G is a graph and $\epsilon \in E(G)$, then $G - \epsilon$ denotes G with ϵ deleted, and more generally if $S \subseteq E(G)$, then $G - S$ denotes G with every $\epsilon \in S$ deleted. Meanwhile for the second of these, if G and ϵ are as before, then G/ϵ denotes G with ϵ contracted and its vertices at either end identified. However, what will be the apex of our attention will be the chromatic symmetric function in NCSym, which we will define after one final notion. Given a graph G , we define a *proper colouring* κ of G to be a function

$$\kappa : V(G) \longrightarrow \{1, 2, \dots\}$$

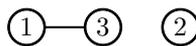
such that if $u, v \in V(G)$ are adjacent, then $\kappa(u) \neq \kappa(v)$.

Definition 2.8. [16, Definition 3.1] For a graph G with distinct vertex labels in $[n]$ let the label i be on vertex v_i . Then the *chromatic symmetric function* in NCSym is defined to be

$$Y_G = \sum_{\kappa} x_{\kappa(v_1)} x_{\kappa(v_2)} \cdots x_{\kappa(v_n)}$$

where the sum is over all proper colourings κ of G . If G is the empty graph then $Y_G = 1$.

Example 2.9. If $G = K_{13/2}$, then our graph is



and $Y_{K_{13/2}} = x_1 x_1 x_2 + x_1 x_2 x_2 + x_2 x_2 x_1 + x_2 x_1 x_1 + \dots + x_1 x_2 x_3 + x_2 x_3 x_4 + \dots = e_{13/2}$.

We conclude this section by recalling

$$\rho(Y_G) = X_G$$

where X_G is the chromatic symmetric function originally defined by Stanley [26] and motivated the definition and study of its analogue Y_G by Gebhard and Sagan [16].

3. New tools for chromatic symmetric functions in NCSym

In this section we generalize some of the classical and contemporary results for the chromatic symmetric function in Sym to that in NCSym. More precisely, given a graph G we derive Y_G in terms of power sum symmetric functions in NCSym whose coefficients are obtained from the Möbius function of the lattice of contractions of G , generalizing [26, Theorem 2.6]. We show that Y_G is multiplicative, generalizing [26, Proposition 2.3]. We also show that Y_G satisfies the properties of X_G known as triple-deletion [22, Theorem 3.1] and more generally k -deletion [11, Proposition 5].

For the first result, given a graph G with distinct vertex labels in $[n]$, we say $\pi = B_1/B_2/\dots/B_{\ell(\pi)} \vdash [n]$ is a *connected partition* of G if the graph induced by the vertices whose labels belong to any block B_i is connected. Let the *lattice of contractions*, L_G , of G be the set of all connected partitions of G partially ordered by refinement, and let μ_L be its Möbius function. Then we have the following generalization of [26, Theorem 2.6], whose statement and proof are analogous to the original.

Theorem 3.1. *For a graph G with distinct vertex labels in $[n]$ we have for any $\pi \in L_G$ that $\mu_L(\hat{0}_n, \pi) \neq 0$ and*

$$Y_G = \sum_{\pi \in L_G} \mu_L(\hat{0}_n, \pi) p_\pi.$$

Proof. By [26, Equation (1)] we have that $|\mu_L(\hat{0}_n, \pi)| > 0$ and hence $\mu_L(\hat{0}_n, \pi) \neq 0$. For a set partition $\pi \in L_G$ define

$$Y_\pi = \sum_{\kappa} x_\kappa$$

to be the sum over all special colourings, κ , of all the vertices that for $u, v \in V(G)$ is given by: (i) if u and v are in the same block of π then $\kappa(u) = \kappa(v)$, and (ii) if instead u and v are in different blocks and there is an edge between u and v , then $\kappa(u) \neq \kappa(v)$.

Note that *any* colouring κ of G contributes uniquely to one Y_π . We can see this by starting with any colouring κ and form each block of its partition π by colours, so that all vertices of the same colour are in the same block. Then we refine these blocks further to respect connected components, so that π is a connected partition of G .

Next using the definition of power sum symmetric functions in NCSym we have for $\sigma \in L_G$ that

$$p_\sigma = \sum_{\pi \geq \sigma, \pi \in L_G} Y_\pi.$$

By Möbius inversion we obtain the theorem, since when $\pi = \hat{0}_n$ the definition of special colouring coincides with that of proper colouring, so $Y_{\hat{0}_n} = Y_G$. \square

This is not the only expression for Y_G in terms of the power sum symmetric functions in NCSym, and it is the next one that will be useful for our second result. For it, given a graph G and edge set $S \subseteq E(G)$ we define $\pi(S)$ to be the set partition whose blocks are determined by the vertex labels of the connected components of G restricted to the edges in S . For example, in Example 2.7 if $S = \{(1, 3), (1, 4), (7, 8)\}$ then $\pi(S) = 134/2/5/6/78$.

Proposition 3.2. [16, Proposition 3.6] For a graph G we have

$$Y_G = \sum_{S \subseteq E(G)} (-1)^{|S|} p_{\pi(S)}.$$

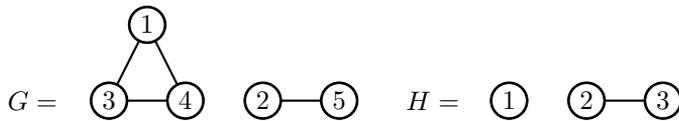
We also need the following result on the multiplicativity of the power sum symmetric functions in NCSym.

Lemma 3.3. [2, Lemma 4.1 (i)] For two set partitions π and σ we have

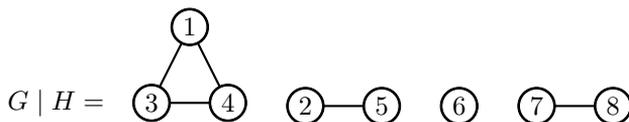
$$p_{\pi|\sigma} = p_\pi p_\sigma.$$

Lastly, given two graphs G , with distinct vertex labels in $[n]$, and H , with distinct vertex labels in $[m]$, define $G | H$ to be the disjoint union of G and H where the vertices corresponding to G have labels in $[n]$ in the same relative order as G , and the vertices corresponding to H have labels $\{n + 1, n + 2, \dots, n + m\}$ in the same relative order as H .

Example 3.4. If



then we get the following.



Then we have the following generalization of [26, Proposition 2.3].

Proposition 3.5. *For two graphs G , with distinct vertex labels in $[n]$, and H , with distinct vertex labels in $[m]$, we have*

$$Y_{G|H} = Y_G Y_H.$$

Proof. Using Proposition 3.2 we have

$$\begin{aligned} Y_{G|H} &= \sum_{S \subseteq E(G) \cup E(H)} (-1)^{|S|} p_{\pi(S)} \\ &= \sum_{S_1 \subseteq E(G), S_2 \subseteq E(H)} (-1)^{|S_1| + |S_2|} p_{\pi(S_1 \cup S_2)}. \end{aligned}$$

Note that because of the vertex labelling of $G | H$ we know that $\pi(S_1 \cup S_2) = \pi(S_1) | \pi(S_2)$. By Lemma 3.3 we have

$$\begin{aligned} Y_{G|H} &= \sum_{S_1 \subseteq E(G), S_2 \subseteq E(H)} (-1)^{|S_1| + |S_2|} p_{\pi(S_1) | \pi(S_2)} \\ &= \sum_{S_1 \subseteq E(G), S_2 \subseteq E(H)} (-1)^{|S_1| + |S_2|} p_{\pi(S_1)} p_{\pi(S_2)} \\ &= \sum_{S_1 \subseteq E(G)} (-1)^{|S_1|} p_{\pi(S_1)} \sum_{S_2 \subseteq E(H)} (-1)^{|S_2|} p_{\pi(S_2)} \\ &= Y_G Y_H. \quad \square \end{aligned}$$

Our final generalization is that the triple-deletion property [22, Theorem 3.1], and more generally the k -deletion property [11, Proposition 5], extend to chromatic symmetric functions in NCSym.

Proposition 3.6. *Let G be a graph such that $\epsilon_1, \epsilon_2, \dots, \epsilon_k \in E(G)$ form a k -cycle for $k \geq 3$. Then*

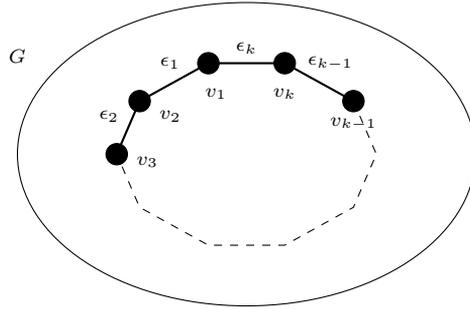
$$\sum_{S \subseteq [k-1]} (-1)^{|S|} Y_{G - \cup_{i \in S} \{\epsilon_i\}} = 0.$$

In particular let $k = 3$ and $\epsilon_1, \epsilon_2, \epsilon_3 \in E(G)$ form a triangle. Then

$$Y_G = Y_{G - \epsilon_1} + Y_{G - \epsilon_2} - Y_{G - \{\epsilon_1, \epsilon_2\}}.$$

Proof. We assume that G has a k -cycle for $k \geq 3$, with edges $\epsilon_1, \epsilon_2, \dots, \epsilon_k$ and vertices v_1, v_2, \dots, v_k as below.

We will prove the formula using a sign-reversing involution without any fixed points. Our signed set will be pairs (κ, S) where $S \subseteq [k - 1]$ and κ is a proper colouring on the



graph $G - \cup_{i \in S} \{\epsilon_i\}$. The weight of this pair will be $(-1)^{|S|} x_\kappa$, that is, the monomial associated to the colouring κ with sign determined by $|S|$. Since the edge ϵ_k is present in all graphs $G - \cup_{i \in S} \{\epsilon_i\}$ for any $S \subseteq [k - 1]$ all colourings will have at least two colours on the vertices v_1, v_2, \dots, v_k . Hence, there will always exist a smallest $j \in [k - 1]$ where the two colours on v_j and v_{j+1} are different. We map (κ, S) to (κ, T) where $T = S \cup \{j\}$ if $j \notin S$ and otherwise $T = S \setminus \{j\}$. The colouring κ is a proper colouring on $G - \cup_{i \in S} \{\epsilon_i\}$ and $G - \cup_{i \in T} \{\epsilon_i\}$ since we are only including or excluding the edge ϵ_j , which has adjacent vertices v_j and v_{j+1} with different colours. We can easily see that this is an involution that only switches the sign of the weight, and since there are no fixed points our desired sum is zero. \square

4. The classification of \mathbf{x} -positive and e -positive chromatic symmetric functions in NCSym

In 2001 Gebhard and Sagan observed that [16, Section 6]

Unfortunately, even for some of the simplest graphs, Y_G is usually not e -positive. The only graphs that are obviously e -positive are the complete graphs on n vertices and their complements.

In this section we confirm their observation by proving in Theorem 4.14 that Y_G is e -positive if and only if G is a union of complete graphs, and is never e -negative.

A closely related question is to classify when Y_G is \mathbf{x} -positive. Remarkably, it is always either \mathbf{x} -positive or \mathbf{x} -negative, depending on the number of vertices and connected components of G , as we shall see in Theorem 4.8. The elegant resolution to when Y_G is e -positive or \mathbf{x} -positive is in stark contrast to the analogous questions for X_G in Sym , which are still open and the subject of much research, as discussed in the introduction. We will first work towards proving when Y_G is \mathbf{x} -positive or \mathbf{x} -negative, and for this we require some tools.

The first tool is the Relabelling Proposition, which considers how, given a graph G , permuting the vertex labels of G affects Y_G . Given the symmetric group \mathfrak{S}_n , a permutation $\delta \in \mathfrak{S}_n$ and $f \in \text{NCSym}$, define $\delta \circ f$ to be the function after we permute the

placements of the variables, rather than the subscripts. For example, having $\delta = 213$ acting on $m_{1/23}$ means we switch the first two variables so $\delta \circ m_{1/23} = m_{13/2}$. Also define for a graph G with distinct vertex labels in $[n]$ a new graph $\delta(G)$, which is G but we permute the labels of the vertices subject to δ . Likewise, define for a set partition $\pi \vdash [n]$ a new set partition $\delta(\pi)$, which is π but we permute the block elements subject to δ . The following is due to Gebhard and Sagan.

Proposition 4.1. (Relabelling Proposition [16, Proposition 3.3]) *For a graph G with distinct vertex labels in $[n]$ and $\delta \in \mathfrak{S}_n$ we have*

$$Y_{\delta(G)} = \delta \circ Y_G.$$

Consequently, for any graph G with more than one component, using the Relabelling Proposition (Proposition 4.1) we can calculate Y_G using the connected components of G .

Corollary 4.2. *For any graph $G = G_1 \cup G_2$ a disjoint union of graphs G_1 and G_2 with distinct vertex labels in $[n]$, let $\delta \in \mathfrak{S}_n$ be a permutation such that $\delta(G) = G_1 \mid G_2$. Then we have*

$$Y_G = \delta^{-1} \circ (Y_{G_1} Y_{G_2}).$$

The second tool is Deletion-Contraction for Y_G by Gebhard and Sagan. Define the *induced* monomial to be

$$x_{i_1} x_{i_2} \cdots x_{i_{n-1}} \uparrow = x_{i_1} x_{i_2} \cdots x_{i_{n-1}}^2$$

where we make an extra copy of the last variable at the end and extend this definition linearly. Given a set partition $\pi \vdash [n - 1]$ we define $\pi \oplus n \vdash [n]$ to be the set partition where we place n in the same block as $n - 1$. For example, $14/23 \oplus 5 = 145/23$. Gebhard and Sagan [16, p. 233] state for $\pi \vdash [n - 1]$ that

$$m_\pi \uparrow = m_{\pi \oplus n} \text{ and } p_\pi \uparrow = p_{\pi \oplus n}. \tag{4.1}$$

Proposition 4.3. (Deletion-Contraction [16, Proposition 3.5]) *For a graph G with distinct vertex labels in $[n]$ and an edge ϵ between vertices labelled n and $n - 1$ we have*

$$Y_G = Y_{G-\epsilon} - Y_{G/\epsilon} \uparrow$$

where the new vertex remaining after the contraction of ϵ is labelled $n - 1$.

Before we show that all graphs G have Y_G be \mathbf{x} -positive or \mathbf{x} -negative we will show that all trees T have Y_G be \mathbf{x} -positive or \mathbf{x} -negative.

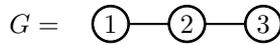
Lemma 4.4. For a tree T with distinct vertex labels in $[n]$ we have

$$Y_T = (-1)^{n-1} \sum_{\substack{\sigma \vdash [n], S_\sigma = E(T) \\ \text{no leaf is alone}}} \mathbf{x}_\sigma$$

where S_σ is the set of all edges in the path from any u to any v in $V(T)$ where u and v are in the same block of σ and “no leaf is alone” means that if vertex v is a leaf, then the block containing v has size at least two.

Before we prove this lemma we give a small but illustrative example.

Example 4.5. If



then $Y_G = (-1)^2(\mathbf{x}_{123} + \mathbf{x}_{13/2}) = \mathbf{x}_{123} + \mathbf{x}_{13/2}$.

Proof. Let T be a tree on n vertices with edge set E . Without loss of generality say T has a leaf at vertex n connected to vertex $n - 1$, by the Relabelling Proposition (Proposition 4.1). Call this edge ϵ and $\bar{E} = E \setminus \{\epsilon\}$. Using Deletion-Contraction (Proposition 4.3), Proposition 3.2 and Equation (2.4) we have

$$\begin{aligned} Y_T &= Y_{T-\epsilon} - Y_{T/\epsilon} \uparrow \\ &= \sum_{S \subseteq \bar{E}} (-1)^{|S|} p_{\pi(S)|1} - \sum_{S \subseteq \bar{E}} (-1)^{|S|} p_{\pi(S) \oplus n} \\ &= \sum_{S \subseteq \bar{E}} \sum_{\sigma \leq \pi(S)|1} (-1)^{|S|} \mathbf{x}_\sigma - \sum_{S \subseteq \bar{E}} \sum_{\sigma \leq \pi(S) \oplus n} (-1)^{|S|} \mathbf{x}_\sigma. \end{aligned}$$

Between these two sums of sums we cancel out all the \mathbf{x}_σ associated to $\sigma \leq \pi(S)|1$ in the first sum with the identical term in the second sum since certainly $\sigma \leq \pi(S) \oplus n$. We then have

$$Y_T = - \sum_{S \subseteq \bar{E}} \sum_{\substack{\sigma \leq \pi(S) \oplus n \\ n \text{ is not alone}}} (-1)^{|S|} \mathbf{x}_\sigma$$

where “ n is not alone” in $\sigma \vdash [n]$ means that the block containing n has at least two elements. Our next step will be to switch the order of summation. To do this, given a σ where n is not alone, we must identify all $S \subseteq \bar{E}$ that satisfy $\sigma \leq \pi(S) \oplus n$. Consider $S_\sigma \subseteq E$ the collection of all edges in the path between u and v for all u and v in the same block of σ .

First we will show that $\sigma \leq \pi(S_\sigma \setminus \{\epsilon\}) \oplus n$. We will show this by considering any a, b in the same block of σ and show a, b are in the same block of $\pi(S_\sigma \setminus \{\epsilon\}) \oplus n$. Say that

$a, b \neq n$ are in the same block of σ , we know that all edges in the unique path between a and b are in S_σ and this path avoids the edge ϵ . This guarantees that a and b are in the same block in $\pi(S_\sigma \setminus \{\epsilon\})$. Say that $a \neq n$ and $b = n$ are in the same block of σ . Then S_σ certainly contains ϵ . This guarantees that all edges in the unique path from n to $n - 1$ to a are in S_σ . This then guarantees that a and $n - 1$ are in the same block in $\pi(S_\sigma \setminus \{\epsilon\})$, which implies a and n are in the same block in $\pi(S_\sigma \setminus \{\epsilon\}) \oplus n$.

Next we will show that $S_\sigma \setminus \{\epsilon\} \subseteq S$ if and only if $\sigma \leq \pi(S) \oplus n$, which identifies all the $S \subseteq \bar{E}$ that we need. Certainly if $S_\sigma \setminus \{\epsilon\} \subseteq S$ we have $\sigma \leq \pi(S_\sigma \setminus \{\epsilon\}) \oplus n \leq \pi(S) \oplus n$. Say instead that $\sigma \leq \pi(S) \oplus n$ and $\bar{u}\bar{v} \neq \epsilon$ is an edge in S_σ , which we want to show is also in S . Because $\bar{u}\bar{v} \in S_\sigma$ we know that u and v are on some path P in T with end points a and b in the same block of σ . If $a, b \neq n$ then this means a and b are in the same block of $\pi(S)$, which only happens when all edges in the path P are in S , so $\bar{u}\bar{v} \in S$. If instead $a \neq n$ and $b = n$, then we still have a and n in the same block of $\pi(S) \oplus n$ meaning that $n - 1$ is in the same block as a similarly implying that all edges between a and $n - 1$, which includes $\bar{u}\bar{v}$, are in S . Switching the order of summation we have

$$Y_T = - \sum_{\substack{\sigma \\ n \text{ is not alone}}} \sum_{S_\sigma \setminus \{\epsilon\} \subseteq S \subseteq \bar{E}} (-1)^{|S|} \mathbf{x}_\sigma.$$

Note that unless $S_\sigma \setminus \{\epsilon\} = \bar{E}$ the inner sum is 0, so we have

$$Y_T = - \sum_{\substack{\sigma, S_\sigma \setminus \{\epsilon\} = \bar{E} \\ n \text{ is not alone}}} (-1)^{|\bar{E}|} \mathbf{x}_\sigma = (-1)^{n-1} \sum_{\substack{\sigma, S_\sigma \setminus \{\epsilon\} = \bar{E} \\ n \text{ is not alone}}} \mathbf{x}_\sigma = (-1)^{n-1} \sum_{\substack{\sigma, S_\sigma = E \\ n \text{ is not alone}}} \mathbf{x}_\sigma$$

after additionally using the fact that T is a tree so $|\bar{E}| = n - 2$. The lemma then follows by observing the same proof applies to every leaf. \square

We will now show that inducing an \mathbf{x} -positive function preserves \mathbf{x} -positivity. For this we need to recall the following.

Lemma 4.6. [2, Lemma 4.2 (i)] *For two set partitions π and σ we have*

$$\mathbf{x}_{\pi|\sigma} = \mathbf{x}_\pi \mathbf{x}_\sigma.$$

Also note that given a set partition $\pi \vdash [n]$ and $\delta \in \mathfrak{S}_n$, since $\delta \circ p_\pi = p_{\delta(\pi)}$ [16, p. 230] we also have that

$$\delta \circ \mathbf{x}_\pi = \mathbf{x}_{\delta(\pi)}. \tag{4.2}$$

Lemma 4.7. *We have*

$$\mathbf{x}_{\hat{1}_{n-1}} \uparrow = \sum_{\sigma} \mathbf{x}_\sigma$$

summed over all $\sigma \vdash [n]$ with at most two blocks where every block contains n or $n - 1$. Further, $\mathbf{x}_\pi \uparrow$ is \mathbf{x} -positive for all $\pi \vdash [n]$ and hence inducing any \mathbf{x} -positive function preserves \mathbf{x} -positivity.

Proof. By change of bases formulas in Equations (2.3) and (2.4) and the formula for inducing in (4.1) we have

$$\begin{aligned} \mathbf{x}_{\hat{1}_{n-1}} &= \sum_{\sigma \vdash [n-1]} \mu_\Pi(\sigma, \hat{1}_{n-1}) p_\sigma \\ \Rightarrow \mathbf{x}_{\hat{1}_{n-1}} \uparrow &= \sum_{\sigma \vdash [n-1]} \mu_\Pi(\sigma, \hat{1}_{n-1}) p_\sigma \uparrow = \sum_{\sigma \vdash [n-1]} \mu_\Pi(\sigma, \hat{1}_{n-1}) p_{\sigma \oplus n} \\ &= \sum_{\sigma \vdash [n-1]} \sum_{\pi \leq \sigma \oplus n} \mu_\Pi(\sigma, \hat{1}_{n-1}) \mathbf{x}_\pi. \end{aligned}$$

Next note that $\{\sigma \vdash [n-1] : \pi \leq \sigma \oplus n\} = \{\sigma \vdash [n-1] : \tilde{\pi} \leq \sigma \leq \hat{1}_{n-1}\}$ where $\tilde{\pi}$ is the set partition π where the blocks containing n and $n - 1$ are merged and n is removed. This is because for a fixed $\pi \vdash [n]$ the minimal set partition $\tilde{\pi} \geq \pi$ where $n - 1$ and n are in the same block is π but the blocks containing n and $n - 1$ are merged. Note that $\sigma \oplus n$ has n and $n - 1$ in the same block so because $\pi \leq \sigma \oplus n$ we have $\tilde{\pi} \leq \sigma \oplus n$ and removing the n for each set partition gives us $\tilde{\pi} \leq \sigma$. Conversely, given $\tilde{\pi} \leq \sigma$ we have $\pi \leq \tilde{\pi} = \tilde{\pi} \oplus n \leq \sigma \oplus n$. We then have

$$\mathbf{x}_{\hat{1}_{n-1}} \uparrow = \sum_{\pi \vdash [n]} \sum_{\tilde{\pi} \leq \sigma \leq \hat{1}_{n-1}} \mu_\Pi(\sigma, \hat{1}_{n-1}) \mathbf{x}_\pi.$$

By the definition of Möbius functions the inner sum is 0 unless $\tilde{\pi} = \hat{1}_{n-1}$, so we have finished proving the equation in this lemma since it is not hard to see all π with $\tilde{\pi} = \hat{1}_{n-1}$ is the set of all $\pi \vdash [n]$ with at most two blocks where every block contains n or $n - 1$.

For the last part, given a set partition π and a suitable permutation δ we know that f is \mathbf{x} -positive if and only if $\delta \circ f$ is \mathbf{x} -positive by Equation (4.2). Also, it is not hard to see that if given a permutation $\delta \in \mathfrak{S}_{n-1}$ with $\delta(n-1) = n-1$, $\tilde{\delta} \in \mathfrak{S}_n$ that is δ with $\tilde{\delta}(n) = n$ and $f \in \text{NCSym}$ is of homogeneous degree $n - 1$ then $\tilde{\delta} \circ (f \uparrow) = (\delta \circ f) \uparrow$. Say particularly that for $\pi \vdash [n - 1]$ that $\delta \in \mathfrak{S}_{n-1}$ is a permutation where $\delta(\pi) = \hat{1}_{k_1} | \hat{1}_{k_2} | \cdots | \hat{1}_{k_L}$ and $\delta(n - 1) = n - 1$. We can see that by Equation (4.2) and Lemma 4.6

$$\delta(\mathbf{x}_\pi) \uparrow = \mathbf{x}_{\delta(\pi)} \uparrow = (\mathbf{x}_{\hat{1}_{k_1}} \mathbf{x}_{\hat{1}_{k_2}} \cdots \mathbf{x}_{\hat{1}_{k_L}}) \uparrow = \mathbf{x}_{\hat{1}_{k_1}} \mathbf{x}_{\hat{1}_{k_2}} \cdots \mathbf{x}_{\hat{1}_{k_{L-1}}} (\mathbf{x}_{\hat{1}_{k_L}} \uparrow)$$

is \mathbf{x} -positive by Lemma 4.6 and by what we have already shown. To complete the proof note that

$$\tilde{\delta}^{-1} \circ (\delta(\mathbf{x}_\pi) \uparrow) = \tilde{\delta}^{-1} \circ \tilde{\delta} \circ (\mathbf{x}_\pi \uparrow) = \mathbf{x}_\pi \uparrow. \quad \square$$

We can now give our desired classification for \mathbf{x} -positivity and \mathbf{x} -negativity.

Theorem 4.8. *For a connected graph G with distinct vertex labels in $[n]$ we have*

$$Y_G = (-1)^{n-1} Z_G$$

where Z_G is \mathbf{x} -positive. Consequently, for a graph G with k connected components and distinct vertex labels in $[n]$ we have

$$Y_G = (-1)^{n-k} Z_G$$

where Z_G is \mathbf{x} -positive. Therefore Y_G is \mathbf{x} -positive if $n - k$ is even, and is \mathbf{x} -negative if $n - k$ is odd.

Proof. We will prove this for a connected graph by doing an induction on the number of vertices n and the number of edges m . Our base case is when $n = 1$ in which case $Y_G = p_1 = \mathbf{x}_1$.

We now assume our result is true for any connected graph with vertices $0 < |V| < n$. We already have the result in the case that a connected graph has the minimal number of edges, that is, a tree by Lemma 4.4, so let $m > n - 1$. We assume the result is true for a connected graph on n vertices with number of edges $n - 1 < |E| < m$.

Since we know that $m > n - 1$ we know that G has an edge that is not a bridge, that is, an edge ϵ such that $G - \epsilon$ is connected. Without loss of generality, by the Relabelling Proposition (Proposition 4.1) assume that ϵ is the edge between vertices $n - 1$ and n . By Deletion-Contraction (Proposition 4.3) we have

$$Y_G = Y_{G-\epsilon} - Y_{G/\epsilon} \uparrow.$$

By induction we know that $Y_{G-\epsilon} = (-1)^{n-1} Z_{G-\epsilon}$ where $Z_{G-\epsilon}$ is \mathbf{x} -positive and $Y_{G/\epsilon} = (-1)^{n-2} Z_{G/\epsilon}$ where $Z_{G/\epsilon}$ is \mathbf{x} -positive. By Lemma 4.7 we know an induced \mathbf{x} -positive function is \mathbf{x} -positive, so $Y_{G/\epsilon} \uparrow = (-1)^{n-2} \tilde{Z}_{G/\epsilon}$ where $\tilde{Z}_{G/\epsilon}$ is \mathbf{x} -positive. Hence the result follows for a connected graph.

The result for a graph with k connected components now follows by Corollary 4.2, Lemma 4.6 and the result for a connected graph. \square

Having classified when Y_G is \mathbf{x} -positive and \mathbf{x} -negative, we now address e -positivity and e -negativity. For this we will need the following three lemmas, the first of which also has an interesting corollary, and the second of which yields a natural multiplicative version of the fundamental theorem of symmetric functions in NCSym.

Lemma 4.9. *For a set partition π we have*

$$Y_{K_\pi} = e_\pi.$$

Proof. This follows immediately by comparing the definitions of Y_{K_π} and e_π . \square

We now have an elegant relationship between the e -basis and the \mathbf{x} -basis.

Corollary 4.10. *For a set partition $\pi = B_1/B_2/\dots/B_{\ell(\pi)} \vdash [n]$ we have e_π is \mathbf{x} -positive if $n - \ell(\pi)$ is even and is \mathbf{x} -negative if $n - \ell(\pi)$ is odd.*

Proof. This follows immediately by Theorem 4.8 and Lemma 4.9. \square

For our second lemma we show that the e -basis satisfies the property exhibited by the p -basis in Lemma 3.3, and the \mathbf{x} -basis in Lemma 4.6. This was also shown by the first author using Möbius functions [10, Lemma 2.1].

Lemma 4.11. *For two set partitions π and σ we have*

$$e_{\pi|\sigma} = e_\pi e_\sigma.$$

Proof. First note that

$$K_\pi | K_\sigma = K_{\pi|\sigma}. \tag{4.3}$$

Consequently by Lemma 4.9 and Proposition 3.5 we have

$$e_{\pi|\sigma} = Y_{K_{\pi|\sigma}} = Y_{K_\pi | K_\sigma} = Y_{K_\pi} Y_{K_\sigma} = e_\pi e_\sigma. \quad \square$$

We now use this second lemma to establish for NCSym the version of the fundamental theorem of symmetric functions that states that the elementary symmetric functions $\{e_i : i \geq 1\}$ in Sym are algebraically independent over \mathbb{Q} and generate Sym.

Theorem 4.12. *The elementary symmetric functions $\{e_\alpha : \alpha \text{ is atomic}\}$ in NCSym are algebraically independent over \mathbb{Q} and freely generate NCSym.*

Proof. Since $\{e_\pi : \pi \vdash [n], n \geq 1\}$ is a basis for NCSym they are all linearly independent. They are also multiplicative along atomic decompositions since Lemma 4.11 shows that for $\pi = \alpha_1 | \alpha_2 | \dots | \alpha_k$, written as its unique atomic decomposition, we have

$$e_\pi = \prod_{i=1}^k e_{\alpha_i}.$$

Thus, all $\{e_\alpha : \alpha \text{ is atomic}\}$ are algebraically independent over \mathbb{Q} and freely generate NCSym. \square

For our third lemma recall that if some function f in a vector space is expressed in terms of a basis $\{b_i\}$ then $[b_i]f$ denotes the coefficient of basis element b_i in f . We will also need the following change of basis formula in NCSym [24, Theorem 3.4] that says that for a set partition $\pi \vdash [n]$

$$p_\pi = \frac{1}{\mu_\Pi(\hat{0}_n, \pi)} \sum_{\sigma \leq \pi} \mu_\Pi(\sigma, \pi) e_\sigma. \tag{4.4}$$

Lemma 4.13. *Given a set partition $\pi = B_1/B_2 \vdash [n]$ with two blocks and a graph G with distinct vertex labels in $[n]$, we have*

$$[e_{\hat{1}_n}]Y_G = \frac{1}{(n-1)!} |[p_{\hat{1}_n}]Y_G|$$

and

$$[e_{B_1/B_2}]Y_G = -\frac{1}{(n-1)!} |[p_{\hat{1}_n}]Y_G| + \frac{(-1)^n}{(|B_1|-1)! (|B_2|-1)!} \sum_{S \subseteq E, \pi(S) = B_1/B_2} (-1)^{|S|}.$$

Proof. Using Proposition 3.2 and Equation (4.4) we have

$$\begin{aligned} Y_G &= \sum_{S \subseteq E(G)} \sum_{\sigma \leq \pi(S)} (-1)^{|S|} \frac{\mu_\Pi(\sigma, \pi(S))}{\mu_\Pi(\hat{0}_n, \pi(S))} e_\sigma \\ &= \sum_{\sigma} \sum_{S \subseteq E(G), \sigma \leq \pi(S)} (-1)^{|S|} \frac{\mu_\Pi(\sigma, \pi(S))}{\mu_\Pi(\hat{0}_n, \pi(S))} e_\sigma. \end{aligned}$$

This means that by Equation (2.1)

$$\begin{aligned} [e_{\hat{1}_n}]Y_G &= \sum_{S \subseteq E(G), \hat{1}_n = \pi(S)} (-1)^{|S|} \frac{\mu_\Pi(\hat{1}_n, \pi(S))}{\mu_\Pi(\hat{0}_n, \pi(S))} \\ &= \frac{1}{(-1)^{n-1} (n-1)!} \sum_{S \subseteq E(G), \hat{1}_n = \pi(S)} (-1)^{|S|}. \end{aligned}$$

Note that the summation above is equal to $[p_{\hat{1}_n}]Y_G$, which has sign $(-1)^{n-1}$ by Theorem 3.1 and Proposition 3.2. This completes the proof for the coefficient $[e_{\hat{1}_n}]Y_G$.

For a set partition B_1/B_2 with two blocks we have by Equations (2.1) and (2.2)

$$\begin{aligned} [e_{B_1/B_2}]Y_G &= \sum_{S \subseteq E(G), B_1/B_2 \leq \pi(S)} (-1)^{|S|} \frac{\mu_\Pi(B_1/B_2, \pi(S))}{\mu_\Pi(\hat{0}_n, \pi(S))} \\ &= \sum_{S \subseteq E(G), \hat{1}_n = \pi(S)} (-1)^{|S|} \frac{\mu_\Pi(B_1/B_2, \pi(S))}{\mu_\Pi(\hat{0}_n, \pi(S))} \\ &+ \sum_{S \subseteq E(G), B_1/B_2 = \pi(S)} (-1)^{|S|} \frac{\mu_\Pi(B_1/B_2, \pi(S))}{\mu_\Pi(\hat{0}_n, \pi(S))} \\ &= \frac{-1}{(-1)^{n-1} (n-1)!} \sum_{S \subseteq E(G), \hat{1}_n = \pi(S)} (-1)^{|S|} \end{aligned}$$

$$+ \frac{1}{(-1)^{n-2}(|B_1| - 1)! (|B_2| - 1)!} \sum_{S \subseteq E(G), B_1/B_2 = \pi(S)} (-1)^{|S|}.$$

Again note that the first summation in the last equality above is equal to $[p_{\hat{1}_n}]Y_G$, which has sign $(-1)^{n-1}$ by Theorem 3.1 and Proposition 3.2. This completes the proof for the coefficient $[e_{B_1/B_2}]Y_G$. \square

We are now ready to classify when Y_G is e -positive or e -negative.

Theorem 4.14. *Given a graph G with distinct vertex labels in $[n]$, Y_G is e -positive if and only if G is a disjoint union of complete graphs. Y_G is never e -negative.*

Proof. By Lemma 4.9 we know that $Y_{K_{\hat{1}_n}} = e_{\hat{1}_n}$, which is certainly e -positive. We now show that if G is connected and not a complete graph then Y_G has both a positive and negative term in the elementary basis.

First by combining Theorem 3.1 and Lemma 4.13 we know that $[e_{\hat{1}_n}]Y_G$ is nonzero and furthermore $[e_{\hat{1}_n}]Y_G > 0$ for all graphs G with n vertices.

Now assume that G is connected and has two vertices u and v with no edge between them. Consider the set partition $\pi = B_1/B_2$ where $B_1 = \{u, v\}$ and $B_2 = [n] \setminus B_1$. Because G lacks an edge from u to v we know that π is not a connected partition of G and there does not exist a $S \subseteq E(G)$ such that $\pi(S) = \pi$. Using Theorem 3.1 and Lemma 4.13 we see that $[e_\pi]Y_G$ is nonzero and

$$[e_\pi]Y_G = -\frac{1}{(n-1)!} |[p_{\hat{1}_n}]Y_G| < 0.$$

Thus, by Proposition 3.5, Corollary 4.2 and Lemma 4.11, the only e -positive graphs are disjoint unions of complete graphs. Proposition 3.5, Corollary 4.2 and Lemma 4.11, also show that there are no e -negative graphs. \square

Note that the argument in the third paragraph of the above proof holds more generally for a connected graph G and two-block set partition $\pi = B_1/B_2$ that is not a connected partition of G . This observation yields the following simple way to locate some negative coefficients.

Corollary 4.15. *For a connected graph G , if the two-block set partition $\pi = B_1/B_2$ is not connected then $[e_\pi]Y_G < 0$.*

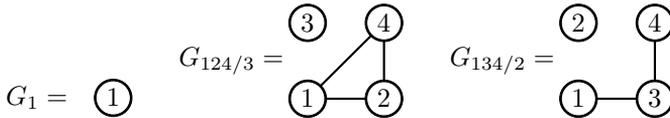
5. Chromatic bases for symmetric functions in NCSym

In this section we discover a multitude of new bases for NCSym arising from chromatic symmetric functions in noncommuting variables, and establish that with the exception of e_π for all $\pi \vdash [n], n \geq 1$ and \mathbf{x}_π for all $\pi \vdash [n], n \geq 1$ with blocks of size 1 and 2, none of the known functions in NCSym can be realized as Y_G for some graph G .

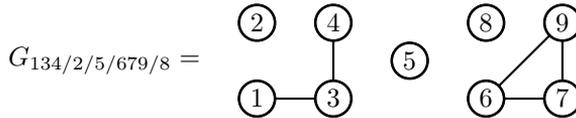
For a graph G , define the set partition $\pi(G)$ to be the set partition such that each block corresponds to the vertices in a connected component of G . Now for each atomic set partition $\alpha \vdash [n]$ choose a graph G_α with n vertices such that $\pi(G_\alpha) = \alpha$. From this, given a generic $\pi \vdash [n]$ written as its unique atomic decomposition $\pi = \alpha_1 \mid \alpha_2 \mid \dots \mid \alpha_k$ define the graph

$$G_\pi = G_{\alpha_1} \mid G_{\alpha_2} \mid \dots \mid G_{\alpha_k}.$$

Example 5.1. Consider the following.



For the above note that $\pi(G_1) = 1$, $\pi(G_{124/3}) = 124/3$ and $\pi(G_{134/2}) = 134/2$. Given these graph choices for atomic partitions 1, 124/3, 134/2 we get the following graph associated to $134/2/5/679/8 = 134/2 \mid 1 \mid 124/3$.



Theorem 5.2. Let $\{G_\alpha : \alpha \text{ is atomic}\}$ be a set of graphs such that G_α has n vertices and $\pi(G_\alpha) = \alpha$ for each α . Then

$$\{Y_{G_\pi} : \pi \vdash [n]\}$$

is a \mathbb{Q} -basis for NCSym^n . Additionally, all $\{Y_{G_\alpha} : \alpha \text{ is atomic}\}$ are algebraically independent over \mathbb{Q} and freely generate NCSym .

Proof. Choose a set $\{G_\alpha : \alpha \text{ is atomic}\}$ such that G_α has n vertices and $\pi(G_\alpha) = \alpha$ for each α . Given $\pi \vdash [n]$ we know by definition that if $\sigma \in L_{G_\pi}$ then $\sigma \leq \pi$. Hence, by Theorem 3.1,

$$Y_{G_\pi} = \sum_{\sigma \leq \pi} c_{\sigma\pi} p_\sigma$$

and $c_{\pi\pi} = \mu_L(\hat{0}, \pi) \neq 0$ also from Theorem 3.1. From this we can conclude that the Y_{G_π} for all $\pi \vdash [n]$ are linearly independent and hence $\{Y_{G_\pi} : \pi \vdash [n]\}$ is a \mathbb{Q} -basis for NCSym^n . They are also multiplicative along atomic decompositions since Proposition 3.5 shows that for $\pi = \alpha_1 \mid \alpha_2 \mid \dots \mid \alpha_k$, written as its unique atomic decomposition, we have

$$Y_{G_\pi} = \prod_{i=1}^k Y_{G_{\alpha_i}}.$$

Thus, all $\{Y_{G_\alpha} : \alpha \text{ is atomic}\}$ are algebraically independent over \mathbb{Q} and freely generate NCSym. \square

Now we move on to identifying whether Y_G , for some graph G , is a known function in NCSym. The first functions we will study are the Schur-like functions of Rosas and Sagan, S_λ where λ is an *integer* partition. They are defined in [24, Section 6], and since we will need them only for the next result, we refer the reader there for the general definition, though note that $S_{(1^n)} = e_{\hat{1}_n}$, where (1^n) is the integer partition consisting of n parts equal to 1. The only other property we recall is that for $\lambda \vdash n$ [24, Theorem 6.2 (iii)]

$$\rho(S_\lambda) = n!s_\lambda \tag{5.1}$$

where s_λ is the classical Schur function.

Proposition 5.3. *Of the functions $\{S_\lambda\}_{\lambda \vdash n \geq 1}$ in NCSym and their scalar multiples only $S_{(1^n)}$ can be realized as Y_G for some graph G . In particular*

$$S_{(1^n)} = Y_{K_n}.$$

Proof. Assume $Y_G = c_\lambda S_\lambda$. By Equation (5.1) we have

$$X_G = \rho(Y_G) = \rho(c_\lambda S_\lambda) = c_\lambda n!s_\lambda$$

that by [9, Theorem 2.8] is possible if and only if $\lambda = (1^n)$, $G = K_n$ and $c_\lambda = 1$. \square

For the functions indexed by set partitions we have one more family to define. The *complete homogeneous symmetric function in NCSym*, h_π where $\pi \vdash [n]$, is given by [24, Theorem 3.4]

$$h_\pi = \sum_{\sigma \leq \pi} |\mu_\Pi(\hat{0}_n, \sigma)| p_\sigma.$$

Due to the common proof technique, we will now consider all except the \mathbf{x} -basis.

Theorem 5.4. *Of the functions $\{e_\pi\}_{\pi \vdash [n], n \geq 1}$, $\{h_\pi\}_{\pi \vdash [n], n \geq 1}$, $\{m_\pi\}_{\pi \vdash [n], n \geq 1}$ and $\{p_\pi\}_{\pi \vdash [n], n \geq 1}$ in NCSym and their scalar multiples only $\{e_\pi\}_{\pi \vdash [n], n \geq 1}$ can be realized as Y_G for some graph G . In particular*

$$e_\pi = Y_{K_\pi}.$$

Proof. Since $e_{\hat{0}_n} = h_{\hat{0}_n} = p_{\hat{0}_n}$ for $n \geq 1$ we consider this to be $e_{\hat{0}_n}$ for convenience, and likewise consider $e_{\hat{1}_n} = m_{\hat{0}_n}$ for $n \geq 1$ to be $e_{\hat{1}_n}$. Assume $Y_G = c_\pi b_\pi$ where $b = h, m$ or p and $\pi \neq \hat{0}_n$ for $n \geq 1$. Then

$$X_G = \rho(Y_G) = \rho(c_\pi b_\pi) = c_\pi \rho(b_\pi).$$

Now [24, Theorem 2.1] states that $\rho(b_\pi) = \lambda(\pi)! b_{\lambda(\pi)}$ if $b = h$, $\rho(b_\pi) = \lambda(\pi)! b_{\lambda(\pi)}$ if $b = m$, and $\rho(b_\pi) = b_{\lambda(\pi)}$ if $b = p$. However, by [9, Theorem 2.8] none of these outcomes are possible. By [24, Theorem 2.1] we also have

$$X_G = \rho(Y_G) = \rho(c_\pi e_\pi) = c_\pi \rho(e_\pi) = c_\pi \lambda(\pi)! e_{\lambda(\pi)}$$

that by [9, Theorem 2.8] is possible if and only if G is a disjoint union of complete graphs and $c_\pi = 1$. The final equation is then by Lemma 4.9. \square

Finally, we consider the \mathbf{x} -basis, for which we need the following lemma.

Lemma 5.5. *For any graph G containing a connected component with at least three vertices, we have Y_G has at least two terms in its expansion in the \mathbf{x} -basis.*

Proof. Consider a connected graph G with $n \geq 3$ vertices. Substituting Equation (2.4) into Theorem 3.1 we get

$$Y_G = \sum_{\pi \in L_G} \mu_L(\hat{0}_n, \pi) \sum_{\sigma \leq \pi} \mathbf{x}_\sigma$$

and $\hat{1}_n \in L_G$ so the coefficient of $\mathbf{x}_{\hat{1}_n}$ is $\mu_L(\hat{0}_n, \hat{1}_n) \neq 0$. If G is not the complete graph then there exist two vertices that are not adjacent. Say their labels are u and v with $u < v$. Consider the set partition

$$\sigma = uv/1 \cdots \hat{u} \cdots \hat{v} \cdots n \notin L_G$$

where \hat{u}, \hat{v} means that u and v are removed. Note the coefficient of \mathbf{x}_σ is also $\mu_L(\hat{0}_n, \hat{1}_n) \neq 0$, and we have two nonzero terms. If G is the complete graph, then the poset $L_G = \Pi_n$, and using Equations (2.1) and (2.2) we can calculate the coefficient of $\pi = 1/23 \cdots n$ is

$$\begin{aligned} \mu_{\Pi}(\hat{0}_n, \pi) + \mu_{\Pi}(\hat{0}_n, \hat{1}_n) &= (-1)^{n-2}(n-2)! + (-1)^{n-1}(n-1)! \\ &= ((-1)^{n-2} + (-1)^{n-1}(n-1))(n-2)! \\ &= (-1)^{n-1}(n-2)(n-2)! \neq 0 \end{aligned}$$

since $n > 2$.

Lastly, consider a graph G with a connected component with at least three vertices, \tilde{G} . By Corollary 4.2 and Lemma 4.6 we know that Y_G has at least as many terms as $Y_{\tilde{G}}$ and we are done. \square

Our final theorem completes the classification of which known functions in NCSym can be realized as Y_G for some graph G .

Theorem 5.6. *Of the functions $\{\mathbf{x}_\pi\}_{\pi \vdash [n], n \geq 1}$ in NCSym and their scalar multiples only $\pm \mathbf{x}_\pi$ where $\pi = B_1/B_2/\dots/B_{\ell(\pi)}$ and $|B_i| = 1, 2$ for $1 \leq i \leq \ell(\pi)$ can be realized as Y_G for some graph G . In particular*

$$\mathbf{x}_\pi = (-1)^{t_\pi} Y_{K_\pi}$$

where $t_\pi =$ the number of $|B_i| = 2$ for $1 \leq i \leq \ell(\pi)$.

Proof. Assume $Y_G = c_\pi \mathbf{x}_\pi$ for some $\pi \vdash [n], n \geq 1$. By Lemma 5.5 we know G only contains connected components with one or two vertices. Therefore $G = K_\pi$ where $\pi = B_1/B_2/\dots/B_{\ell(\pi)}$ and $|B_i| = 1, 2$ for $1 \leq i \leq \ell(\pi)$. By using Lemma 4.9, Proposition 3.2 and Equation (2.3), respectively, twice we have

$$\begin{aligned} e_1 &= Y_{K_1} = p_1 = \mathbf{x}_1 \\ e_{12} &= Y_{K_{12}} = -p_{12} + p_{1/2} = -\mathbf{x}_{12}. \end{aligned}$$

Thus by choosing $\delta \in \mathfrak{S}_n$ such that

$$\delta(K_\pi) = K_{|B_1|} | K_{|B_2|} | \dots | K_{|B_{\ell(\pi)}|} = K_{\hat{1}_{|B_1|}} | K_{\hat{1}_{|B_2|}} | \dots | K_{\hat{1}_{|B_{\ell(\pi)}|}}$$

we have the following equalities, given respectively by Corollary 4.2, Lemma 4.9, the above substitutions to the \mathbf{x} -basis, Lemma 4.6, Equation (4.2) and the definition of δ

$$\begin{aligned} Y_{K_\pi} &= \delta^{-1} \circ (Y_{K_{|B_1|}} Y_{K_{|B_2|}} \dots Y_{K_{|B_{\ell(\pi)}|}}) = \delta^{-1} \circ (e_{\hat{1}_{|B_1|}} e_{\hat{1}_{|B_2|}} \dots e_{\hat{1}_{|B_{\ell(\pi)}|}}) \\ &= \delta^{-1} \circ ((-1)^{t_\pi} \mathbf{x}_{\hat{1}_{|B_1|}} \mathbf{x}_{\hat{1}_{|B_2|}} \dots \mathbf{x}_{\hat{1}_{|B_{\ell(\pi)}|}}) = (-1)^{t_\pi} \delta^{-1} \circ (\mathbf{x}_{\hat{1}_{|B_1|}|\hat{1}_{|B_2|}|\dots|\hat{1}_{|B_{\ell(\pi)}|}}) \\ &= (-1)^{t_\pi} (\mathbf{x}_{\delta^{-1}(\hat{1}_{|B_1|}|\hat{1}_{|B_2|}|\dots|\hat{1}_{|B_{\ell(\pi)}|})}) = (-1)^{t_\pi} \mathbf{x}_\pi \end{aligned}$$

where $t_\pi =$ the number of $|B_i| = 2$ for $1 \leq i \leq \ell(\pi)$. \square

Example 5.7. Returning to Example 2.9 we have $Y_{K_{13/2}} = -\mathbf{x}_{13/2}$.

Acknowledgments

We would like to thank the referee for the care they took with our paper.

Funding

Both authors were supported in part by the National Sciences and Engineering Research Council of Canada.

References

- [1] M. Aguiar, et al., Supercharacters, symmetric functions in noncommuting variables, and related Hopf algebras, *Adv. Math.* 229 (2012) 2310–2337.
- [2] N. Bergeron, C. Hohlweg, M. Rosas, M. Zabrocki, Grothendieck bialgebras, partition lattices, and symmetric functions in noncommutative variables, *Electron. J. Combin.* 13 (2006) R75.
- [3] N. Bergeron, C. Reutenauer, M. Rosas, M. Zabrocki, Invariants and coinvariants of the symmetric group in noncommuting variables, *Canad. J. Math.* 60 (2008) 266–296.
- [4] N. Bergeron, M. Zabrocki, The Hopf algebras of symmetric functions and quasi-symmetric functions in non-commutative variables are free and co-free, *J. Algebra Appl.* 8 (2009) 581–600.
- [5] G. Birkhoff, A determinant formula for the number of ways of coloring a map, *Ann. of Math.* 14 (1912) 42–46.
- [6] M. Can, B. Sagan, Partitions, rooks, and symmetric functions in noncommuting variables, *Electron. J. Combin.* 18 (2011) P3.
- [7] S. Cho, J. Huh, On e -positivity and e -unimodality of chromatic quasisymmetric functions, arXiv:1711.07152v2.
- [8] S. Cho, S. van Willigenburg, Chromatic bases for symmetric functions, *Electron. J. Combin.* 23 (2016) P1.15.
- [9] S. Cho, S. van Willigenburg, Chromatic classical symmetric functions, *J. Comb.* 9 (2018) 401–409.
- [10] S. Dahlberg, Triangular ladders $P_{d,2}$ are e -positive, arXiv:1811.04885v1.
- [11] S. Dahlberg, S. van Willigenburg, Lollipop and lariat symmetric functions, *SIAM J. Discrete Math.* 32 (2018) 1029–1039.
- [12] A. Foley, C. Hoàng, O. Merkel, Classes of graphs with e -positive chromatic symmetric function, arXiv:1808.03391v1.
- [13] A. Foley, J. Kazdan, L. Kröll, S. Martínez Alberga, O. Melnyk, A. Tenenbaum, Spiders and their kin (K_n), arXiv:1812.03476v1.
- [14] V. Gasharov, Incomparability graphs of $(3+1)$ -free posets are s -positive, *Discrete Math.* 157 (1996) 193–197.
- [15] V. Gasharov, On Stanley’s chromatic symmetric function and clawfree graphs, *Discrete Math.* 205 (1999) 229–234.
- [16] D. Gebhard, B. Sagan, A chromatic symmetric function in noncommuting variables, *J. Algebraic Combin.* 13 (2001) 227–255.
- [17] M. Guay-Paquet, A modular law for the chromatic symmetric functions of $(3+1)$ -free posets, arXiv:1306.2400v1.
- [18] A. Hamel, C. Hoàng, J. Tuero, Chromatic symmetric functions and H -free graphs, arXiv:1709.03354v1.
- [19] M. Harada, M. Precup, The cohomology of Abelian Hessenberg varieties and the Stanley-Stembridge conjecture, arXiv:1709.06736v2.
- [20] J. Huh, S.-Y. Nam, M. Yoo, Melting lollipop chromatic quasisymmetric functions and Schur expansion of unicellular LLT polynomials, arXiv:1812.03445v1.
- [21] A. Lauve, M. Mastnak, The primitives and antipode in the Hopf algebra of symmetric functions in noncommuting variables, *Adv. in Appl. Math.* 47 (2011) 536–544.
- [22] R. Orellana, G. Scott, Graphs with equal chromatic symmetric function, *Discrete Math.* 320 (2014) 1–14.
- [23] B. Pawłowski, Chromatic symmetric functions via the group algebra of S_n , arXiv:1802.05470v3.
- [24] M. Rosas, B. Sagan, Symmetric functions in noncommuting variables, *Trans. Amer. Math. Soc.* 358 (2006) 215–232.
- [25] J. Shreshian, M. Wachs, Chromatic quasisymmetric functions, *Adv. Math.* 295 (2016) 497–551.
- [26] R. Stanley, A symmetric function generalization of the chromatic polynomial of a graph, *Adv. Math.* 111 (1995) 166–194.
- [27] R. Stanley, Graph colorings and related symmetric functions: ideas and applications. A description of results, interesting applications, & notable open problems, *Discrete Math.* 193 (1998) 267–286.
- [28] R. Stanley, J. Stembridge, On immanants of Jacobi-Trudi matrices and permutations with restricted position, *J. Combin. Theory Ser. A* 62 (1993) 261–279.
- [29] N. Thiem, Branching rules in the ring of superclass functions of unipotent upper-triangular matrices, *J. Algebraic Combin.* 31 (2010) 267–298.

- [30] S. Tsujie, The chromatic symmetric functions of trivially perfect graphs and cographs, arXiv:1707.04058v1.
- [31] M. Wolf, Symmetric functions of non-commutative elements, *Duke Math. J.* 2 (1936) 626–637.
- [32] M. Wolfe, Symmetric chromatic functions, *Pi Mu Epsilon J.* 10 (1998) 643–757.