

MATCHING POLYTOPES AND SPECHT MODULES

RICKY INI LIU

ABSTRACT. We prove that the dimension of the Specht module of a forest G is the same as the normalized volume of the matching polytope of G . We also associate to G a symmetric function s_G (analogous to the Schur symmetric function s_λ for a partition λ) and investigate its combinatorial and representation-theoretic properties in relation to the Specht module and Schur module of G . We then use this to define notions of standard and semistandard tableaux for forests.

1. INTRODUCTION

The (*fractional*) *matching polytope* of a graph G is the space of all nonnegative edge weightings of G such that the sum of the weights around any vertex is at most 1. This polytope, along with the related *perfect matching polytope*, have been well studied with respect to combinatorial optimization ([4], [5]).

In this paper, we will first address a natural question, namely, how do we compute the volume of the matching polytope of a forest? We will present several simple recurrences for this volume in Section 2.

We then consider a seemingly unrelated question in the representation theory of the symmetric group. The *Specht module* construction takes the Young diagram of a partition λ of n and produces an irreducible representation S^λ of the symmetric group Σ_n . However, this construction can be applied to diagrams of boxes other than Young diagrams of partitions, though it will generally not yield an irreducible representation. For instance, given a skew Young diagram λ/μ , one can apply the same construction to get a skew Specht module. Moreover, when one writes the decomposition of this skew Specht module in terms of irreducible representations, the coefficients that appear are the famous Littlewood-Richardson coefficients [7]. The general question of describing the Specht module obtained for an arbitrary starting diagram has been the subject of much study ([3], [6], [8], [9]).

Our main result is to show that the two questions described above are in fact related. First we note that the Specht module construction naturally applies to bipartite graphs. We then prove our main result: for a forest G , the dimension of the Specht module S^G is exactly the normalized volume $V(G)$ of the matching polytope of G .

This result allows us to draw a combinatorial relationship between matchings and Specht modules. In particular, we will ask how to compute the restriction of the Specht module of a forest with n edges from Σ_n to Σ_{n-1} . In the case of partitions, this amounts to a “branching rule”, in which one only needs to identify

Received by the editors October 7, 2010 and, in revised form, November 29, 2010.
2010 *Mathematics Subject Classification*. Primary 05E10.

corner boxes of a Young diagram. We will show that for forests there is a similar rule, in which the role of corner boxes is replaced by that of edges in a so-called *almost perfect matching*.

We will also define a symmetric function s_G analogous to the Schur symmetric function s_λ for a partition λ and describe its combinatorial and representation-theoretic significance. In particular, we will show that it somehow universally describes one of the recurrences satisfied by $\dim S^G = V(G)$.

Finally, we will discuss the Schur module construction, which is the analogue of the Specht module construction for $GL(N)$ -modules. We will then discuss its relation to the previous objects of study. This will allow us to define analogues of standard and semistandard tableaux for forests.

In Section 2, we will discuss the basic properties and recurrences for matching polytopes that we will need for the remainder of the paper. In Section 3, we will introduce Specht modules and prove our main result, as well as deduce a “branching rule” for Specht modules of forests. In Section 4, we will define a Schur symmetric function for forests and describe its universality. In Section 5, we will discuss Schur modules and define a notion of semistandard tableaux for forests. Finally, in Section 6 we will formulate some remaining questions and concluding remarks.

2. MATCHING POLYTOPES

In this section, we will describe several recurrences for the normalized volume of the matching polytope of a forest. We will also provide a characterization for this volume via these recurrences.

We begin with some key definitions.

Definition. Let $G = (V, E)$ be a graph. The *matching polytope* M_G of G is the space of all nonnegative edge weightings $w: E \rightarrow \mathbf{R}_{\geq 0}$ such that for all $v \in V$,

$$\sum_{e \ni v} w(e) \leq 1.$$

Note that M_G is a rational convex polytope of full dimension in \mathbf{R}^n , where $n = |E|$ is the number of edges in G . The reason that M_G is called a matching polytope is due to the following definition.

Definition. A *matching* M of a graph G is a collection of edges of G such that no two edges of M share a vertex.

Given a matching M of G , let us write χ_M for the edge weighting of G defined by $\chi_M(e) = 1$ if $e \in M$ and 0 otherwise. Then clearly $\chi_M \in M_G$ for all matchings M . Since M_G lies in the hypercube $[0, 1]^n$, each of the χ_M is a vertex of M_G , and all lattice points in M_G are of the form χ_M for some M .

The following well-known proposition characterizes when M_G is a lattice polytope. (See, for instance, [1].)

Proposition 2.1. *The matching polytope M_G is a lattice polytope if and only if G is bipartite. In this case, M_G is the convex hull of χ_M , where M ranges over all matchings of G .*

It follows that when G is bipartite, the volume of M_G is an integer multiple of $\frac{1}{n!}$. Therefore, we let

$$V(G) = n! \cdot \text{vol}(M_G) \in \mathbf{Z}.$$

When considering bipartite graphs, it will be important later that we distinguish the two parts of our graph. Therefore, we will henceforth assume that all bipartite graphs are equipped with a bipartition or 2-coloring of the vertices; that is, each vertex will be colored either black or white such that no two adjacent vertices have the same color. We will also assume that our graphs contain no isolated vertices. (Removing any isolated vertices will not change the matching polytope of the graph.)

In particular, let \mathcal{F} be the set of all finite (2-colored) forests. We will now present the recurrences that will allow us to compute $V(G)$ for any $G \in \mathcal{F}$.

First, we require the following base case.

Proposition 2.2. *Let T_n be the star with n edges and white center vertex. Then $V(T_n) = 1$.*

Proof. If the edges have weights w_1, \dots, w_n , then M_G is defined by $0 \leq w_i \leq 1$ for all i and $\sum_{i=1}^n w_i \leq 1$. This is just the elementary n -simplex in \mathbf{R}^n , which has volume $\frac{1}{n!}$. □

Next, we use the following recurrence to reduce to connected graphs.

Proposition 2.3. *Let G be a disjoint union of graphs $G_1 + G_2$, where G_1 has m edges and G_2 has $n - m$ edges. Then $V(G) = \binom{n}{m} V(G_1) V(G_2)$.*

Proof. Clearly, $M_G = M_{G_1} \times M_{G_2}$. Hence,

$$\begin{aligned} V(G) &= n! \cdot \text{vol}(M_G) \\ &= n! \cdot \text{vol}(M_{G_1}) \cdot \text{vol}(M_{G_2}) \\ &= \frac{n!}{m! \cdot (n - m)!} \cdot V(G_1) \cdot V(G_2). \end{aligned} \quad \square$$

We now present a more interesting recurrence for $V(G)$.

Let H be a graph, and let v_1 and v_2 be distinct vertices of H . We construct three graphs G , G_1 , and G_2 as follows. Let G be the graph obtained from H by adding pendant edges $\overline{v_1 v'_1}$ and $\overline{v_2 v'_2}$. (Then v'_1 and v'_2 are leaves in G .) Let G_1 be the graph obtained from H by adding a pendant edge $\overline{v_1 v'_1}$ and an edge $\overline{v_1 v_2}$, and let G_2 be obtained from H by adding a pendant edge $\overline{v_2 v'_2}$ and an edge $\overline{v_1 v_2}$ (see Figure 1).

Proposition 2.4 (Leaf recurrence). *Let G , G_1 , and G_2 be as described above. Then $V(G) = V(G_1) + V(G_2)$.*

Proof. Consider the matching polytope M_G , and write w_i for the weight of $\overline{v_i v'_i}$. Let M_G^1 be the intersection of M_G with the halfspace $w_1 \geq w_2$, and let M_G^2 be the intersection of M_G with the halfspace $w_2 \geq w_1$. Clearly, $\text{vol}(M_G) = \text{vol}(M_G^1) + \text{vol}(M_G^2)$.

Consider the matching polytope M_{G_1} , and write z_1 for the weight of $\overline{v_1 v'_1}$ and z_2 for the weight of $\overline{v_1 v_2}$. Then for any $z \in M_{G_1}$, note that by letting $w_1 = z_1 + z_2$, $w_2 = z_2$, and keeping all other weights the same, we obtain a point $f(z) \in M_G^1$. Moreover, f is a bijection from M_{G_1} to M_G^1 : the inverse map is given by letting $z_1 = w_1 - w_2$, $z_2 = w_2$, and keeping all other weights the same. Since f is a volume-preserving linear transformation, it follows that $\text{vol}(M_{G_1}) = \text{vol}(M_G^1)$.

An analogous argument gives that $\text{vol}(M_{G_2}) = \text{vol}(M_G^2)$. The result follows easily. □




FIGURE 1. Three graphs as related by the leaf recurrence. They differ only in the edges marked; inside the circles the graphs are arbitrary as long as they are the same in all three cases.

Given any function f on forests, we will say that f satisfies the leaf recurrence if $f(G) = f(G_1) + f(G_2)$ for any three forests G , G_1 , and G_2 as described above.

We claim that the previous three propositions suffice to calculate $V(G)$ for any forest G .

Proposition 2.5. *There is a unique function $f: \mathcal{F} \rightarrow \mathbf{R}$ satisfying the following properties:*

- (1) For the star T_n with n edges and white center vertex, $f(T_n) = 1$.
- (2) If G_1 and G_2 have m and $n - m$ edges, respectively, then $f(G_1 + G_2) = \binom{n}{m} f(G_1) f(G_2)$.
- (3) The function f satisfies the leaf recurrence.

In this case, $f(G) = V(G)$, the normalized volume of the matching polytope of G .

Proof. Since we have seen that $V(G)$ satisfies these three properties, it suffices to show that for any forest G , one can determine $f(G)$ using these properties alone.

We induct first on n , the number of edges of G . By (2), we may then assume that G is connected and hence a tree. Choose a white vertex v_0 to be the root of G . We will then induct on s , the sum of the distances from all vertices to v_0 . The base case occurs when $s = n$, which occurs only when G is T_n , in which case $f(G) = 1$ by (1).

Suppose $s > n$. Then G must have a leaf v_1' whose neighbor is $v_1 \neq v_0$. Let v_2 be the neighbor of v_1 closest to v_0 , and let G' be the graph obtained from G by removing the edge $\overline{v_1 v_1'}$ and adding a pendant edge $\overline{v_2 v_2'}$. Then if H is the forest obtained from G by removing the edge $\overline{v_1 v_2}$ and adding a pendant edge $\overline{v_2 v_2'}$, we have that $f(H) = f(G) + f(G')$ by (3). Since H is disconnected, we can calculate $f(H)$ by induction. But $s(G') = s(G) - 1$, so we can determine $f(G')$ by induction as well. Thus we can determine $f(G) = f(H) - f(G')$, completing the proof. \square

Note from the proof of Proposition 2.5 that even if we only stipulate that the three conditions hold when all graphs involved have at most n edges, we still find that $f(G) = V(G)$ for all graphs with at most n edges.

Although the leaf recurrence suffices to calculate $V(G)$, it is not very efficient for calculations. Moreover, it does not give a subtraction-free rule for determining $V(G)$, in the sense that while it allows one to write $V(G)$ as a linear combination of

multinomial coefficients, it does not give a nonnegative linear combination. Therefore, we will introduce an alternative recurrence for calculating $V(G)$ for $G \in \mathcal{F}$. In any case, we will use both of these recurrences in Section 3. We first make the following definition.

Definition. We say that a matching M of G is *almost perfect* if every isolated edge of G lies in M and every nonleaf vertex of G is contained in an edge of M .

The importance of almost perfect matchings lies in the following proposition.

Proposition 2.6. *Let G be a bipartite graph, and let M be an almost perfect matching of G . Then*

$$V(G) = \sum_{e \in M} V(G \setminus e).$$

Proof. We claim that it is possible to partition M_G into cones C_e of volume $\frac{1}{n!}V(G \setminus e)$ for $e \in M$. Note that $M_{G \setminus e}$ embeds naturally in M_G : in fact, $M_{G \setminus e} = M_G \cap \{w \mid w(e) = 0\}$. For $e \in M$, let C_e for the cone with vertex χ_M and base $M_{G \setminus e}$. Since C_e has height 1, it has volume $\frac{1}{n} \text{vol}(M_{G \setminus e}) = \frac{1}{n!}V(G \setminus e)$, so it suffices to show that the cones C_e partition M_G (up to a measure zero set).

Let $w \in M_G$, and let e_0 be such that $w(e_0) = t$ is minimum among all $w(e)$ for $e \in M$. Let

$$w' = \frac{1}{1-t}(w - t \cdot \chi_M),$$

so that $w = t \cdot \chi_M + (1-t) \cdot w'$. Clearly, $0 \leq t \leq 1$. By our choice of e_0 , w' is a nonnegative weighting of $G \setminus e_0$. We claim that it lies in M_G (and hence in $M_{G \setminus e_0}$), which will imply that $w \in C_{e_0}$.

Note that although M_G is defined by an inequality for each vertex of G , the condition at a leaf that is not part of an isolated edge is redundant, for it is superseded by the condition at the adjacent vertex. Therefore, to check that $w' \in M_G$, it suffices to check that the sum of the weights of the edges incident to any nonleaf is at most 1 and that the weight on any isolated edge is at most 1. If v is a nonleaf of G , then

$$\sum_{e \ni v} w'(e) = \frac{1}{1-t} \left(\sum_{e \ni v} w(e) - t \cdot \sum_{e \ni v} \chi_M(e) \right) = \frac{1}{1-t} \left(\sum_{e \ni v} w(e) - t \right) \leq 1.$$

Similarly, if e is an isolated edge of G , then $w'(e) = \frac{1}{1-t}(w(e) - t) \leq 1$. It follows that $w' \in M_{G \setminus e_0}$, so $w \in C_{e_0}$. Therefore the cones C_e for $e \in M$ cover M_G .

We now show that the C_e have disjoint interiors. Suppose that w lies in the interior of C_{e_0} . Then we can write $w = t \cdot \chi_M + (1-t) \cdot w'$, where $0 < t < 1$ and w' lies in the interior of $M_{G \setminus e_0}$, so $w(e) = t + (1-t) \cdot w'(e)$ for $e \in M$. Since w' lies in the interior of $M_{G \setminus e_0}$, we have that $w'(e_0) = 0$, but $w'(e) > 0$ for all other $e \in M$. Therefore, e_0 is uniquely determined: $w(e_0)$ is the unique minimum among all $w(e)$ for $e \in M$. This proves the result. \square

Not all graphs have almost perfect matchings, but every forest does.

Proposition 2.7. *Every forest has an almost perfect matching.*

Proof. It suffices to show that every rooted tree G with at least one edge has an almost perfect matching M such that the root lies in an edge of M . We induct on the number of edges of G , with the base case being trivial. Choose any edge e

incident to the root, and consider the forest G' obtained from G by removing both endpoints of e as well as any edge incident to either endpoint. Root each component of G' at the vertex that was closest to e in G . By induction, any component with at least one edge has an almost perfect matching containing an edge that contains the root. It is then easy to check that the union of these matchings together with e is an almost perfect matching of G . \square

It follows that we can use Proposition 2.6 to recursively compute $V(G)$ for any $G \in \mathcal{F}$. While Proposition 2.6 has the advantage of giving a subtraction-free rule, it has the disadvantage that unlike the leaf recurrence, it is not homogeneous. In other words, here $V(G)$ is written in terms of $V(H)$, where H does not have the same number of edges as G .

One can express the resulting rule for $V(G)$ by defining the following notion of standard labelings of G .

Definition. Fix an almost perfect matching $M(G)$ for every $G \in \mathcal{F}$. For any $G \in \mathcal{F}$, we say that an edge labeling $z: E \rightarrow \mathbf{N}$ is *standard* if z is a bijection between E and $[n] = \{1, \dots, n\}$ such that $z^{-1}(n) \in M(G)$ and, if $n > 1$, $z|_{G \setminus z^{-1}(n)}$ is a standard labeling of $G \setminus z^{-1}(n)$.

The following proposition is then immediate.

Proposition 2.8. *For any $G \in \mathcal{F}$, $V(G)$ is the number of standard labelings of G .*

Proof. Both $V(G)$ and the number of standard labelings of G satisfy the recurrence in Proposition 2.6. \square

These standard labelings will turn out to be analogues of standard Young tableaux. We will also define analogues of semistandard Young tableaux in Section 5.

Armed with the results of this section, we will now be able to draw the connection between matching polytopes and Specht modules.

3. SPECHT MODULES

In this section, we first define the notion of a Specht module for diagrams and bipartite graphs. We will also summarize some known results about them. We will then state and prove our main theorem. As an immediate corollary, we will also derive a branching rule for Specht modules of forests. (For more information on the basic background needed for this section, see, for instance, [10].)

Consider an array of unit lattice boxes in the plane. We will write (i, j) to denote the box in the i th row from the top and the j th column from the left, where i and j are positive integers. By a *diagram*, we will mean any subset of these boxes.

For instance, a *partition* $\lambda = (\lambda_1, \lambda_2, \dots)$ of $|\lambda| = n$ is a sequence of weakly decreasing nonnegative integers summing to n . (We may add or ignore trailing zeroes to λ as convenient.) Then the *Young diagram* of λ consists of all boxes (i, j) with $j \leq \lambda_i$. We may refer to a partition and its Young diagram interchangeably. If λ and μ are partitions such that $\mu_i \leq \lambda_i$ for all i , then the *skew Young diagram* λ/μ consists of all boxes in λ not in μ .

Given a diagram D , a *tableau* of shape D is a filling of the boxes of D with positive integers. In the case of a (skew) Young diagram, a tableau T is called a *semistandard Young tableau* if the rows are weakly increasing and the columns are strictly increasing. A semistandard Young tableau is *standard* if it contains only the numbers $1, 2, \dots, n$, each exactly once.

Consider any diagram D of n boxes. Order the boxes of D arbitrarily, and let the symmetric group Σ_n act on them in the obvious way. Let R_D be the subgroup containing those $\sigma \in \Sigma_n$ that stabilize each row of D , and likewise define C_D for columns of D . Let $\mathbf{C}[\Sigma_n]$ denote the group algebra over Σ_n , and consider elements

$$R(D) = \sum_{\sigma \in R_D} \sigma \quad \text{and} \quad C(D) = \sum_{\sigma \in C_D} \text{sgn}(\sigma)\sigma.$$

Definition. The *Specht module* (over \mathbf{C}) of D is the left ideal

$$S^D = \mathbf{C}[\Sigma_n]C(D)R(D).$$

Clearly, S^D is a representation of Σ_n by left multiplication. We may interpret basis elements of $\mathbf{C}[\Sigma_n]$ as bijections from D to $[n]$, or, equivalently, as tableaux of shape D with labels $1, 2, \dots, n$. Then multiplying on the right corresponds to applying a permutation to the boxes of D , while the action of Σ_n on the left corresponds to applying a permutation to the labels $[n]$.

It is well known that the irreducible representations of Σ_n are exactly the Specht modules S^λ , where λ ranges over all partitions of n . Moreover, the dimension of S^λ is f^λ , the number of standard Young tableaux of shape λ .

Given two diagrams D and E , let $D \oplus E$ be the diagram obtained by drawing them so that no box of D lies in the same row or column as a box of E . In other words, choose some positive integer N such that no box of D has a coordinate larger than N , and then let $D \oplus E = D \cup \tilde{E}$, where \tilde{E} is the same diagram as E but shifted by the vector (N, N) . If D has m boxes and E has $n - m$ boxes, then we can think of Σ_m and Σ_{n-m} acting on D and E independently (where we think of $\Sigma_m \subset \Sigma_n$ as acting on the first m letters and $\Sigma_{n-m} \subset \Sigma_n$ as acting on the last $n - m$ letters). Then it follows that

$$S^{D \oplus E} = \text{Ind}_{\Sigma_m \times \Sigma_{n-m}}^{\Sigma_n} (S^D \otimes S^E).$$

Let us write c_λ^D for the multiplicity of S^λ in S^D , so that

$$S^D \cong \bigoplus (S^\lambda)^{\oplus c_\lambda^D}.$$

In the case when D is a skew shape, the coefficients c_λ^D are the famous *Littlewood-Richardson coefficients*, and there are a number of combinatorial ways to calculate them, collectively called *Littlewood-Richardson rules*. However, the question of how to compute $\dim S^D$ or c_λ^D for an arbitrary diagram D is still open. The most general known result, due to Reiner and Shimozono [9], applies to so-called *percentage-avoiding* diagrams.

Although Specht modules are traditionally defined for diagrams, we can also define the Specht module of a bipartite graph. Indeed, observe that permuting the rows or permuting the columns of D does not change the structure of S^D up to isomorphism. Therefore, for any (2-colored) bipartite graph G , we can define the Specht module S^G as follows.

Number the white vertices of G with positive integers, and likewise number the black vertices. Then let $D = D(G)$ be the diagram that contains the box (i, j) if and only if G contains an edge connecting white vertex i and black vertex j . By our observation, the structure of S^D does not depend on the labeling of the vertices of G , so we will write S^G for $S^{D(G)}$. We will also sometimes refer to the graph G and

the diagram D interchangeably. We will say two diagrams are *equivalent* if their underlying graphs are the same. (We may also implicitly identify them at times as convenient.)

In other words, we let Σ_n be the symmetric group on the edges of G . Then C_G corresponds to the subgroup that stabilizes the incident edges to any black vertex and R_G to the subgroup that stabilizes the incident edges to any white vertex. We can then form the Specht module $S^G = \mathbf{C}[\Sigma_n]C(G)R(G)$ exactly as before.

We can now state our main theorem on the Specht module of a forest.

Theorem 3.1. *For all $G \in \mathcal{F}$, $\dim S^G = V(G)$.*

It is worth noting that the general forest G is not covered by Reiner and Shimozono's result on percentage-avoiding diagrams [9].

To prove this theorem, we will show that $\dim S^G$ satisfies the properties of Proposition 2.5. While the first two properties are easy, the third will require a bit of work.

The first ingredient is the following lemma by James and Peel [7].

Given a diagram D , choose two boxes $(i_1, j_1), (i_2, j_2) \in D$ such that $(i_1, j_2), (i_2, j_1) \notin D$. We form the diagram D^A as follows. For every row except i_1 and i_2 , D^A contains exactly the same boxes as D . Also, D^A contains the box (i_1, j) if and only if D contains both (i_1, j) and (i_2, j) , while D^A contains the box (i_2, j) if and only if D contains either (i_1, j) or (i_2, j) .

Informally, D^A is obtained from D by shifting boxes from row i_1 to row i_2 if possible. Let us similarly define D^B using columns j_1 and j_2 . (We will refer to this process as *splitting* D into D^A and D^B using boxes (i_1, j_1) and (i_2, j_2) .) We then have the following result, which we present without proof.

Lemma 3.2 ([7]). *There exists a surjective $\mathbf{C}[\Sigma_n]$ -module homomorphism of Specht modules $\varphi: S^D \rightarrow S^{D^A}$. Moreover, $S^{D^B} \subset S^D$ and $\varphi(S^{D^B}) = 0$.*

This allows us to easily prove part of the leaf recurrence for $\dim S^G$.

Proposition 3.3. *Let G, G_1 , and G_2 be three forests as appearing in Proposition 2.4 (the leaf recurrence). Then $\dim S^G \geq \dim S^{G_1} + \dim S^{G_2}$.*

Proof. It is easy to check that the diagram of G splits into G_1 and G_2 . Then by Lemma 3.2, S^G contains a subrepresentation isomorphic to $S^{G_1} \oplus S^{G_2}$. The result follows. \square

The second ingredient will be a lower bound on $\dim S^G$ obtained by considering S^G as a representation of Σ_{n-1} instead of Σ_n . We first give a definition.

Definition. We say that a matching M of a bipartite graph G is *special* if there does not exist a cycle of G , half of whose edges lie in M .

Note that every matching in a forest is special since there are no cycles.

Definition. A subset U of boxes in a diagram D is called a *transversal* if no two boxes of U lie in the same row or column.

Clearly, a transversal of a diagram corresponds to a matching of the corresponding bipartite graph.

Proposition 3.4. *Let U be a transversal of a diagram D , and let E be the subdiagram of D given by the intersection of the rows and columns of D containing boxes*

in U . Then the following are equivalent:

- (1) The edges corresponding to U in the graph of D form a special matching.
- (2) The set U is the unique transversal in E of size $|U|$.
- (3) There exists a diagram equivalent to E such that the boxes corresponding to U lie along the main diagonal and all other boxes of E lie below the main diagonal.

Proof. For the first equivalence, let M be the matching corresponding to U in the graph G corresponding to D . If M is not special, then there exists a cycle, half of whose edges lie in M . Replacing those edges in M with the other edges in the cycle gives another matching M' , and the corresponding transversal U' lies in the same rows and columns of D as U . Conversely, suppose two subsets U and U' are both transversals lying in the same rows and columns. If M and M' are the corresponding matchings, then $(M \setminus M') \cup (M' \setminus M)$ consists of two disjoint matchings on the same set of vertices. It therefore has the same number of edges as vertices, so it contains a cycle. Since both M and M' are matchings, exactly half of the edges in the cycle must lie in each of M and M' , so M is not special.

For the second equivalence, suppose U is the unique transversal in E of size $|U|$. We may assume that U consists of the boxes (i, i) for $1 \leq i \leq u$. Consider the directed graph H on $[u]$ with an edge from i to j if $(i, j) \in E \setminus U$. In fact H is acyclic: if $(i_0, i_1, \dots, i_s = i_0)$ formed a cycle in H , then we could replace (i_j, i_j) in U by (i_j, i_{j+1}) for $0 \leq j < s$ and obtain another transversal in E of size $|U|$. It follows that we can reorder the vertices of H such that there is an edge from i to j only if $i > j$. Applying this reordering to the rows and columns of E gives an equivalent diagram with the desired property. Since the other direction of the equivalence is trivial, this completes the proof. \square

We will say that a transversal satisfying the conditions of Proposition 3.4 is *special*. The importance of special transversals comes from the following proposition.

Proposition 3.5. *Let U be a special transversal of a diagram D . Then the restricted representation $S^D|_{\Sigma_{n-1}}$ contains a subrepresentation isomorphic to $\bigoplus_{x \in U} S^{D \setminus \{x\}}$. In particular,*

$$\dim S^D \geq \sum_{x \in U} \dim S^{D \setminus \{x\}}.$$

Proof. By Proposition 3.4, we may assume that U consists of boxes (i, i) for $1 \leq i \leq u$ and that $(i, j) \notin D$ for $1 \leq i < j \leq u$. Write $D_i = D \setminus \{(i, i)\}$.

Recall that we interpret the basis elements of $\mathbf{C}[\Sigma_n]$ as tableaux of shape D with entries $1, \dots, n$. For $1 \leq i \leq u$, let T_i be the set of tableaux of shape D such that n appears in box (i, i) . By our choice of U , every term of $TC(D)R(D)$ for $T \in T_i$ contains the entry n in row at least i .

Let φ_i be the linear map on $\mathbf{C}[\Sigma_n]$ defined on a tableau T of shape D by $\varphi_i(T) = T$ if $T \in T_i$ and 0 otherwise. Since the action of Σ_{n-1} on tableaux in $\mathbf{C}[\Sigma_n]$ does not change the position of n , it follows that φ_i is a Σ_{n-1} -homomorphism.

Let V_i be the subrepresentation of $S^D|_{\Sigma_{n-1}}$ that is generated by $(T_i \cup T_{i+1} \cup \dots \cup T_u)C(D)R(D)$. By above, we have that for $1 \leq i \leq u$, $V_{i+1} \subset \ker \varphi_i$. We then have that $\varphi_i(V_i)$ is generated by $\varphi_i(T_i C(D)R(D)) = T_i C(D_i)R(D_i)$, so it is

naturally isomorphic to S^{D_i} . Therefore, the successive quotients of the filtration $0 \subset V_u \subset V_{u-1} \subset \dots \subset V_1$ contain $S^{D_u}, S^{D_{u-1}}, \dots, S^{D_1}$ in succession. The result follows easily. \square

Comparing Proposition 3.5 with Proposition 2.6, we obtain the following.

Proposition 3.6. *For all $G \in \mathcal{F}$, $\dim S^G \geq V(G)$.*

Proof. We induct on the number of edges in G . Let M be an almost perfect matching of G . Then since G is a forest, M is special, so by Proposition 3.5, the induction hypothesis, and Proposition 2.6,

$$\dim S^G \geq \sum_{e \in M} \dim S^{G \setminus e} \geq \sum_{e \in M} V(G \setminus e) = V(G). \quad \square$$

Armed with these facts, it is a simple matter to deduce Theorem 3.1.

Proof of Theorem 3.1. We wish to show that for any $G \in \mathcal{F}$, $\dim S^G = V(G)$. It suffices to show that $\dim S^G$ satisfies the three conditions given in Proposition 2.5. We induct on n , the number of edges of G , noting as before that if we prove the conditions for $n \leq k$, then $\dim S^G = V(G)$ for all graphs with at most k edges. Let us then assume that we have proven the claim for all graphs with fewer than n edges.

The first condition states that if T_n is the tree with n edges and white center vertex, then $\dim S^{T_n} = 1$. But the diagram of T_n is a row of n boxes, so its Specht module is the trivial representation, which has dimension 1. (This also proves the base case when $n = 1$.)

The second condition states that if G_1 and G_2 have m and $n - m$ edges, respectively, then the disjoint union $G_1 + G_2$ satisfies $\dim S^{G_1 + G_2} = \binom{n}{m} \cdot \dim S^{G_1} \cdot \dim S^{G_2}$. But as diagrams, $G_1 + G_2 = G_1 \oplus G_2$, so

$$\begin{aligned} \dim S^{G_1 + G_2} &= \dim \text{Ind}_{\Sigma_m \times \Sigma_{n-m}}^{\Sigma_n} (S^{G_1} \otimes S^{G_2}) \\ &= [\Sigma_n : \Sigma_m \times \Sigma_{n-m}] \cdot \dim(S^{G_1} \otimes S^{G_2}) \\ &= \binom{n}{m} \cdot \dim S^{G_1} \cdot \dim S^{G_2}. \end{aligned}$$

In particular, this implies by induction that $\dim S^G = V(G)$ for any disconnected graph G with n edges.

The third and final condition states that $\dim S^G$ satisfies the leaf recurrence. In other words, we need to show that for $G, G_1, G_2 \in \mathcal{F}$ related as in Proposition 2.4, $\dim S^G = \dim S^{G_1} + \dim S^{G_2}$. Since G is disconnected, we have from above that $\dim S^G = V(G)$. But by Propositions 3.3, 3.6, and 2.4, we have

$$V(G) = \dim S^G \geq \dim S^{G_1} + \dim S^{G_2} \geq V(G_1) + V(G_2) = V(G).$$

Therefore we must have equality everywhere. The result follows immediately. \square

Note that equality must hold everywhere in the proof above, and therefore we must also have equality in Propositions 3.3 and 3.6. This gives us the following two corollaries.

Corollary 3.7. *Let $G, G_1, G_2 \in \mathcal{F}$ be related as in Proposition 2.4 (the leaf recurrence). Then*

$$S^G \cong S^{G_1} \oplus S^{G_2}.$$

Corollary 3.8. *Let $G \in \mathcal{F}$, and let M be an almost perfect matching of G . Then*

$$S^G|_{\Sigma_{n-1}} \cong \bigoplus_{e \in M} S^{G \setminus e}.$$

Recall that for a partition λ , the restriction of the Specht module S^λ from Σ_n to Σ_{n-1} can be described as a direct sum of S^μ , where μ ranges over all partitions that can be obtained from λ by removing a corner box from the Young diagram of λ . Therefore Corollary 3.8 shows that the edges in an almost perfect matching of a forest are analogous to corner boxes of a Young diagram.

4. SYMMETRIC FUNCTIONS

In this section, we will associate to each forest G a symmetric function s_G . We will also show that the map sending G to s_G is, in a sense, the most general or universal function satisfying the leaf recurrence.

Recall that to a Young diagram λ , we may associate a symmetric function

$$s_\lambda(x_1, x_2, \dots) = \sum_T x^T,$$

where T ranges over all semistandard tableaux of shape λ , and $x^T = x_1^{\alpha_1} x_2^{\alpha_2} \dots$ if T contains α_i occurrences of i . The s_λ form an additive integer basis for the ring Λ of symmetric functions over the x_i with coefficients in \mathbf{Z} . Moreover, the multiplicative structure constants for the s_λ are given by the Littlewood-Richardson coefficients.

Due to the importance of symmetric functions to the representation theory of the symmetric group, we make the following definition.

Definition. Given a diagram D , the *Schur function* $s_D \in \Lambda$ is given by

$$s_D = \sum c_\lambda^D s_\lambda.$$

This can also be defined as the image of the character of S^D under the Frobenius characteristic map. When D is a skew Young diagram this definition coincides with the usual definition of a skew Schur function in terms of semistandard tableaux because the c_λ^D are again Littlewood-Richardson coefficients.

In the case when D corresponds to a forest, we can now show the following proposition.

Proposition 4.1. *There is a unique function $s: \mathcal{F} \rightarrow \Lambda$ that satisfies the following properties:*

- (1) *For the star T_n with n edges and white center vertex, $s(T_n) = h_n$, the n th complete homogeneous symmetric function.*
- (2) *For $G_1, G_2 \in \mathcal{F}$, $s(G_1 + G_2) = s(G_1)s(G_2)$.*
- (3) *The function s satisfies the leaf recurrence.*

In this case, $s(G) = s_G$, the Schur function associated to G .

Proof. The proof of uniqueness is identical to that of Proposition 2.5, so it suffices to check that the map $G \mapsto s_G$ satisfies these properties. The first follows because the Schur function corresponding to the partition (n) is h_n . The second follows because the analogous statement holds for partitions. The third follows immediately from Corollary 3.7. □

Proposition 4.1 essentially allows us to write s_G for $G \in \mathcal{F}$ in the basis $\{h_\lambda\}$ of Λ , where $h_\lambda = h_{\lambda_1} h_{\lambda_2} \cdots$. (This is similar to the *Jacobi-Trudi identity*; see, for instance, [11].)

We can even say something a bit stronger than Proposition 4.1: the map s is universal in the following sense.

Theorem 4.2. *Let R be a commutative ring with unit, and suppose that $\bar{f}: \{T_n \mid n > 0\} \rightarrow R$ is any function. Then f can be uniquely extended to a function $f: \mathcal{F} \rightarrow R$ that satisfies the following.*

- (1) For $G_1, G_2 \in \mathcal{F}$, $f(G_1 + G_2) = f(G_1)f(G_2)$.
- (2) The function f satisfies the leaf recurrence.

Moreover, f factors uniquely as $\varphi \circ s$, where $s: G \mapsto s_G$ and $\varphi: \Lambda \rightarrow R$ is a ring homomorphism.

Proof. Recall that the h_n are algebraically independent, so that $\Lambda = \mathbf{Z}[h_1, h_2, \dots]$. Then setting $\varphi(h_n) = \bar{f}(T_n)$ defines φ uniquely, and the claim follows easily from Proposition 4.1. □

For example, recall that the map $G \mapsto \text{vol } M_G$ satisfies the conditions of Theorem 4.2. It follows that this map corresponds to a ring homomorphism $\Lambda \rightarrow \mathbf{Q}$. Indeed, this is the map known as *exponential specialization*, which maps a homogeneous function $f \in \Lambda$ of degree n to the coefficient $\frac{1}{n!}[x_1 x_2 \cdots x_n]f$. (See, for instance, [11].)

As another example, the evaluation of a symmetric function $f(x_1, x_2, \dots)$ when N of the variables equal 1 and the rest are 0 is a polynomial function of N . We find that this has the following combinatorial interpretation.

Proposition 4.3. *Let $G \in \mathcal{F}$. For any nonnegative integer N , let $m_G(N)$ be the number of nonnegative integer edge labelings $w: E \rightarrow \mathbf{Z}$ such that for all $v \in V$,*

$$\sum_{e \ni v} w(e) \leq \begin{cases} N - 1, & \text{if } v \text{ is white, and} \\ N - \text{deg}(v), & \text{if } v \text{ is black.} \end{cases}$$

Then $m_G(N)$ is a polynomial in N , the map $G \mapsto m_G$ satisfies the conditions of Theorem 4.2, and $m_G(N) = s_G(\underbrace{1, 1, \dots, 1}_N)$.

Proof. First suppose $G = T_n$. Then $m_G(N)$ is the number of n -tuples of nonnegative integers summing to at most $N - 1$, and a simple counting argument shows that this is $\binom{n+N-1}{n} = h_n(\underbrace{1, \dots, 1}_N)$.

It suffices to check that the map $G \mapsto m_G$ satisfies the conditions of Theorem 4.2. The first condition is obvious. For the second, we mimic the proof of Proposition 2.4. Let $G, G_1, G_2, v_1, v'_1, v_2$, and v'_2 be as before, and let us assume without loss of generality that v_1 is white and v_2 is black. Let w be a suitable edge weighting for G with $w_i = w(\overline{v_i v'_i})$. If $w_1 \geq w_2$, then let z be the weighting of G_1 such that $z(\overline{v_1 v_2}) = w_2$, $z(\overline{v_1 v'_1}) = w_1 - w_2$, and $z(e) = w(e)$ for all other edges; if $w_1 < w_2$, let z' be the weighting of G_2 such that $z'(\overline{v_1 v_2}) = w_1$, $z'(\overline{v_2 v'_2}) = w_2 - w_1 - 1$, and $z'(e) = w(e)$ for all other edges. It is simple to check that this gives a bijection between suitable edge weightings for G and suitable edge weightings for either G_1 or G_2 . The result follows. □

The polynomial $m_G(N)$ counts lattice points in a way that is similar to the Ehrhart polynomial of the matching polytope. We shall see a representation-theoretic interpretation of $m_G(N)$ in the next section.

5. SCHUR MODULES

Just as $V(G)$ has an interpretation in terms of the dimension of a representation of the symmetric group, $m_G(N)$ has an interpretation as the dimension of a representation of the general linear group.

Given a diagram D with n boxes and at most N rows, we can define a $GL(N)(=GL(N, \mathbf{C}))$ -module as follows. Let $V = \mathbf{C}^N$ be the defining representation of $GL(N)$. Then, defining $C_D, R_D, C(D)$, and $R(D)$ as before, we define the Schur module \mathcal{S}^D by

$$\mathcal{S}^D = V^{\otimes n}C(D)R(D),$$

where we associate the copies of V in the tensor power to the boxes of the diagram D , and $C(D)$ and $R(D)$ act in the obvious way. Note that even if D has more than N rows, we can still make this definition, except that we may find $\mathcal{S}^D = 0$.

Let e_1, \dots, e_N be the standard basis of V . Then we may identify a tableau of shape D with entries at most N with the corresponding tensor product of standard basis elements of V , an element of $V^{\otimes n}$. Then \mathcal{S}^D is spanned by $TC(D)R(D)$, where T ranges over all tableaux of shape D with entries at most N .

When D is the Young diagram of a partition λ with at most N parts, \mathcal{S}^λ is the irreducible polynomial representation of $GL(N)$ with highest weight λ . In general, \mathcal{S}^D is a polynomial representation of $GL(N)$ of degree n , and it follows from Schur-Weyl duality [2] that

$$\mathcal{S}^D \cong \bigoplus (\mathcal{S}^\lambda)^{\oplus c_\lambda^D},$$

where the coefficients c_λ^D are the same as those encountered in the decomposition of the Specht module S^D .

Recall that the character of \mathcal{S}^D is the trace of the action of the diagonal matrix $\text{diag}(x_1, \dots, x_N)$ on \mathcal{S}^D as a polynomial in the x_i . Then it follows from the fact that $\text{ch}(\mathcal{S}^\lambda) = s_\lambda$ and the definition of s_D that

$$\text{ch}(\mathcal{S}^D) = s_D(x_1, \dots, x_N).$$

We therefore find the following.

Proposition 5.1. *Let $G \in \mathcal{F}$. Then $m_G(N) = \dim \mathcal{S}^G$.*

Proof. Follows from Proposition 4.3, since $\dim \mathcal{S}^G$ is the trace of the identity matrix in $GL(N)$. □

Recall that for Specht modules, Corollary 3.8 gives a description of the restriction of S^G from Σ_n to Σ_{n-1} by identifying a set of edges that act analogously to corner boxes of a Young diagram. We shall now obtain a similar result for the restriction of \mathcal{S}^G from $GL(N)$ to $GL(N-1)$ (where $GL(N-1) \subset GL(N)$ acts in the first $N-1$ dimensions and as the identity on the last). This amounts to finding certain subsets of edges that act analogously to horizontal strips of a Young diagram.

Let D be the diagram of a forest, and fix once and for all a transversal corresponding to an almost perfect matching U . Let us assume, as per Proposition 3.4, that U consists of the boxes (i, i) for $1 \leq i \leq u$ and that $(i, j) \notin D$ for $1 \leq i < j \leq u$

(so that the set U is also equipped with an ordering). We will say that such a diagram is in *standard form*.

Lemma 5.2. *Let D be the diagram of a forest with transversal U as above. Then the diagram D' obtained by removing column u from D has a transversal U' corresponding to an almost perfect matching that contains a box from each of the first $u - 1$ columns of D' (and possibly others).*

Proof. Let x_0 be the box (u, u) , and choose distinct boxes x_i and y_i recursively such that y_i lies in the same row as x_i (if such a box exists), and x_{i+1} lies in the same column as y_i and in U (if such a box exists). Since D is the diagram of a forest, this must terminate. Then simply replace the set of all x_i in U with the set of all y_i to get U' . \square

In the definition below, we will write D' for the diagram obtained from D by removing column u , and we associate to it the transversal U' as in the lemma. Also, let D'' be the diagram obtained from D by removing row and column u and associate to it the transversal $U'' = U \setminus \{(u, u)\}$.

Definition. A subset $Y \subset D$ is called a *horizontal strip* (with respect to U) if either $Y = \{(u, u)\} \cup Y'$, where Y' is a horizontal strip of D' with respect to U' , or $Y = Y''$, where Y'' is a horizontal strip of D'' with respect to U'' . (The empty set is a horizontal strip for any diagram D .)

Here we have abused notation slightly: D' may not be in standard form with respect to U' , but it is equivalent to some diagram E' that is in standard form with respect to the image of U' . By a horizontal strip of D' we mean the boxes of D' that correspond to a horizontal strip of E' .

It is easy to check that the horizontal strips of size 1 exactly correspond to elements of U , just as for partitions, a horizontal strip of length 1 is a corner box. It is also easy to check that a horizontal strip can have at most one box in any column but can have more than one box in a row.

Note that the set of horizontal strips of D can vary greatly according to the choices and orderings of almost perfect matchings. We will therefore fix once and for all a single such choice for each diagram D , yielding a single set $\mathcal{Y}(D)$ of horizontal strips.

Although the definition of a horizontal strip is a bit unusual, its value lies in the following result, the analogue of Proposition 2.6.

Proposition 5.3. *Let D be the diagram of a forest and $N > 1$ a positive integer. Then*

$$m_D(N) = \sum_{Y \in \mathcal{Y}(D)} m_{D \setminus Y}(N - 1).$$

Proof. Let $W_D(N)$ be the set of weightings of D as enumerated by $m_D(N)$ in Proposition 4.3. In other words, $W_D(N)$ is the number of labelings of the boxes of D with nonnegative integers such that the sum of each row is at most $N - 1$ and the sum of each column is at most N minus the number of boxes in that column.

Note that the conditions on any row or column after the first u are superfluous: since the diagonal corresponds to an almost perfect matching, any row after the first u contains at most one box, and any such box lies in one of the first u columns. Then the condition from the row states that this box is labeled at most $N - 1$, which

is automatically implied by the condition on its column. (The same argument holds for columns after the first u .)

We will construct a bijection $f: \bigcup_{Y \in \mathcal{Y}(D)} W_{D \setminus Y}(N - 1) \rightarrow W_D(N)$ inductively on the number of boxes of D . This bijection will satisfy the following properties: for $w \in W_{D \setminus Y}(N - 1)$,

- $f(w)$ is obtained from w by inserting boxes labeled 0 at the positions of Y and increasing the labels of certain other boxes of $D \setminus Y$ by 1;
- each of the first u columns will have either a box inserted or a box whose label is increased (not both);
- no row will contain two boxes with labels that are increased.

In particular, this bijection will increase the sum of any row or column by either 0 or 1, and, more specifically, the sum of any of the first u columns will either increase by 1, or else the sum will stay the same but the column will gain a box. Note that this will imply that the image of f is contained in $W_D(N)$.

For the forward direction, let $w \in W_{D \setminus Y}(N - 1)$. If $(u, u) \in Y$, then let z be the labeling of D such that $z(u, u) = 0$, $z|_{D'} = f(w|_{D' \setminus Y})$ by induction, and z and w agree elsewhere. If $(u, u) \notin Y$, then let z be such that $z(u, u) = w(u, u) + 1$, $z|_{D''} = f(w|_{D'' \setminus Y})$, and z and w agree elsewhere. It is easy to check by induction that all the conditions above are satisfied.

It remains to verify that f is a bijection. Choose $z \in W_D(N)$, and suppose $z(u, u) = 0$. Then if $f(w) = z$, then w must be unique. Indeed, first note that we must have $(u, u) \in Y$. Since $z|_{D'} \in W_{D'}(N)$, by induction we can find $f^{-1}(z|_{D'}) \in W_{D' \setminus Y'}(N - 1)$ for some horizontal strip Y' of D' . Then we must have $Y = \{(u, u)\} \cup Y'$, $w|_{D' \setminus Y'} = f^{-1}(z|_{D'})$, and w and z agree elsewhere on $D \setminus Y$, so w is uniquely determined. We then need to verify that such a w exists, that is, that it lies in $W_{D \setminus Y}(N - 1)$. Since $w|_{D' \setminus Y'} \in W_{D' \setminus Y'}(N - 1)$, the sum of any of the first u rows is at most $N - 2$. Similarly, the sum of any of the first $u - 1$ columns is at most $N - 1$ minus the length of the column. This also holds for column u , for the sum of this column in z was at most N minus the length, and while the sum in w is the same, we removed a box. Thus $w = f^{-1}(z)$ is well defined and unique in this case.

If $z(u, u) \neq 0$, then we similarly find that we must have $(u, u) \notin Y$ and $f^{-1}(z|_{D''}) \in W_{D'' \setminus Y''}(N - 1)$ for $Y'' = Y$, so w is determined by $w|_{D'' \setminus Y''} = f^{-1}(z|_{D''})$, $w(u, u) = z(u, u) - 1$, and w and z agree elsewhere. We again need to check that $w \in W_{D \setminus Y}(N - 1)$. Since $w|_{D'' \setminus Y''} \in W_{D'' \setminus Y''}(N - 1)$, the sum of any of the first u rows is at most $N - 2$ (even row u , for it was at most $N - 1$ and we subtracted 1 from (u, u)). Similarly, the sum of any of the first u columns is at most $N - 1$ minus the length of the column, for in each of these columns we either decreased a label by 1 or deleted a box (by the properties above). Thus $w = f^{-1}(z)$ is well defined and unique in this case as well. The result now follows from Proposition 4.3. \square

We may now use this result to prove the following analogue of Corollary 3.8.

Theorem 5.4. *Let D be the diagram of a forest, and let \mathcal{S}^D be the corresponding $GL(N)$ -module. Then*

$$\mathcal{S}^D|_{GL(N-1)} \cong \bigoplus_{Y \in \mathcal{Y}(D)} \mathcal{S}^{D \setminus Y}.$$

Proof. By Proposition 5.3, we have that the dimension of both sides of this equation are equal. It therefore suffices to show that the left side contains the right side as a submodule.

For a horizontal strip Y , let T_Y be the set of all tableaux T with boxes in Y labeled N and all other boxes labeled at most $N - 1$. Define the map φ_Y by $\varphi_Y(T) = T$ if $T \in T_Y$ and 0 otherwise. We may identify φ_Y with a linear map on $V^{\otimes n}$. Since $GL(N - 1)$ stabilizes e_1, \dots, e_{N-1} and fixes e_N , φ_Y is a $GL(N - 1)$ -homomorphism.

We construct a linear order \prec on $\mathcal{Y}(D)$ inductively on the number of boxes of D as follows:

- (1) If $(u, u) \in Y$ and $(u, u) \notin Z$, then $Z \prec Y$.
- (2) If $(u, u) \notin Y, Z$, then $Y'' = Y, Z'' = Z \in \mathcal{Y}(D'')$, so write $Z \prec Y$ if and only if $Z'' \prec Y''$.
- (3) If $(u, u) \in Y, Z$, then $Y' = Y \setminus \{(u, u)\}, Z' = Z \setminus \{(u, u)\} \in \mathcal{Y}(D')$, so write $Z \prec Y$ if and only if $Z' \prec Y'$.

Let $T \in T_Y$. We claim that \prec has the property that if $Z \prec Y$, then $TC(D)R(D)$ does not contain a term in T_Z . In other words, if $Z \prec Y$, then $Tqp \notin T_Z$ for $q \in C_D, p \in R_D$. To see this, we check that it holds for the three defining conditions above.

For (1), if $(u, u) \in Y$, then Tqp contains a box labeled N in either row u or column u . (When D is in standard form, all boxes in column u other than (u, u) are alone in their row.) Hence by the definition of horizontal strips, it cannot lie in T_Z if $(u, u) \notin Z$.

For (2), note that if $(u, u) \notin Y, Z$, then Y and Z do not contain any boxes in row u or column u . Thus we may assume that q and p do not affect row u and column u . If $Tqp \in T_Z$, then some term of $T|_{D''}\bar{q}\bar{p}$ lies in $T_{Z''}$, where \bar{q} and \bar{p} correspond to q and p in $C_{D''}$ and $R_{D''}$. Then we cannot have $Z'' \prec Y''$ and hence neither can we have $Z \prec Y$.

For (3), suppose $Tqp \in T_Z$. Note that Y and Z contain only one box in column u , namely (u, u) . Then q must send box (u, u) to itself, and p sends (u, u) to another box (u, v) . Let σ be the transposition in R_D switching (u, u) and (u, v) . Then by replacing p by $p\sigma$, we may assume that neither q nor p affects column u . Then restricting to D' as in the case above shows this case as well.

It follows that if $T \in T_Y$ and $Z \prec Y$, then $\varphi_Z(TC(D)R(D)) = 0$. Note that the analysis above also shows that if $Tqp \in T_Y$, then q acts as the identity on Y , so that we may associate q to an element of $C_{D \setminus Y}$. From this, it follows that $\varphi_Y(TC(D)R(D))$ is a scalar multiple of $TC(D \setminus Y)R(D \setminus Y)$ (by the order of the subgroup of R_D that stabilizes Y).

With this, we are ready to prove the result. Index the horizontal strips such that $Y_1 \prec Y_2 \prec \dots \prec Y_m$, and let V_i be the subrepresentation of $\mathcal{S}^D|_{GL(N-1)}$ generated by $(T_{Y_i} \cup T_{Y_{i+1}} \cup \dots \cup T_{Y_m})C(D)R(D)$. By the discussion above, $V_{i+1} \subset \ker \varphi_{Y_i}$, while $\varphi_{Y_i}(V_i)$ is generated by $\varphi_{Y_i}(TC(D)R(D))$ for $T \in T_{Y_i}$ and hence is naturally isomorphic to $\mathcal{S}^{D \setminus Y_i}$. Therefore the successive quotients of $0 \subset V_m \subset V_{m-1} \subset \dots \subset V_1$ contain $\mathcal{S}^{D \setminus Y_m}, \mathcal{S}^{D \setminus Y_{m-1}}, \dots, \mathcal{S}^{D \setminus Y_1}$ in succession. It follows from the complete reducibility of $GL(N - 1)$ -representations that we have the desired inclusion, and the result follows. \square

In fact, this shows that if for each $Y \in \mathcal{Y}(D)$, there exists a set of tableaux $SS(D \setminus Y, N - 1)$ such that $SS(D \setminus Y, N - 1)C(D \setminus Y)R(D \setminus Y)$ forms a basis of the

$GL(N - 1)$ -module $\mathcal{S}^{D \setminus Y}$, then we can construct a basis of the $GL(N)$ -module \mathcal{S}^D simply by adding boxes labeled N to the tableaux in $SS(D \setminus Y, N - 1)$ along Y for all Y . This leads us to the following definition.

Definition. Let D be the diagram of a forest, and let T be a tableau of shape D with entries at most N . We say that T is *semistandard* if the boxes labeled by N in T form a horizontal strip in $\mathcal{Y}(D)$, and the tableau formed from T by removing all boxes labeled N is also semistandard. (The empty tableau is semistandard by default.) We say that T is *standard* if it is semistandard and contains only the entries $1, \dots, n$, each exactly once.

(Note that since horizontal strips of length 1 on D precisely correspond to edges in an almost perfect matching of the graph of G , this definition agrees with the definition of standard labelings given at the end of Section 2.)

It is easy to see that a semistandard tableau is equivalent to a sequence of diagrams $\emptyset = D^{(0)} \subset D^{(1)} \subset \dots \subset D^{(N)} = D$ such that $D^{(i)} \setminus D^{(i-1)} \in \mathcal{Y}(D^{(i)})$ for $1 \leq i \leq N$. We denote the set of semistandard tableaux of shape D with entries at most N by $SS(D, N)$ and the set of all semistandard tableaux of shape D by $SS(D)$.

Proposition 5.5. *Let D be the diagram of a forest. Then*

$$s_D(x_1, x_2, \dots) = \sum_{T \in SS(D)} x_1^{\alpha_1} x_2^{\alpha_2} \dots,$$

where α_i is the number of occurrences of i in T . In particular, $\dim(S^D) = V(D)$ is the number of standard tableaux of shape D .

Proof. From the discussion after the proof of Theorem 5.4, we find that $TC(D)R(D)$, for $T \in SS(D, N)$, forms a basis for \mathcal{S}^D . Thus, since $s_D(x_1, \dots, x_N)$ is the character of \mathcal{S}^D , and $\text{diag}(x_1, \dots, x_N)$ acts on $TC(D)R(D)$ by multiplication by $x_1^{\alpha_1} x_2^{\alpha_2} \dots$, the result follows easily. \square

This therefore provides a combinatorial description of the coefficients of s_D .

6. CONCLUSION

In this paper, we have shown that there is a surprising relationship between matchings of a forest and the representation theory of the symmetric and general linear groups. In particular, we have shown how properties of the Specht and Schur modules of a forest can arise from considering its matchings and its matching polytope.

However, there is still one piece of this picture that remains. Although we have given a method of determining combinatorially the Schur function s_G for $G \in \mathcal{F}$, we have not yet given a subtraction-free combinatorial rule for the coefficients c_λ^G akin to the Littlewood-Richardson rule. In other words, we would like to identify c_λ^G as enumerating a certain type of combinatorial object. Since diagrams corresponding to forests are not in general “percentage-avoiding,” this would give a new class of diagrams for which we can determine c_λ^D that is not already covered by existing results.

Even though the proofs above are fairly short, in some sense they remain a bit mysterious. In particular, it seems almost coincidental that both $\dim S^G$ and $V(G)$ happen to satisfy the leaf recurrence. First, we can ask whether there are any other interesting functions that satisfy the leaf recurrence, and if so, what their role in this story is. We can also ask whether there is a more natural explanation for the relationship between Specht modules and matching polytopes presented here, and in particular whether it is possible to formulate similar results that are true for more general classes of diagrams.

If a more general picture does in fact exist, it will necessarily be more complicated than in the case of forests. For instance, it is not true in general that the normalized volume of the matching polytope of a bipartite graph is equal to the dimension of its Specht module. (As an example, if G is a cycle of length 4, which corresponds to the partition $(2, 2)$, then $\dim S^G = 2$ but $V(G) = 4$. For Young diagrams, the analogous polytope to consider seems not to be its matching polytope but rather the chain polytope on its poset of boxes.) One might also ask whether there exists a nice combinatorial interpretation for the volume of the matching polytope of a general bipartite graph or indeed of a general graph.

It is also not true that for a general diagram D one can define “corner boxes” and “horizontal strips” that satisfy the properties of Corollary 3.8 and Theorem 5.4. (The smallest counterexample corresponds to the cycle of length 6.) However, this is possible for many diagrams, including all northwest shapes [9]. It would therefore be interesting to classify those diagrams for which one can make such definitions.

ACKNOWLEDGMENTS

The author would like to thank Alexander Postnikov and Richard Stanley for their input in the development of this paper. This material is based upon work supported under a National Science Foundation Graduate Research Fellowship.

REFERENCES

- [1] Jack Edmonds. Maximum matching and a polyhedron with 0,1-vertices. *J. Res. Nat. Bur. Standards Sect. B*, 69B:125–130, 1965. MR0183532 (32:1012)
- [2] William Fulton and Joe Harris. *Representation theory*, volume 129 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991. A first course, Readings in Mathematics. MR1153249 (93a:20069)
- [3] G. D. James. A characteristic-free approach to the representation theory of \mathfrak{S}_n . *J. Algebra*, 46(2):430–450, 1977. MR0439924 (55:12805)
- [4] Bernhard Korte and Jens Vygen. *Combinatorial optimization*, volume 21 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, fourth edition, 2008. Theory and algorithms. MR2369759 (2008j:90002)
- [5] L. Lovász and M. D. Plummer. *Matching theory*, volume 121 of *North-Holland Mathematics Studies*. North-Holland Publishing Co., Amsterdam, 1986. Annals of Discrete Mathematics, 29. MR859549 (88b:90087)
- [6] Peter Magyar. Borel-Weil theorem for configuration varieties and Schur modules. *Adv. Math.*, 134(2):328–366, 1998. MR1617793 (2000e:14087)
- [7] M. H. Peel and G. D. James. Specht series for skew representations of symmetric groups. *J. Algebra*, 56(2):343–364, 1979. MR528580 (80h:20021)
- [8] Victor Reiner and Mark Shimozono. Specht series for column-convex diagrams. *J. Algebra*, 174(2):489–522, 1995. MR1334221 (96m:20020)
- [9] Victor Reiner and Mark Shimozono. Percentage-avoiding, northwest shapes and peelable tableaux. *J. Combin. Theory Ser. A*, 82(1):1–73, 1998. MR1616579 (2000a:05220)

- [10] Bruce E. Sagan. *The symmetric group: representations, combinatorial algorithms, and symmetric functions*, volume 203 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 2001. MR1824028 (2001m:05261)
- [11] Richard P. Stanley. *Enumerative combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999. With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin.

DEPARTMENT OF MATHEMATICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE,
MASSACHUSETTS 02139

E-mail address: `riliu@math.mit.edu`

Current address: Department of Mathematics, University of Michigan, Ann Arbor, Michigan
48109

E-mail address: `riliu@umich.edu`