

PERMUTED COMPOSITION TABLEAUX, 0-HECKE ALGEBRA AND Labeled BINARY TREES

V. TEWARI AND S. VAN WILLIGENBURG

ABSTRACT. We introduce a generalization of semistandard composition tableaux called permuted composition tableaux. These tableaux are intimately related to permuted basement semistandard augmented fillings studied by Haglund, Mason and Remmel. Our primary motivation for studying permuted composition tableaux is to enumerate all possible ordered pairs of permutations (σ_1, σ_2) that can be obtained by standardizing the entries in two adjacent columns of an arbitrary composition tableau. We refer to such pairs as compatible pairs. To study compatible pairs in depth, we define a 0-Hecke action on permuted composition tableaux. This action naturally defines an equivalence relation on these tableaux. Certain distinguished representatives of the resulting equivalence classes in the special case of two-columned tableaux are in bijection with compatible pairs. We provide a bijection between two-columned tableaux and labeled binary trees. This bijection maps a quadruple of descent statistics for 2-columned tableaux to left and right ascent-descent statistics on labeled binary trees introduced by Gessel, and we use it to prove that the number of compatible pairs is $(n + 1)^{n-1}$.

CONTENTS

1. Introduction	2
2. Preliminaries	3
2.1. The symmetric group \mathfrak{S}_n	4
2.2. The 0-Hecke algebra $H_n(0)$	4
2.3. Compositions, partitions and diagrams	5
2.4. Tableaux	5
3. 0-Hecke modules on SPCTs	9
3.1. Source and sink tableaux	11
4. Descents in 2-columned tableaux and ascent-descents on labeled binary trees	13
4.1. Dyck paths, unlabeled and labeled	14

2010 *Mathematics Subject Classification*. Primary 05E10, 20C08; Secondary 05A05, 05A19, 05C20, 05E05, 05E15.

Key words and phrases. 0-Hecke algebra, allowable pair, composition tableau, descent, Dyck path, labeled tree, pattern avoidance, reverse tableau.

The second author was supported in part by the National Sciences and Engineering Research Council of Canada.

Email: vasut@math.washington.edu, Dept. of Mathematics, University of Washington, Seattle, WA 98195, USA, Telephone: +1-206-543-9165.

4.2. Bijection between \mathcal{LD}_n and $\text{SPCT}((2^n))$	15
4.3. Plane binary trees, unlabeled and labeled	16
4.4. Labeled Dyck paths and labeled binary trees	18
5. Allowable pairs of permutations and acyclic graphs	24
6. Final remarks	30
Acknowledgements	30
References	30

1. INTRODUCTION

Composition tableaux were introduced in [19] to define the basis of quasisymmetric Schur functions for the Hopf algebra of quasisymmetric functions. These functions are analogues of the ubiquitous Schur functions [41], have been studied in substantial detail recently [8, 19, 20, 30, 31, 45], and have consequently been the genesis of an active new branch of algebraic combinatorics discovering Schur-like bases in quasisymmetric functions [1, 5], type B quasisymmetric Schur functions [27, 37], quasi-key polynomials [2, 42] and quasisymmetric Grothendieck polynomials [35]. Just as Young tableaux play a crucial role in the combinatorics of Schur functions [40, 43], composition tableaux are key to understanding the combinatorics of quasisymmetric Schur functions.

The aim of this article is to shed light on certain enumerative aspects of composition tableaux by studying actions of the 0-Hecke algebra of type A . Recall that the number of standard Young tableaux is given by the famous hook-length formula of Frame-Robinson-Thrall [15]. While there is no known ‘hook-length formula’ for enumerating standard composition tableaux, and the presence of large prime factors in the data suggests there might not be a simple closed formula, we hope that our article convinces the reader of other interesting enumerative aspects of composition tableaux and motivates further exploration.

The following question serves as the primary motivation for this article: Can we characterize and enumerate the possible ordered pairs of permutations (σ_1, σ_2) obtained by considering the relative order of entries in two adjacent columns of an arbitrary composition tableau? This question in the context of Young tableaux is uninteresting. On the other hand, it is natural with regards to composition tableaux as the relative order of entries in, say, column i of a composition tableau τ is governed by the relative order of the entries in column $i - 1$. This is because of the so-called triple condition/rule [19, Definition 4.1]. If the columns under consideration have the same number of entries, say n , then we say that the resulting pair (σ_1, σ_2) is a *compatible pair* of permutations in \mathfrak{S}_n . We remark here that the assumption that the two columns possess the same number of entries is mild as this can always be achieved by supplementing the composition tableau under consideration with 0s as suggested in [19, Definition 4.1 part 3)]. One of our main results is that the number of compatible pairs of permutations in \mathfrak{S}_n is $(n + 1)^{n-1}$. To study the combinatorics of an arbitrary pair of adjacent

columns in a composition tableau, we introduce a generalization thereof that we call permuted composition tableaux. Just as semistandard composition tableaux are intimately tied to semistandard augmented fillings [19], our permuted composition tableaux are connected to permuted basement semistandard augmented fillings introduced by Haglund-Mason-Remmel [21]. Given our interest in the relative order of entries in columns, the authors' previous work [45] hints at constructing an action of the 0-Hecke algebra $H_n(0)$ on permuted composition tableaux. To this end, we generalize the operators from [45] to our current setting and also obtain analogues of the results therein.

To answer our original question, we focus on standard permuted composition tableaux of shape (2^n) . We refer to such tableaux as 2-columned tableaux. The 0-Hecke action that we construct naturally leads to a deeper study of descents in the first column in 2-columned tableaux and they naturally come in four flavors. We construct a bijection between the set of 2-columned tableaux and labeled binary trees that maps certain descent statistics on tableaux to ascent-descent statistics on labeled binary trees. Note that the ascent-descent statistics on labeled binary trees were first studied in unpublished work by Gessel, and functional equations for their distribution were established by Kalikow [28] and Drake [12]. Given the striking observations of Gessel [16] relating the distribution of these statistics to enumerative questions in the theory of hyperplane arrangements, much work has been done recently [7, 11, 14, 17, 44]. That said, we will not focus on the hyperplane arrangements perspective in this article.

Outline of the article. The paper essentially has two halves: the first half focuses on studying the combinatorics of permuted composition tableaux and defining a 0-Hecke action on them. The second half concerns itself with enumerating and characterizing compatible pairs by way of Dyck paths and binary trees.

In Section 2, we introduce our main combinatorial objects, and develop most of the notation we need. Our main result here is Theorem 2.4 that makes explicit the link between reverse tableaux and permuted composition tableaux. In Section 3, we construct a 0-Hecke action on the set of permuted composition tableaux of a given shape. Section 4 studies certain descents in 2-columned tableaux and relates them to ascents-descents in labeled binary trees via bijections that use labeled Dyck paths as intermediate objects. Our main result in this section is Theorem 4.8. Section 5 gives a characterization for compatible pairs in terms of pattern-avoidance. We demonstrate that compatible pairs are essentially allowable pairs introduced in [4] and investigated further in [3, 18, 22]. To answer our original question, we show in Corollary 5.8 that every allowable pair can be obtained by standardizing the entries in the last two columns of some standard composition tableau.

2. PRELIMINARIES

We denote the set of positive integers by \mathbb{N} . Given $n \in \mathbb{N}$, we define $[n]$ to be the set of the first n positive integers $\{1, \dots, n\}$. The set of all words in the alphabet \mathbb{N} is denoted by \mathbb{N}^* . The empty word is denoted by ε . A *ud-word* is a word in the alphabet $\{u, d\}$. A related notion is that of a *labeled ud-word*, which is a word in the alphabet $\{u_i, d_i \mid i \in \mathbb{N}\}$. There

is a natural projection χ from the set of labeled ud-words to the set of ud-words defined by mapping u_i to u and d_i to d for all $i \in \mathbb{N}$.

2.1. The symmetric group \mathfrak{S}_n . The symmetric group \mathfrak{S}_n is generated by the elements s_1, \dots, s_{n-1} subject to the following relations

$$\begin{aligned} s_i^2 &= 1 \text{ for } 1 \leq i \leq n-1 \\ s_i s_{i+1} s_i &= s_{i+1} s_i s_{i+1} \text{ for } 1 \leq i \leq n-2 \\ s_i s_j &= s_j s_i \text{ if } |i-j| \geq 2. \end{aligned}$$

In the first relation above, the 1 denotes the identity element in \mathfrak{S}_n . In practice, one identifies \mathfrak{S}_n with the group of all bijections from $[n]$ to itself, otherwise known as *permutations*, by setting s_i to be the *simple transposition* interchanging i and $i+1$ for $1 \leq i \leq n-1$. An expression for $\sigma \in \mathfrak{S}_n$ of the form $s_{i_1} \cdots s_{i_p}$ that uses the minimal number of simple transpositions is called a *reduced word* for σ .

We write permutations in one-line notation and, on occasion, treat them as words in \mathbb{N}^* . Given a word $w = w_1 \cdots w_n \in \mathbb{N}^*$, we define the *standardization* of w , denoted by $\text{stan}(w)$, to be the unique permutation $\sigma \in \mathfrak{S}_n$ such that $\sigma(i) > \sigma(j)$ if and only if $w_i > w_j$ for $1 \leq i < j \leq n$. For example, $\text{stan}(3122) = 4123$. An *inversion* of $\sigma \in \mathfrak{S}_n$ is an ordered pair (p, q) such that $1 \leq p < q \leq n$ and $\sigma(p) > \sigma(q)$. The set of inversions in σ is denoted by $\text{Inv}(\sigma)$. This given, define the (left) *weak Bruhat order* \leq_L on \mathfrak{S}_n by $\sigma_1 \leq_L \sigma_2$ if and only if $\text{Inv}(\sigma_1) \subseteq \text{Inv}(\sigma_2)$ [9, Proposition 3.1.3]. We denote the cover relation in the weak Bruhat order by \prec_L . The symmetric group \mathfrak{S}_n endowed with the weak Bruhat order \leq_L inherits the structure of a graded lattice with a unique minimum element (given by the identity permutation) and a unique maximum element (given by the reverse of the identity permutation). To emphasize the dependence on n , the identity permutation in \mathfrak{S}_n and its reverse will henceforth be denoted by ϵ_n and $\bar{\epsilon}_n$ respectively.

2.2. The 0-Hecke algebra $H_n(0)$. The 0-Hecke algebra $H_n(0)$ is the \mathbb{C} -algebra generated by the elements T_1, \dots, T_{n-1} subject to the following relations

$$\begin{aligned} T_i^2 &= T_i \text{ for } 1 \leq i \leq n-1 \\ T_i T_{i+1} T_i &= T_{i+1} T_i T_{i+1} \text{ for } 1 \leq i \leq n-2 \\ T_i T_j &= T_j T_i \text{ if } |i-j| \geq 2. \end{aligned}$$

If $s_{i_1} \cdots s_{i_p}$ is a reduced word for a permutation $\sigma \in \mathfrak{S}_n$, then we define an element $T_\sigma \in H_n(0)$ as

$$T_\sigma = T_{i_1} \cdots T_{i_p}.$$

The Word Property [9, Theorem 3.3.1] of \mathfrak{S}_n implies that T_σ is independent of the choice of reduced word. Moreover, the set $\{T_\sigma \mid \sigma \in \mathfrak{S}_n\}$ is a linear basis for $H_n(0)$. Thus, the dimension of $H_n(0)$ is $n!$. In general, $H_n(0)$ is not semisimple and possesses a rich combinatorial representation theory [10, 36]. It is intimately related with the Hopf algebras of quasisymmetric functions and noncommutative symmetric functions respectively [13], in

the same way as symmetric group representation theory is intimately connected with the Hopf algebra of symmetric functions. More information about $H_n(0)$ and its representations can be found in [10, 34], and contemporary results can be found in [6, 23, 24, 25, 26, 29]. Our interest in the 0-Hecke algebra stems from the authors' previous work in the context of providing a representation-theoretic interpretation for quasisymmetric Schur functions [45].

2.3. Compositions, partitions and diagrams. A *composition* α of a positive integer n is a finite ordered list of positive integers $(\alpha_1, \dots, \alpha_k)$ satisfying $\sum_{i=1}^k \alpha_i = n$. The α_i are called the *parts* of α and their sum is called the *size* of α (denoted by $|\alpha|$). The number of parts is called the *length* of α and is denoted by $\ell(\alpha)$, and we often denote m consecutive parts j by j^m . The *empty composition*, denoted by \emptyset , is the unique composition of size and length 0. If α is a composition of size n , we denote this by $\alpha \vDash n$. From $\alpha = (\alpha_1, \dots, \alpha_k) \vDash n$, we can obtain another composition $\hat{\alpha} \vDash 2n$ of length n by first incrementing each part of α by 1 and subsequently adding $n - k$ parts equaling 1 at the end. For instance, if $\alpha = (2, 1, 3)$, then $\hat{\alpha} = (3, 2, 4, 1, 1, 1)$. If $\alpha = (\alpha_1, \dots, \alpha_k) \vDash n$ is such that $\alpha_1 \geq \dots \geq \alpha_k$, then we say that α is a *partition* of n and denote this by $\alpha \vdash n$. The partition obtained by sorting the parts of α in weakly decreasing order is denoted by $\tilde{\alpha}$.

We depict a composition $\alpha = (\alpha_1, \dots, \alpha_k) \vDash n$ by its *composition diagram*, also referred to as α , which is a left-justified array of n cells where the i -th row from the top has α_i cells. The *augmented composition diagram* of α , denoted by $\hat{\alpha}$, is the composition diagram of $\hat{\alpha}$. We refer to the first column of $\hat{\alpha}$ as its *basement*. Note that if α is a partition, then the composition diagram of α is the Young diagram of α in English notation.

2.4. Tableaux. Given a partition $\lambda \vdash n$, a *reverse tableau* (henceforth abbreviated to RT) T of *shape* λ is a filling of the Young diagram of λ with positive integers such that the rows decrease weakly from left to right, whereas the columns decrease strictly from top to bottom. Let $\text{RT}(\lambda)$ denote the set of reverse tableaux of shape λ all of whose entries are weakly less than $|\lambda|$. If $T \in \text{RT}(\lambda)$ is such that its entries are all distinct, then we say that T is *standard*. We denote the set of standard RTs of shape λ by $\text{SRT}(\lambda)$. Additionally, we refer to a standard RT as an SRT.

Given $\alpha \vDash n$ and $\sigma \in \mathfrak{S}_{\ell(\alpha)}$, we define a *permuted composition tableau* (henceforth abbreviated to PCT) of *shape* α and *type* σ to be a filling τ of the composition diagram α with positive integers such that the following conditions hold.

- (1) The entries in the first column are all distinct and the standardization of the word obtained by reading the first column from top to bottom is σ .
- (2) The entries along the rows decrease weakly when read from left to right.
- (3) For any configuration in τ of the type in Figure 1, if $a \geq c$ then $b > c$. We call this condition the *triple condition*.

The first condition and the triple condition guarantee that the entries in any given column of a PCT are all distinct. Figure 2 gives a PCT of shape $(1, 3, 2, 4)$ and type 1324.




FIGURE 1. A triple configuration.

Remark 2.1. Note that our version of the triple condition follows [31, Definition 4.2.6]. The version in [19, Definition 4.1 part 3]), though stated differently, is equivalent. It is the latter that we employ in Section 5.

Let $\text{PCT}^\sigma(\alpha)$ denote the set of all PCTs of shape α and type σ all of whose entries are weakly less than $|\alpha|$. Additionally, let

$$\text{PCT}(\alpha) := \coprod_{\sigma \in \mathfrak{S}_{\ell(\alpha)}} \text{PCT}^\sigma(\alpha),$$

where \coprod indicates disjoint union. If $\tau \in \text{PCT}^\sigma(\alpha)$ is such that its entries are all distinct, then we say that τ is *standard*. We denote the set of standard PCTs of shape α and type σ by $\text{SPCT}^\sigma(\alpha)$. Additionally, we refer to a standard PCT as an SPCT. Finally, let

$$\text{SPCT}(\alpha) := \coprod_{\sigma \in \mathfrak{S}_{\ell(\alpha)}} \text{SPCT}^\sigma(\alpha).$$

Remark 2.2. From this point on, whenever we write $\text{PCT}^\sigma(\alpha)$ or $\text{SPCT}^\sigma(\alpha)$ without explicitly specifying σ , it is implicit that σ is a permutation in $\mathfrak{S}_{\ell(\alpha)}$. Furthermore, all the RTs (respectively PCTs) we consider in this article belong to $\text{RT}(\lambda)$ (respectively $\text{PCT}(\alpha)$) for the appropriate λ (respectively α).

FIGURE 2. A PCT (left) and an SPCT (right) of shape $(1, 3, 2, 4)$ and type 1324.

Given $\alpha \vDash n$ and $\tau \in \text{SPCT}(\alpha)$, we say that an integer $1 \leq i \leq n - 1$ is a *descent* of τ if $i + 1$ lies weakly right of i in τ . The *descent set* of τ , denoted by $\text{Des}(\tau)$, consists of all descents in τ . For instance, the descent set of the SPCT in Figure 2 is $\{1, 2, 4, 6, 7\}$.

Remark 2.3. Note that if σ equals $\epsilon_{\ell(\alpha)}$, then a PCT of shape α and type σ is in fact a *composition tableau* (abbreviated henceforth to CT) introduced in [19, Definition 4.1] of the same shape α . Similarly, an SPCT of shape α and type $\sigma = \epsilon_{\ell(\alpha)}$ corresponds to a *standard* CT (abbreviated henceforth to SCT) of shape α .

Let $\alpha \vDash n$ whose largest part is α_{max} and let $\tau \in \text{PCT}^\sigma(\alpha)$. Suppose that the entries in column i for $1 \leq i \leq \alpha_{max}$ read from top to bottom form some word w^i in \mathbb{N}^* . We refer to w^i as the i -th column word of τ . Furthermore, we define the *standardized i -th column word* of τ , denoted by $\text{st}_i(\tau)$, to be $\text{stan}(w^i)$. Note that $\text{st}_1(\tau)$ is σ since $\tau \in \text{PCT}^\sigma(\alpha)$. We define the *standardized column word* of τ , denoted by $\text{st}(\tau)$, to be the word

$$\text{st}(\tau) = \text{st}_1(\tau) \text{st}_2(\tau) \cdots \text{st}_{\alpha_{max}}(\tau).$$

For the PCT in Figure 2, the standardized column word is 1324 213 12 1.

We now discuss two procedures connecting RTs and PCTs. The reader may interpret this correspondence as the analogue to that between CTs and semistandard augmented fillings [33]. Consider a composition α and permutation $\sigma \in \mathfrak{S}_{\ell(\alpha)}$. Let α_{max} be the largest part of α . Our first procedure, $\bar{\rho}_\sigma$, takes $\tau \in \text{PCT}^\sigma(\alpha)$ as input and outputs a filling T of shape $\lambda := \tilde{\alpha}$ by considering the entries of column i of τ in decreasing order and putting them in column i of λ from top to bottom for all $1 \leq i \leq \alpha_{max}$.

For our second procedure $\bar{\phi}_\sigma$, let λ be a partition and let $\sigma \in \mathfrak{S}_{\ell(\lambda)}$. This procedure takes $T \in \text{RT}(\lambda)$ as input and outputs a filling τ as follows.

- (1) Consider the entries in the first column of T and write them in rows $1, 2, \dots, \ell(\lambda)$ so that the standardization of the word obtained by reading from top to bottom is σ .
- (2) Consider the entries in column 2 in decreasing order and place each of them in the row with the smallest index so that the cell to the immediate left of the number being placed is filled and the row entries weakly decrease when read from left to right.
- (3) Repeat the previous step with the set of entries in column k for $k = 3, \dots, \lambda_1$.

Figure 3 illustrates the procedures just introduced and motivates the theorem that follows.




FIGURE 3. An RT and a PCT related via the maps $\bar{\rho}_\sigma$ and $\bar{\phi}_\sigma$ where $\sigma = 3142$.

Theorem 2.4. *Let λ be a partition and let $\sigma \in \mathfrak{S}_{\ell(\lambda)}$. The map*

$$\bar{\rho}_\sigma : \coprod_{\tilde{\alpha}=\lambda} \text{PCT}^\sigma(\alpha) \rightarrow \text{RT}(\lambda)$$

is a bijection and its inverse is the map $\bar{\phi}_\sigma$.

Proof. First, we show that $\bar{\rho}_\sigma$ and $\bar{\phi}_\sigma$ are both injections, which suffices to conclude that they are in fact bijections. To this end, we will use *permuted basement semistandard augmented fillings* (henceforth abbreviated to PBFs) introduced in [21] generalizing the notion of semi-standard augmented fillings from [32]. We refer the reader to [21] for details on

the terminology used in our proof. The reader will benefit from referring to Figure 4 while reading this proof.

Let $\lambda \vdash n$ and let $k = \ell(\lambda)$. Consider $\tau \in \text{PCT}^\sigma(\alpha)$ where $\tilde{\alpha} = \lambda$. Suppose that the first column word of τ is $w = w_1 \cdots w_k$. Clearly $\text{stan}(w) = \sigma$. Let $\hat{\sigma} \in \mathfrak{S}_n$ be the permutation obtained by concatenating the entries in $[n] \setminus \{w_1, \dots, w_k\}$ at the end of w in increasing order. Consider the filling $\hat{\tau}$ of the augmented composition diagram $\hat{\alpha}$ constructed as follows.

- The basement contains $\hat{\sigma}(i)$ for $i = 1$ through n from top to bottom.
- The rest of the diagram, which is essentially the composition diagram of α , is filled exactly as τ .

Given its construction, one may check that $\hat{\tau}$ is a PBF with basement permutation $\hat{\sigma}$. Indeed, the triple condition satisfied by τ ensures that, in $\hat{\tau}$, all type A and B triples are inversion triples and the B -increasing condition is satisfied. We omit the details. Crucial for us is the fact that the association of $\hat{\tau}$ to τ is one-to-one. Next, we describe a map that allows us to associate a reverse tableau with $\hat{\tau}$.

Given a permutation π , Haglund-Mason-Remmel [21, Section 4] define a map ρ_π that generalizes Mason's shift map [32] (also known as the ρ map). This map takes as input a PBF T_1 with basement permutation π and outputs the unique PBF T_2 such that the entries in any column of T_2 read from top to bottom are the entries in the corresponding column of T_1 read in decreasing order. Note that the basement permutation of T_2 is $\bar{\epsilon}_n$. Haglund-Mason-Remmel (see discussion after [21, Corollary 9]) establish that T_2 is in fact a reverse tableau. The map ρ_π^{-1} [21, Page 309] is the same as our map $\bar{\phi}_\pi$. In Figure 4, the two fillings in the middle are PBFs related by ρ_π where $\pi = 10\ 2\ 11\ 9\ 1\ 3\ 4\ 5\ 6\ 7\ 8$ in one-line notation.

In light of the preceding discussion, using the map $\rho_{\hat{\sigma}}$ we can associate the reverse tableau $\hat{T} := \rho_{\hat{\sigma}}(\hat{\tau})$ with $\hat{\tau}$. Furthermore, given how $\rho_{\hat{\sigma}}$ operates, we conclude that the reverse tableau T obtained by omitting the basement $\bar{\epsilon}_n$ from \hat{T} may be otherwise obtained by sorting the entries in individual columns of τ in decreasing order and writing them along columns in the Young diagram of $\lambda = \tilde{\alpha}$. Thus, we have $\bar{\rho}_\sigma(\tau) = T$.

The injectivity of $\bar{\rho}_\sigma$ is explained next. Suppose that τ_1 and τ_2 satisfy $\bar{\rho}_\sigma(\tau_1) = \bar{\rho}_\sigma(\tau_2)$. Let $\hat{\sigma}_1, \hat{\sigma}_2 \in \mathfrak{S}_n$ be obtained by extending the first column words in τ_1 and τ_2 respectively as before. Our hypothesis implies that $\rho_{\hat{\sigma}_1}(\hat{\tau}_1) = \rho_{\hat{\sigma}_2}(\hat{\tau}_2)$. Thus, we have that the set of entries in the corresponding columns of τ_1 and τ_2 are the same. Since the standardized first column word of both τ_1 and τ_2 is σ , the first column of τ_1 is the same as the first column of τ_2 . This implies $\hat{\sigma}_1 = \hat{\sigma}_2$, which, in view of [21, Theorem 10 part 2)], implies that $\hat{\tau}_1 = \hat{\tau}_2$. It follows immediately that $\tau_1 = \tau_2$. Thus, $\bar{\rho}_\sigma$ is an injection from $\coprod_{\tilde{\alpha}=\lambda} \text{PCT}^\sigma(\alpha) \rightarrow \text{RT}(\lambda)$.

To establish the injectivity of $\bar{\phi}_\sigma$, we again use results in [21], beginning with an alternative description of $\bar{\phi}_\sigma$. Let $T \in \text{RT}(\lambda)$. Construct a PBF \hat{T} of shape $\hat{\lambda}$ with basement permutation $\bar{\epsilon}_n$ by filling the remaining diagram, essentially the composition diagram of λ , exactly as T . To obtain $\bar{\phi}_\sigma(T)$, construct a permutation $\hat{\sigma}$ such that

- its first k letters are a rearrangement of the entries in the first column of T ,
- the standardization of the word formed by the first k letters is σ ,

- the last $n - k$ entries increase when read from left to right.

Then $\hat{\tau} := \rho_{\hat{\sigma}}^{-1}(\hat{T}) = \bar{\phi}_{\hat{\sigma}}(\hat{T})$ is a PBF with basement permutation $\hat{\sigma}$. Given how $\hat{\sigma}$ relates to σ and to the first column of T , we infer that the first k rows in the second column of $\hat{\tau}$ are the same as the first k rows of the first column. This implies that the filling obtained by removing the basement from $\hat{\tau}$ is in fact $\bar{\phi}_{\sigma}(T)$.

To establish the injectivity of $\bar{\phi}_{\sigma}$, argue as follows. Suppose $T_1, T_2 \in \text{RT}(\lambda)$ are such that $\bar{\phi}_{\sigma}(T_1) = \bar{\phi}_{\sigma}(T_2)$. The preceding discussion implies that there exist permutations $\hat{\sigma}_1, \hat{\sigma}_2 \in \mathfrak{S}_n$ such that $\rho_{\hat{\sigma}_1}^{-1}(\hat{T}_1) = \rho_{\hat{\sigma}_2}^{-1}(\hat{T}_2)$, where \hat{T}_1 and \hat{T}_2 are obtained by appending the basement permutation $\bar{\epsilon}_n$ as before. Again, we conclude that the sets of entries in corresponding columns of T_1 and T_2 are equal, which implies that $T_1 = T_2$.

To finish the proof, we need to establish that $\bar{\phi}_{\sigma}$ and $\bar{\rho}_{\sigma}$ are mutually inverse. This follows from our alternate description for $\bar{\rho}_{\sigma}$ and $\bar{\phi}_{\sigma}$ in terms of $\rho_{\hat{\sigma}}$ and $\rho_{\hat{\sigma}}^{-1}$ for an appropriately constructed $\hat{\sigma}$. \square

Figure 4 describes the various steps in the proof of Theorem 2.4 in the case of the RT and PCT from Figure 3. Let $\sigma = 3142$, $\alpha = (4, 1, 3, 3)$ and $\lambda = (4, 3, 3, 1)$. The PCT τ on the left belongs to $\text{PCT}^{\sigma}(\alpha)$. The second filling from the left is the PBF $\hat{\tau}$ with basement permutation $\hat{\sigma}$ constructed in the preceding proof. We have $\hat{\sigma} = 10\ 2\ 11\ 9\ 1\ 3\ 4\ 5\ 6\ 7\ 8$ in one-line notation. The second filling from the right is $\rho_{\hat{\sigma}}(\hat{\tau})$ and the rightmost filling is $\bar{\rho}_{\sigma}(\tau)$.

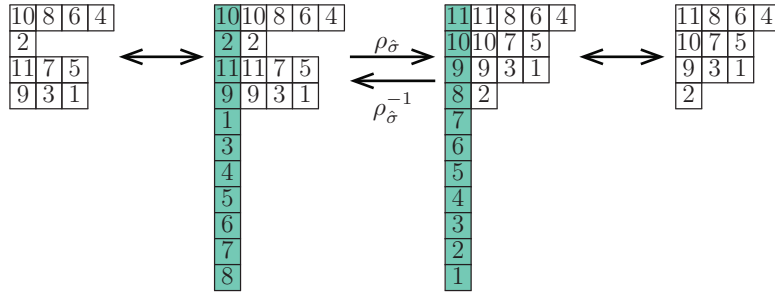


FIGURE 4. The bijection between $\coprod_{\tilde{\alpha}=\lambda} \text{PCT}^{\sigma}(\alpha)$ and $\text{RT}(\lambda)$ via PBFs.

We conclude this section by noting that if $\sigma = \epsilon_{\ell(\alpha)}$, then the map $\bar{\rho}_{\sigma}$ maps a CT τ of shape α to an RT of shape $\tilde{\alpha}$.

3. 0-HECKE MODULES ON SPCTS

In order to define an $H_n(0)$ -action on SPCTs we recall the concept of attacking defined in [45]. In fact, we employ notions similar to those in [45, Section 3] throughout and essentially all results there are true in our case as well. We focus solely on deriving analogues relevant for this article.

Given $\tau \in \text{SPCT}^{\sigma}(\alpha)$ for some $\alpha \vDash n$, and a positive integer i such that $1 \leq i \leq n - 1$, we say that i and $i + 1$ are *attacking* if either

- (1) i and $i+1$ are in the same column in τ (in which case we call them *strongly attacking*),
or
(2) i and $i+1$ are in adjacent columns in τ , with $i+1$ positioned southeast of i (in which case we call them *weakly attacking*).

Given an integer i satisfying $1 \leq i \leq n-1$, let $s_i(\tau)$ denote the filling obtained by interchanging the positions of entries i and $i+1$ in τ . Define operators π_i for $1 \leq i \leq n-1$ as follows.

$$(3.1) \quad \pi_i(\tau) = \begin{cases} \tau & i \notin \text{Des}(\tau) \\ 0 & i \in \text{Des}(\tau), i \text{ and } i+1 \text{ attacking} \\ s_i(\tau) & i \in \text{Des}(\tau), i \text{ and } i+1 \text{ nonattacking} \end{cases}$$

If $i \in \text{Des}(\tau)$ is such that i and $i+1$ are attacking (respectively nonattacking) then i is an *attacking descent* (respectively *nonattacking descent*).

Theorem 3.1. *The operators $\{\pi_i\}_{i=1}^{n-1}$ satisfy the same relations as $H_n(0)$. Thus, given a composition $\alpha \vDash n$, they define a 0-Hecke action on $\text{SPCT}^\sigma(\alpha)$.*

The proof of the above theorem in the context of SCTs involves a lengthy verification and relies on some preliminary lemmas [45]. The techniques we need vary slightly and we discuss the differences next. The key lemma for proving that the operators give rise to a 0-Hecke action on SCTs is [45, Lemma 3.7] and it, in turn, relies on [45, Lemmas 3.4, 3.5, 3.6]. Since we avoid an approach involving box-adding operators here, we cannot use an analogue of [45, Lemmas 3.4 and 3.5] for SPCTs. In fact, [45, Lemma 3.6] does not hold in our setting. In spite of this, we do have the following analogue of [45, Lemma 3.7].

Lemma 3.2. (1) *If $\alpha \vDash n$, $\tau \in \text{SPCT}^\sigma(\alpha)$ and $j \in \text{Des}(\tau)$ such that j and $j+1$ are nonattacking, then $s_j(\tau) \in \text{SPCT}^\sigma(\alpha)$.*
(2) *If $\alpha \vDash n$, $\tau \in \text{SPCT}^\sigma(\alpha)$ and $j \notin \text{Des}(\tau)$ such that j is not in the cell to the immediate right of $j+1$, then $s_j(\tau) \in \text{SPCT}^\sigma(\alpha)$.*

Proof. It is clear in both cases that $s_j(\tau)$ has rows that strictly decrease when read from left to right and that the standardized first column word is σ . Thus, to establish that $s_j(\tau) \in \text{SPCT}^\sigma(\alpha)$, we need to check that the triple condition holds. Our proof proceeds by checking the triple condition locally. That is, we focus on a triple configuration involving a fixed set of cells in τ and show that if the triple condition held for this set of cells, then it continues to hold for the same set of cells in $s_j(\tau)$. Consider throughout a fixed triple configuration in τ as in Figure 1.

Assume now that j is a nonattacking descent in τ . If neither j nor $j+1$ belongs to $\{a, b, c\}$, then the triple condition obviously holds in $s_j(\tau)$. Hence suppose this is not the case. Our premise implies that exactly one of j or $j+1$ belongs to $\{a, b, c\}$. Thus, replacing the j by $j+1$ in τ does not alter the relative order of entries in the cells in the triple configuration under consideration. Therefore, we conclude that $s_j(\tau) \in \text{SPCT}^\sigma(\alpha)$.

We establish the second part of the lemma using an analysis similar to before. Assume that $j \notin \text{Des}(\tau)$ such that j is not in the cell to the immediate right of $j+1$. As before,

we infer that both j and $j + 1$ cannot belong to $\{a, b, c\}$. In the case where neither belongs, the triple condition clearly holds in $s_j(\tau)$. Hence assume that one of j or $j + 1$ belongs to $\{a, b, c\}$. Replacing the j by $j + 1$ in τ does not alter the relative order of entries in the cells in the triple configuration under consideration. Thus the triple condition holds in $s_j(\tau)$. \square

We now state lemmas which show that the relations satisfied by the 0-Hecke algebra $H_n(0)$ are also satisfied by the operators $\{\pi_i\}_{i=1}^{n-1}$. The proofs are omitted given their similarity to [45, Lemmas 3.9, 3.10 and 3.11] respectively. We emphasize here that only the first half of Lemma 3.2 is needed for establishing the three lemmas that follow, and thereby, Theorem 3.1. The second half of Lemma 3.2 is crucial towards establishing Lemma 3.7.

Lemma 3.3. *For $1 \leq i \leq n - 1$, we have $\pi_i^2 = \pi_i$.*

Lemma 3.4. *For $1 \leq i, j \leq n - 1$ such that $|i - j| \geq 2$, we have $\pi_i \pi_j = \pi_j \pi_i$.*

Lemma 3.5. *For $1 \leq i \leq n - 2$, we have $\pi_i \pi_{i+1} \pi_i = \pi_{i+1} \pi_i \pi_{i+1}$.*

The proof of Theorem 3.1 follows from Lemmas 3.3, 3.4 and 3.5. Let \mathbf{S}_α denote the \mathbb{C} -linear span of all SPCTs of shape α . We have established that \mathbf{S}_α is an $H_n(0)$ -module. We now construct a direct sum decomposition of \mathbf{S}_α by defining an equivalence relation based on standardized column words of SPCTs.

3.1. Source and sink tableaux. Given a composition $\alpha \vDash n$, define an equivalence relation \sim_α on $\text{SPCT}(\alpha)$ by defining $\tau_1 \sim_\alpha \tau_2$ if and only if $\text{st}(\tau_1) = \text{st}(\tau_2)$, that is, the entries in each column of τ_1 are in the same relative order as the entries in the corresponding column of τ_2 . Suppose that the equivalence classes under \sim_α are E_1, \dots, E_k . Let \mathbf{S}_{α, E_i} denote the \mathbb{C} -linear span of all SPCTs in E_i for $i = 1, \dots, k$. Then we get the following isomorphism of vector spaces

$$(3.2) \quad \mathbf{S}_\alpha \cong \bigoplus_{i=1}^k \mathbf{S}_{\alpha, E_i},$$

which is in fact an $H_n(0)$ -module isomorphism as the following lemma implies. We omit the proof as it is the same as [45, Lemma 6.6].

Lemma 3.6. *Let E_j for $j = 1, \dots, k$ be the equivalence classes under \sim_α for $\alpha \vDash n$. Then for all i such that $1 \leq i \leq n - 1$, we have that $\pi_i(\mathbf{S}_{\alpha, E_j}) \subseteq \mathbf{S}_{\alpha, E_j}$ for any $1 \leq j \leq k$.*

Next, we discuss two important classes of SPCTs that form special representatives of each equivalence class. Let $\alpha \vDash n$. An SPCT τ of shape α is said to be a *source* tableau if for every $i \notin \text{Des}(\tau)$ where $i \neq n$, we have that $i + 1$ lies to the immediate left of i . An SPCT τ of shape α is said to be a *sink* tableau if for every $i \in \text{Des}(\tau)$, we have that i and $i + 1$ are attacking. Figure 5 shows a source tableau (left) and a sink tableau (right) of shape $(1, 3, 2, 4)$ and type 1324. Note that the standardized column word of each tableau is 1324 213 12 1.

1			
6	5	4	
3	2		
10	9	8	7

3			
8	5	1	
7	4		
10	9	6	2

FIGURE 5. An example of a source tableau and a sink tableau.


By performing an analysis very similar to that in [45, Section 6] we obtain the following analogue of [45, Corollary 6.15], which is pertinent for our purposes. See Remark 3.8 for a brief discussion regarding its proof.

Lemma 3.7. *There is a unique source tableau and a unique sink tableau in every equivalence class under \sim_α .*

Remark 3.8. The proofs of the various lemmas leading up to the proof of [45, Corollary 6.15] employ the notion of removable nodes of a composition diagram followed by distinguished removable nodes in SCTs. The definition of removable nodes in the aforementioned case is inspired by the cover relation for the (left) composition poset [8, Section 2.2]. Saturated chains in this poset are in bijection with SCTs. Although we have not described SPCTs as saturated chains in a poset on compositions, we may still define a removable node as follows. Given the composition diagram α , we call a cell in position (i, j) a *removable node* if one of the two conditions below is met.

- (1) $\alpha_i = j \geq 2$ and there is no part of length $j - 1$ strictly north of α_i .
- (2) $\alpha_i = j = 1$.

Figure 6 shows $\alpha = (2, 3, 1, 4)$ with removable nodes corresponding to cells filled with bullets. With this definition of removable node, one can define distinguished removable nodes in

FIGURE 6. The removable nodes of $\alpha = (2, 3, 1, 4)$.

precisely the same way as in [45, Section 6]. In particular, the analogues of [45, Propositions 6.13 and 6.14], which are crucial for deriving our Lemma 3.7, hold in our setting. Thus concludes our remark.

Recall that our motivation as outlined in the introduction is to enumerate compatible pairs of permutations. Consider a CT τ of shape α whose i -th and $i + 1$ -th column, for some $i \geq 1$ contain the same number of entries, say n . Then $(st_i(\tau), st_{i+1}(\tau))$ is a compatible pair of permutations in \mathfrak{S}_n . Note that if we consider the entries in columns i and $i + 1$ only, we can construct a filling of shape (2^n) . As the triple condition is clearly satisfied, this filling

is in fact a PCT. We infer that compatible pairs of permutations in \mathfrak{S}_n are a subset of the set of pairs of permutations $(st_1(\tau), st_2(\tau))$ corresponding to tableaux $\tau \in \text{PCT}((2^n))$. In Section 5, we establish that this containment is in fact an equality. We proceed towards studying $\text{SPCT}((2^n))$ in depth.

4. DESCENTS IN 2-COLUMNED TABLEAUX AND ASCENT-DESCENTS ON LABELED BINARY TREES

We refer to SPCTs of shape (2^n) for some $n \geq 1$ as *2-columned tableaux* of size $2n$. Consider the following refined classification of certain descents in a 2-columned tableau τ . Suppose i is an entry in the first column of τ . Then $i \in \text{Des}(\tau)$. Based on the relative position of $i + 1$ in τ , we may define the following sets.

$$\begin{aligned} \text{Des}_N(\tau) &= \{i \mid i + 1 \text{ is in the first column north of } i\} \\ \text{Des}_S(\tau) &= \{i \mid i + 1 \text{ is in the first column south of } i\} \\ \text{Des}_{NE}(\tau) &= \{i \mid i + 1 \text{ is in the second column northeast of } i\} \\ \text{Des}_{SE}(\tau) &= \{i \mid i + 1 \text{ is in the second column southeast of } i\} \end{aligned}$$

Additionally, we say that i an *N*-descent (respectively *S*-descent) of τ if $i + 1$ is in the first column and north (respectively south) of i . Similarly, we say that i an *NE*-descent (respectively *SE*-descent) of τ if $i + 1$ is in the second column and northeast (respectively southeast) of i . Figure 7 shows a $\tau \in \text{SPCT}((2^n))$ where $n = 10$. The arcs on the left connect i and $i + 1$ where both belong to the first column. In particular, the number of red arcs on the left is equal to $|\text{Des}_N(\tau)|$ whereas the number of purple arcs on the left is equal to $|\text{Des}_S(\tau)|$. The arcs on the right connect rows where the i is in the first column and the $i + 1$ is in the second column. In particular, the number of red arcs on the right is equal to $|\text{Des}_{NE}(\tau)|$ whereas the number of purple arcs on the left is equal to $|\text{Des}_{SE}(\tau)|$.




FIGURE 7. The classification of descents in the first column of τ .

To enumerate ordered pairs $(st_1(\tau), st_2(\tau))$ where τ ranges over 2-columned tableaux of size $2n$, we need to enumerate equivalence classes under $\sim_{(2^n)}$. By Lemma 3.7, we may

equivalently enumerate sink tableaux of shape (2^n) . Such a sink τ can be alternatively characterized as one that satisfies $|\text{Des}_{NE}(\tau)| = 0$. This motivates our study of the distribution of the quadruple of statistics

$$(|\text{Des}_N(\tau)|, |\text{Des}_S(\tau)|, |\text{Des}_{NE}(\tau)|, |\text{Des}_{SE}(\tau)|)$$

as τ varies over all tableaux in $\text{SPCT}((2^n))$. We establish that this distribution coincides with the distribution of left ascents, left descents, right ascents and right descents over the set of labeled binary trees on n nodes. To this end, we need certain combinatorial notions attached to well-known ‘Catalan objects’ known as Dyck paths.

4.1. Dyck paths, unlabeled and labeled. Given a nonnegative integer n , a *Dyck path of semi-length n* is a lattice path in the plane beginning at $(0, 0)$ and ending at $(2n, 0)$ consisting of *up-steps*, which correspond to a translation by $(1, 1)$, and *down-steps*, which correspond to a translation by $(1, -1)$. We denote the set of Dyck paths of semi-length n by \mathcal{D}_n . The cardinality of \mathcal{D}_n is well-known to be Cat_n , the n -th Catalan number equaling $\frac{1}{n+1} \binom{2n}{n}$. We can identify a Dyck path $D \in \mathcal{D}_n$ with a ud-word of length $2n$ by writing u (respectively d) for every up-step (respectively down-step) in D from left to right. We denote this word by $\text{dw}(D)$ and refer to it as the *Dyck word* of D . A Dyck path is *prime* if it touches the x -axis at its initial and terminal points only. Every Dyck path D can be factorized uniquely as a concatenation of prime Dyck paths $D_1 \cdot D_2 \cdots D_k$. We call each D_i a *factor* of D . Figure 8 depicts an element $D \in \mathcal{D}_9$ with 4 factors. We have $\text{dw}(D) = ud\ ud\ uuududdudd\ uudd$, where the spaces separate the Dyck words of the factors of D .




FIGURE 8. A Dyck path D of semi-length 9.

Consider $D \in \mathcal{D}_n$. If D is endowed with a labeling on its down-steps with distinct positive integers drawn from $[n]$, then we obtain a *labeled Dyck path* D' of semi-length n . We denote the set of labeled Dyck paths of semi-length n by \mathcal{LD}_n . The cardinality of \mathcal{LD}_n equals $n! \text{Cat}_n$. We define prime labeled Dyck paths and factors of labeled Dyck paths in the obvious manner following the description in the unlabeled case. Given $D \in \mathcal{LD}_n$, we abuse notation and denote by $\text{dw}(D)$ the Dyck word of the underlying unlabeled Dyck path. Constructing the *labeled Dyck word* $\text{ldw}(D)$ is slightly more involved and is described next. First we recursively assign labels to the up-steps from right to left. Consider the i -th up-step in D for i from n down to 1. Let \mathfrak{D}_i (respectively \mathfrak{U}_i) be the set of labels that belong to down-steps (respectively up-steps) after the i -th up-step. Clearly $\mathfrak{U}_n = \emptyset$. Assign the minimum of $\mathfrak{D}_i \setminus \mathfrak{U}_i$ as the label of this i -th up-step. Note that as the algorithm executes, we always have that \mathfrak{U}_i is a strict subset of \mathfrak{D}_i . Here we are crucially using the fact that D is a (labeled) Dyck path. Thus, the minimum of $\mathfrak{D}_i \setminus \mathfrak{U}_i$ is well-defined and our procedure terminates when

i attains the value 1. Next, read the resulting Dyck path, with labels on both up-steps and down-steps, from left to right and write u_i (respectively d_i) for an up-step (respectively down-step) labeled i . The word thus obtained is the labeled Dyck word $\text{ldw}(D)$.

It is easy to see that for a labeled Dyck path D with prime factorization $D_1 \cdots D_k$, we have

$$\text{ldw}(D) = \text{ldw}(D_1) \cdots \text{ldw}(D_k).$$

Figure 9 shows a labeled Dyck path D of semi-length 10. The arcs join the up-steps and down-steps that share the same label during the procedure for computing $\text{ldw}(D)$. We have

$$\text{ldw}(D) = u_7 u_3 d_7 d_3 \quad u_{10} u_9 u_8 u_4 d_4 u_1 d_1 d_{10} u_6 d_6 d_8 u_5 u_2 d_9 d_2 d_5.$$




FIGURE 9. A labeled Dyck path of semi-length 10.

4.2. Bijection between \mathcal{LD}_n and $\text{SPCT}((2^n))$. Recall that $|\mathcal{LD}_n| = n! \text{Cat}(n)$. By Theorem 2.4, we have that $|\text{SPCT}^\sigma((2^n))| = |\text{SRT}((2^n))|$ for all $\sigma \in \mathfrak{S}_n$. The following bijection between $\text{SRT}((2^n))$ and \mathcal{D}_n is folklore. For $T \in \text{SRT}((2^n))$, the i -th step for $1 \leq i \leq 2n$ in the associated Dyck path is an up-step if i is in the second column of T and it is a down-step otherwise. Thus, we infer that $|\text{SPCT}((2^n))| = |\coprod_{\sigma \in \mathfrak{S}_n} \text{SPCT}^\sigma((2^n))| = \sum_{\sigma \in \mathfrak{S}_n} |\text{SRT}((2^n))| = n! \text{Cat}_n$.

By a simple extension of the aforementioned bijection, we obtain one between \mathcal{LD}_n and $\text{SPCT}((2^n))$ that associates a $D \in \mathcal{LD}_n$ with $\tau \in \text{SPCT}((2^n))$ as follows. For $1 \leq i \leq 2n$, the i -th step is an up-step if i belongs to the second column. Otherwise, the i -th step is a down-step labeled by the index of the row occupied by i in τ . We denote this map from $\text{SPCT}((2^n))$ to \mathcal{LD}_n by SPCTtoLDyck and defer the proof that this is indeed a bijection to after we describe the inverse map, denoted by LDycktoSPCT , next.



FIGURE 10. Transforming an SPCT to a labeled Dyck path.

Let $D \in \mathcal{LD}_n$ with $\text{ldw}(D) = w_1 \cdots w_{2n}$. Construct a filling τ of shape (2^n) as follows. For $1 \leq i \leq n$, let $1 \leq p < q \leq 2n$ be the positive integers such that $w_p = u_i$ and $w_q = d_i$. Then the i -th row of τ contains q and p in the first and second columns respectively. We define

τ to be $\text{LDycktoSPCT}(D)$. The next lemma establishes that LDycktoSPCT is the inverse of SPCTtoLDyck .

Lemma 4.1. *LDycktoSPCT is a bijection from \mathcal{LD}_n to $\text{SPCT}((2^n))$ with inverse given by SPCTtoLDyck .*

Proof. Consider $D \in \mathcal{LD}_n$ and let $\tau = \text{LDycktoSPCT}(D)$. To establish that τ is an SPCT, we must show that the triple condition holds. Suppose that this is not the case. Then we have a configuration in τ of the type shown in Figure 11 such that $a > c > b$. We interpret




FIGURE 11. A triple configuration in $\text{LDycktoSPCT}(D)$.

this configuration in terms of D . Henceforth, in this proof, we assume that the up-steps in D are labeled using the procedure for computing $\text{ldw}(D)$.

Assume that a and b belong to row i and that c belongs to row j , where $1 \leq i < j \leq n$. Then the a -th step (respectively b -th step) of D is a down-step (respectively up-step) labeled i . Let S_b and S_c be the sets consisting of labels on the down-steps after the b -th and c -th steps respectively, ordered in increasing order. We have $S_c \subseteq S_b$ as $c > b$ and $i \in S_c$ as $a > c$. As the c -th step, which is an up-step, gets the label j , and $j > i$, we conclude that there exists an up-step after the c -th step that has the label i . This follows from the description of our procedure for constructing the labeled Dyck word. Thus, we arrive at a contradiction as the b -th step cannot be labeled i . We conclude that the triple condition holds in τ and that $\tau \in \text{SPCT}((2^n))$.

To show that LDycktoSPCT is a bijection, we establish its injectivity. If $\text{LDycktoSPCT}(D_1) = \text{LDycktoSPCT}(D_2)$ for labeled Dyck paths D_1, D_2 , then we must have $\text{ldw}(D_1) = \text{ldw}(D_2)$. This implies that the unlabeled Dyck paths underlying D_1 and D_2 are the same. Additionally, the labels on the corresponding down-steps in D_1 and D_2 are equal. Thus, $D_1 = D_2$ and we conclude that LDycktoSPCT is a bijection.

It is clear that $\text{SPCTtoLDyck}(\tau) = D$ as the entries in the first column of τ and the indices of the rows they belong to completely determine the down-steps and their labels. We conclude that SPCTtoLDyck and LDycktoSPCT are bijections that are mutually inverse. \square

We proceed to describe unlabeled and labeled binary trees and subsequently introduce descent statistics that are equidistributed with the quadruple $(\text{Des}_N, \text{Des}_S, \text{Des}_{NE}, \text{Des}_{SE})$.

4.3. Plane binary trees, unlabeled and labeled. A *rooted tree* is a finite acyclic graph with a distinguished node called the *root*. A *plane binary tree* is a rooted tree in which every node has at most two children, of which at most one is called a *left child* and at most one is called a *right child*. Henceforth, we reserve the term binary tree for a plane binary tree. We

denote the set of binary trees on $n \geq 1$ nodes by \mathcal{T}_n . The set of nodes of $T \in \mathcal{T}_n$ is denoted by $V(T)$, and the root node is referred to as $\text{root}(T)$. We abuse notation on occasion and write $v \in T$ when we mean $v \in V(T)$. A node $v \in T$ is an *internal node* if it has at least one child. Otherwise it is a *leaf*.

A *left path* is a binary tree T such that no node has a right child. We can decompose any binary tree T into a set of left paths $\{T_1, \dots, T_k\}$ by omitting all edges that connect an internal node to its right child, provided it exists. We refer to this set of left paths as the *maximal left path decomposition* of T and denote it by $\text{mlpd}(T)$. Figure 12 depicts a binary tree T on 10 nodes with $V(T) = \{v_1, \dots, v_{10}\}$. The nodes that belong to the same shaded region constitute a left path in $\text{mlpd}(T)$.

Given $T \in \mathcal{T}_n$, let $\text{mlpd}(T)$ be $\{T_1, \dots, T_m\}$ where T_1 contains $\text{root}(T)$. For every T_i where $2 \leq i \leq m$, there exists a unique node $v \in T$ such that $\text{root}(T_i)$ is the right child of v in T . We define v to be the *parent* of T_i in T and denote it by $\text{parent}(T_i)$. In Figure 12, consider the left path on nodes v_7 and v_8 . Then its parent is v_6 . Note the crucial fact that a binary tree T is completely determined by specifying its maximal left path decomposition, the left path that contains the root, and the parents of the other left paths.




FIGURE 12. A binary tree with its maximal left path decomposition.

A *labeled plane binary tree* (or simply a *labeled binary tree*) on $n \geq 1$ nodes is a binary tree whose nodes have distinct labels drawn from $[n]$. We denote the set of labeled binary trees on n nodes for $n \geq 1$ by \mathcal{T}_n^ℓ . Given a node $u \in T$, we refer to its label as u^ℓ . For a labeled binary tree, we have the following refined classification for its edges. Suppose that q is the right child of p . If $p^\ell < q^\ell$, then the edge joining p and q is a *right ascent*. Otherwise it is a *right descent*. Now suppose that q is the left child of p . If $p^\ell < q^\ell$, then the edge joining p and q is a *left ascent*. Otherwise it is a *left descent*. Let $\text{rasc}(T)$ (respectively, $\text{rdes}(T)$, $\text{lasc}(T)$ and $\text{ldes}(T)$) be the number of right ascents (respectively, right descents, left ascents and left descents) in T . For the labeled binary tree T in Figure 13, we have that $\text{rasc}(T) = 3$, $\text{rdes}(T) = 1$, $\text{lasc}(T) = 1$, $\text{ldes}(T) = 3$. The thickened edges in the figure correspond to descents, both left and right.

We will abuse notation and use the terms left path and maximal left path decomposition in the context of their natural labeled analogues respectively.




FIGURE 13. A labeled binary tree with $\text{rasc}(T)=3$, $\text{rdes}(T)=1$, $\text{lasc}(T)=1$, $\text{lides}(T)=3$.

4.4. Labeled Dyck paths and labeled binary trees. Note that since $|\mathcal{T}_n| = \text{Cat}_n$ the cardinalities of \mathcal{LD}_n and \mathcal{T}_n^ℓ are equal. We seek a bijection between these two sets that tracks the quadruple of statistics $(\text{lasc}, \text{lides}, \text{rasc}, \text{rdes})$ on labeled binary trees in a manner that will eventually allow us to relate them to the quadruple $(|\text{Des}_N|, |\text{Des}_S|, |\text{Des}_{NE}|, |\text{Des}_{SE}|)$ via the map LDycktoSPCT . First, we need some more terminology pertaining to Dyck paths.

A *run* of $D \in \mathcal{D}_n$ is a maximal sequence of down-steps. A *run* of $D \in \mathcal{LD}_n$ is defined similarly, except that we retain the labels. By considering all runs in $D \in \mathcal{LD}_n$ from right to left, we obtain the *run sequence* of D . Consider a run in D composed of a steps. We can transform this run into a left path on a nodes where the label of the nodes read from root to leaf are the labels in the run read from right to left. Thus, from the run sequence of D , say (R_1, \dots, R_m) , we obtain a sequence of left paths (T_1, \dots, T_m) . This datum can be used to construct a labeled binary tree as described in Algorithm 1. As usual, by the label on an up-step of a labeled Dyck path, we mean the label it acquires when computing its labeled Dyck word.

Algorithm 1 Transforming a labeled Dyck path to a labeled binary tree

Input: $D \in \mathcal{LD}_n$.

Output: $T \in \mathcal{T}_n^\ell$.

```

1: function LDYCKTOLTREE( $D$ )
2:   Let  $(R_1, \dots, R_m)$  be the run sequence of  $D$ ;
3:   Let  $(T_1, \dots, T_m)$  be the sequence of left paths corresponding to the run sequence;
4:   Let  $T := T_1, i := 2$ ;
5:   while  $i < m$  do
6:      $j :=$  label on the up-step immediately after the rightmost down-step in  $R_i$ ;
7:     Make the left path  $T_i$  the right subtree of the node labeled  $j$  in  $T$ ;
8:     Let  $T$  be this updated tree;
9:      $i := i + 1$ ;
10:  return  $T$ ;
  
```

Algorithm 2 Transforming a labeled binary tree to a labeled Dyck path

Input: $T \in \mathcal{T}_n^\ell$.**Output:** A labeled Dyck word corresponding to a labeled Dyck path in \mathcal{LD}_n .

```

1: function LTREETOLDYCK( $T$ )
2:   Let  $Q := \emptyset$ , DyckWord :=  $\varepsilon$ , NumSteps := 0, flag := 1;
3:   Let  $C$  be the unique left path in  $\text{mlpd}(T)$  containing  $\text{root}(T)$ ;
4:   while NumSteps <  $2n$  do
5:     if flag is equal to 1 then
6:       Let  $v_1, \dots, v_m$  be the nodes in  $C$  from root to leaf;
7:       for  $1 \leq j \leq m$  do
8:          $Q := Q \cup \{v_j\}$ ; ▷ Push  $i$  into  $Q$ 
9:          $i := \text{label of node } v_j$ ;
10:        DyckWord :=  $d_i \cdot \text{DyckWord}$ ;
11:        NumSteps := NumSteps + 1;
12:      flag := 0;
13:    else
14:      Let  $v$  be the node in  $Q$  whose label  $m$  is minimal;
15:       $Q := Q \setminus \{v\}$ ; ▷ Pop minimum element from  $Q$ 
16:      DyckWord :=  $u_m \cdot \text{DyckWord}$ ;
17:      NumSteps := NumSteps + 1;
18:      Let  $v$  be node in  $T$  with label  $m$ ;
19:      if  $v$  has a right child  $v'$  in  $T$  then
20:        Let  $C$  be the unique left path in  $\text{mlpd}(T)$  rooted at  $v'$ ;
21:        flag := 1;
22:    return DyckWord;

```

Furthermore, once it is pushed, it must take part in exactly one pop operation. In fact, the following stronger claim holds: Every node gets pushed and popped exactly once during the execution of our algorithm. This observation, which hinges on the fact that T is connected, is present implicitly in the proof of the next lemma.

Lemma 4.3. *In Algorithm 2, there can never be a pop operation when Q is empty.*

Proof. Suppose Q becomes empty at some point during our algorithm upon popping the node v . Let $T_1 \in \text{mlpd}(T)$ be the left path that contains v . As v is the last node to be popped from Q , we conclude that every other node in T_1 except v has been popped from Q before v . If v has a right child v_0 in T , then the next step of our algorithm, after setting flag equal to 1, pushes all nodes that belong to the left path in $\text{mlpd}(T)$ containing v_0 into Q . This ensures that the next pop is performed on a nonempty Q .

Assume henceforth that v does not have a right child in T . We claim that NumSteps equals $2n$ upon performing the pop that removes v from Q , thereby implying that the algorithm

NumSteps	Q	Operation
1	{5}	Push 5 into Q
2	{2, 5}	Push 2 into Q
3	{2, 5, 9}	Push 9 into Q
4	{5, 9}	Pop 2 from Q
5	{9}	Pop 5 from Q
6	{8, 9}	Push 8 into Q
7	{6, 8, 9}	Push 6 into Q
8	{8, 9}	Pop 6 from Q
9	{8, 9, 10}	Push 10 into Q
10	{1, 8, 9, 10}	Push 1 into Q
11	{8, 9, 10}	Pop 1 from Q
12	{4, 8, 9, 10}	Push 4 into Q
13	{8, 9, 10}	Pop 4 from Q
14	{9, 10}	Pop 8 from Q
15	{10}	Pop 9 from Q
16	\emptyset	Pop 10 from Q
17	{3}	Push 3 into Q
18	{3, 7}	Push 7 into Q
19	{7}	Pop 3 from Q
20	\emptyset	Pop 7 from Q

TABLE 1. Algorithm 2 applied to the labeled binary tree in Figure 14.

terminates subsequently. To establish this claim, we proceed by contradiction. Suppose NumSteps is strictly less than $2n$ upon popping v . Then there exists a node $v' \in T$ that has not been pushed into Q yet. For if there was no such node, then every node in T would have been pushed and popped from Q , since Q becomes empty upon popping v , contradicting our premise that NumSteps $< 2n$.

Let $T'_1 \in \text{mlpd}(T)$ be the left path containing v' . As v' has not been pushed into Q yet, we infer that no node in T'_1 has been pushed into Q . In particular, root(T) does not belong to T'_1 and $T'_1 \neq T_1$. Let $T'_2 \in \text{mlpd}(T)$ be the left path that contains parent(T'_1). As no node of T'_1 has been pushed into Q , we infer that no element in T'_2 has been pushed into Q either. As before, we have that root(T) is not in T'_2 and $T'_2 \neq T_1$. Continuing this inductively we obtain an infinite sequence of distinct left paths T'_1, T'_2, \dots such that $T'_{i+1} \in \text{mlpd}(T)$ contains parent(T'_i) for $i \geq 1$. Each of these left paths contains nodes that have not been pushed into Q . As $V(T)$ is finite, we have a contradiction. \square

We note an important consequence of Lemma 4.3. At every stage during the execution of Algorithm 2, the number of pushes weakly exceeds the number of pops. Additionally, every time a pop empties Q , the number of pushes equals the number of pops up until that stage.

Conversely, if the number of pops equals the number of pushes at any stage, then Q must be empty. These observations imply that in every suffix of DyckWord, there are at least many d 's as there are u 's. In summary, we have the following.

Corollary 4.4. $\chi(\text{DyckWord})$ is the Dyck word of a Dyck path.

Our next aim is to establish the stronger assertion that DyckWord is the labeled Dyck word of some labeled Dyck path. Let D^u denote the Dyck path corresponding to $\chi(\text{DyckWord})$. Let D^ℓ denote the labeled Dyck path whose underlying unlabeled Dyck path is D^u and whose labels on the down-steps from left to right are derived from the indices of the d 's in DyckWord from left to right.

Lemma 4.5. DyckWord equals $\text{ldw}(D^\ell)$.

Proof. Note that $\chi(\text{DyckWord})$ equals $\text{dw}(D^\ell)$. Furthermore, by construction, the indices of the d 's read from left to right in $\text{ldw}(D^\ell)$ coincide with the indices of the d 's read from left to right in DyckWord. Therefore, all we need to show is that the index of a u in DyckWord matches that of the corresponding u in $\text{ldw}(D^\ell)$.

Let W_{old} be the value of DyckWord at an intermediate stage during the execution of the algorithm just prior to a pop operation. Let W_{new} be the value of DyckWord immediately after a single pop. Then $W_{new} = u_i \cdot W_{old}$ where i is the minimum label amongst all nodes in Q before the pop was performed. We discuss an alternate way to describe this value i . Let A (respectively B) be the set of indices of d s (respectively u s) in W_{old} . Then A consists of labels of the nodes that have been pushed into Q (up until the intermediate stage we are considering), whereas B consists of labels of nodes that have been both pushed and popped. Then $A \setminus B$ contains all labels corresponding to nodes in Q before the pop. Thus, we infer that i is the minimum element in $A \setminus B$.

Now consider the sub-path of D^ℓ that corresponds to W_{old} . The labels on the down-steps are precisely the elements of A . The sub-path of D^ℓ corresponding to W_{new} begins with an up-step. The preceding procedure where we computed the labels on up-steps coincides precisely with our description of the procedure that endows up-steps with labels during the computation of the labeled Dyck word. This finishes the proof. \square

In other words, Algorithm 2 outputs the labeled Dyck word of D^ℓ given a labeled binary tree T . We note certain features of this algorithm that are immediate from its description.

Lemma 4.6. Algorithm 2 possesses the following three features.

- (1) If the label of $\text{root}(T)$ is i , then the last letter in DyckWord is d_i .
- (2) If i is the label of the left child of the node labeled j in T , then DyckWord has a contiguous subword of the type $d_i d_j$.
- (3) If i is the label of the right child of the node labeled j in T , then DyckWord has a contiguous subword of the type $d_i u_j$.

The reader is invited to compare the statement of Lemma 4.6 to that of Lemma 4.2. This comparison naturally leads us to our next theorem. From this point on, we identify $\text{LTreetoLDyck}(T)$ with the labeled Dyck path D^ℓ instead of $\text{ldw}(D^\ell)$. Our next result establishes that LTreetoLDyck is a bijection and its inverse is the map LDycktoLTree .

Theorem 4.7. $\text{LTreetoLDyck} : \mathcal{T}_n^\ell \rightarrow \mathcal{LD}_n$ is a bijection and its inverse is given by LDycktoLTree .

Proof. As \mathcal{LD}_n and \mathcal{T}_n^ℓ have the same cardinality, to establish the claim it suffices to show that $\text{LDycktoLTree} \circ \text{LTreetoLDyck}$ is the identity map. Let D^ℓ be the labeled Dyck path corresponding to the labeled Dyck word output by Algorithm 2 for some $T \in \mathcal{T}_n^\ell$. From Lemmas 4.2 and 4.6, it follows that $\text{LDycktoLTree}(D^\ell) = T$. This finishes the proof. \square

Figure 15 should aid the reader in understanding our next theorem.

Theorem 4.8. $\text{LDycktoLTree} \circ \text{SPCTtoLDyck} : \text{SPCT}((2^n)) \rightarrow \mathcal{T}_n^\ell$ is a bijection that maps the ordered quadruple of statistics $(|\text{Des}_N|, |\text{Des}_S|, |\text{Des}_{NE}|, |\text{Des}_{SE}|)$ to $(\text{lasc}, \text{ldes}, \text{rasc}, \text{rdes})$.

Proof. It is clear that $\text{LDycktoLTree} \circ \text{SPCTtoLDyck}$ is a bijection. We focus on proving the second half of the theorem. Consider $\tau \in \text{SPCT}((2^n))$. Let D be $\text{SPCTtoLDyck}(\tau)$ and let T be $\text{LDycktoLTree}(D)$. Suppose first that i and $i + 1$ both belong to the first column in the p -th and q -th row in τ respectively. Then the i -th step of D is a down-step labeled p and the $i + 1$ -th step is down-step labeled q . In terms of $\text{ldw}(D)$, this corresponds to an instance of a contiguous subword of the type $\mathbf{d}_p \mathbf{d}_q$. Lemma 4.2 implies that p is the label of the left child of the node labeled q in T . If $i \in \text{Des}_N(\tau)$, then $p > q$ and the edge in T joining the nodes labeled p and q is a left ascent. Thus we infer that $|\text{Des}_N(\tau)| = \text{lasc}(T)$. A similar argument proves that $|\text{Des}_S(\tau)| = \text{ldes}(T)$.

Now suppose that i belongs to the first column in the p -th row while $i + 1$ belongs to the second column in the q -th row in τ . Then the i -th step of D is a down-step labeled p and the $i + 1$ -th step is an up-step labeled q . In terms of $\text{ldw}(D)$, this corresponds to an instance of a contiguous subword $\mathbf{d}_p \mathbf{u}_q$. Lemma 4.2 implies that the node labeled p is the right child of the node labeled q in T . If $i \in \text{Des}_{NE}(\tau)$, then $p > q$ and the edge joining the nodes labeled p and q is a right ascent. Thus, we infer that $|\text{Des}_{NE}(\tau)| = \text{rasc}(T)$. A similar argument proves that $|\text{Des}_{SE}(\tau)| = \text{rdes}(T)$. \square

We return to our question at the beginning of this section, that of enumerating sink tableaux of shape (2^n) . Recall that a 2-columned tableau τ is a sink if $|\text{Des}_{NE}(\tau)| = 0$. By Theorem 4.8, the set of sink tableaux of shape (2^n) is in bijection with the set of labeled binary trees on n nodes that have no right ascents. By a result originally due to Pak-Postnikov [39] (see also [38]), the number of such trees is $(n + 1)^{n-1}$. Thus, we arrive at the following result.

Theorem 4.9. The set of ordered pairs $(\text{st}_1(\tau), \text{st}_2(\tau))$ as τ ranges over all 2-columned SPCTs of size $2n$ is in bijection with the set of labeled binary trees on n nodes with no right ascents. Thus, its cardinality is $(n + 1)^{n-1}$.




FIGURE 15. The statistic-preserving bijection between $\text{SPCT}((2^n))$ and \mathcal{T}_n^ℓ by way of labeled Dyck paths.

In the next section, we show that the set of ordered pairs in Theorem 4.9 does in fact coincide with the set of compatible pairs. To this end, we establish that for every $(\text{st}_1(\tau), \text{st}_2(\tau))$ obtained from some $\tau \in \text{SPCT}((2^n))$, we can construct an SCT τ' that has two adjacent columns, say i and $i + 1$, satisfying $\text{st}_i(\tau') = \text{st}_1(\tau)$ and $\text{st}_{i+1}(\tau') = \text{st}_2(\tau)$. An alternative characterization for $(\text{st}_1(\tau), \text{st}_2(\tau))$ using pattern-avoidance is key.

5. ALLOWABLE PAIRS OF PERMUTATIONS AND ACYCLIC GRAPHS

We begin with a lemma characterizing standardized first and second column words of 2-columned tableaux in terms of pattern-avoidance. The straightforward proof relies on the triple condition and is omitted.

Lemma 5.1. *Consider $\tau \in \text{SPCT}((2^n))$. Denote the entry in the i -th row and j -th column by $\tau_{(i,j)}$. Then the following properties hold.*

- (1) *For $1 \leq i < j \leq n$, If $\tau_{(i,1)} > \tau_{(j,1)}$, then $\tau_{(i,2)} > \tau_{(j,2)}$.*
- (2) *For $1 \leq i < j < k \leq n$, if $\tau_{(i,1)} < \tau_{(j,1)} < \tau_{(k,1)}$, then we cannot have that $\tau_{(i,2)} > \tau_{(k,2)} > \tau_{(j,2)}$.*

Consider $\sigma, \gamma \in \mathfrak{S}_n$. We say that (σ, γ) is $(21, 12)$ -avoiding if $\sigma(i) > \sigma(j)$ for integers $1 \leq i < j \leq n$ implies that $\gamma(i) > \gamma(j)$. Similarly, we say that it is $(123, 312)$ -avoiding if $\text{stan}(\sigma(i)\sigma(j)\sigma(k)) = 123$ for integers $1 \leq i < j < k \leq n$ implies that $\text{stan}(\gamma(i)\gamma(j)\gamma(k)) \neq 312$. The statement that (σ, γ) is $(21, 12)$ -avoiding is equivalent to saying that $\sigma \leq_L \gamma$ as it implies that $\text{Inv}(\sigma) \subseteq \text{Inv}(\gamma)$. We call an ordered pair of permutations (σ, γ) an *allowable pair*

provided it is both (21, 12) and (123, 312)-avoiding. Lemma 5.1 states that $(st_1(\tau), st_2(\tau))$ is an allowable pair for any 2-columned tableau τ .

Remark 5.2. The choice of terminology may be explained by noticing that a pair (σ, γ) is allowable in our sense if and only if the pair (σ^c, γ^c) is allowable in the sense of [4]. Here σ^c refers to the *complement* of σ , that is, the permutation obtained by replacing i with $n - i + 1$ if $\sigma \in \mathfrak{S}_n$. Under the complementation map, our pattern-avoidance characterization for allowable pairs is equivalent to the one noted in the remark preceding [3, Lemma 3].


Our next aim is to show that the set of allowable pairs is equal to the set of ordered pairs of permutations considered in Theorem 4.9. We achieve this by associating a directed acyclic graph to an allowable pair. The acyclicity of this graph implies that its node set can be endowed with the structure of a poset as follows. Given nodes v, w , we define $v \geq w$ if there exists a directed path from v to w . For the graph we construct, a linear extension of the associated poset, known as a *topological sort* in theoretical computer science, is a 2-columned tableau. We provide the details next.

Given permutations $\sigma_1, \dots, \sigma_k \in \mathfrak{S}_n$ where $k \geq 2$, we say that the sequence $(\sigma_1, \dots, \sigma_k)$ is an *allowable sequence* of length k if (σ_j, σ_{j+1}) is allowable for all $1 \leq j \leq k - 1$. With an allowable sequence $(\sigma_1, \dots, \sigma_k)$ we associate a directed simple graph $\mathbf{G} := \mathbf{G}(\sigma_1, \dots, \sigma_k)$ as follows. Consider the composition diagram of shape (k^n) . Place a node in the middle of each cell of this diagram. These nodes together form the node (or vertex) set of \mathbf{G} . We reference each node by the pair (i, j) where i is the row and j is the column to which this node belongs. The edge set of \mathbf{G} is determined by $(\sigma_1, \dots, \sigma_k)$ as follows.

- (1) For all $1 \leq i \leq n$ and $1 \leq j \leq k - 1$, draw a directed edge from (i, j) to $(i, j + 1)$. We refer to these edges as *horizontal edges*.
- (2) For all $1 \leq j \leq k$, draw a directed edge from (p, j) to (i, j) where $1 \leq i \neq p \leq n$ are such that $\sigma_j(i) < \sigma_j(p)$. We refer to these edges as *vertical edges*.
- (3) For all $2 \leq j \leq k$, draw a directed edge from (p, j) to $(i, j - 1)$ where $1 \leq i < p \leq n$ are such that $\sigma_j(i) < \sigma_j(p)$. We refer to these edges as *diagonal edges*.

For any $1 \leq j \leq k$ and $1 \leq i \leq n$, the nodes (i, j) are said to belong to the j -th column. Note that all horizontal edges are oriented from west to east, whereas all diagonal edges are oriented from southeast to northwest. Furthermore, they connect nodes that belong to consecutive columns.

Figure 16 depicts the graph \mathbf{G} associated to the allowable sequence (123, 132, 321). Note that it does not contain any directed cycles. Thus, its nodes can be assigned distinct numbers from 1 through 9 such that all arrows point from nodes with larger labels to those with smaller labels. While this labeling is not necessarily unique, the conditions on the edge set guarantee that we obtain an SPCT of shape (3^3) in this manner. We emphasize here that the condition for the existence of a diagonal edge is such that the triple condition, using the version stated in [19, Definition 4.1], holds.

FIGURE 16. The directed graph $\mathbf{G}((123, 132, 321))$.

Remark 5.3. Henceforth, whenever we use the term cycle, we mean a simple directed cycle. Thus we do not allow repeating edges or nodes (except the initial and terminal nodes that are the same). Also, an edge directed from node A towards node B is denoted by \overrightarrow{AB} .

Theorem 5.4. *Given an allowable pair (σ_1, σ_2) , we have that $\mathbf{G}(\sigma_1, \sigma_2)$ is acyclic.*

Proof. Throughout this proof, let $\mathbf{G} := \mathbf{G}(\sigma_1, \sigma_2)$. Additionally, let the subgraph of \mathbf{G} induced by the vertical edges be \mathbf{G}' . Note that \mathbf{G}' is acyclic. Thus, any cycle in \mathbf{G} must use at least one diagonal edge. Furthermore, if there are edges \overrightarrow{AB} and \overrightarrow{BC} in \mathbf{G}' , then there is a directed edge \overrightarrow{AC} . This implies that two consecutive vertical edges in a cycle in \mathbf{G} may be replaced by a single vertical edge. Applying this reduction repeatedly to any cycle results in a (potentially) new cycle where no two consecutive edges are vertical. We call such a cycle *reduced*. Note that \mathbf{G} is acyclic if and only if it does not contain a reduced cycle.

We first show by contradiction that \mathbf{G} does not contain a reduced 3-cycle. Assume that there is a reduced 3-cycle in \mathbf{G} . This cycle must involve exactly one diagonal edge. Suppose this diagonal edge is directed from $A := (i_1, 2)$ to $B := (i_2, 1)$ where $i_1 > i_2$. Let C and D be the nodes $(i_1, 1)$ and $(i_2, 2)$ respectively. Then we have the configuration shown in Figure 17 in \mathbf{G} . As the edge \overrightarrow{AB} is diagonal, we conclude there is a vertical edge \overrightarrow{AD} . As




FIGURE 17. Configuration in a potential 3-cycle.

the pair (σ_1, σ_2) is $(21, 12)$ -avoiding, there exists a vertical edge \overrightarrow{CB} . A reduced 3-cycle that involves \overrightarrow{AB} must have the third node participating in the cycle being either C or D . As the horizontal edges are oriented eastwards, \mathbf{G} cannot contain a reduced 3-cycle.

Next, we show that \mathbf{G} does not contain a reduced 4-cycle. Assume that \mathbf{G} contains a reduced 4-cycle. Such a cycle must involve exactly one diagonal edge, one horizontal edge and two non-consecutive vertical edges. Suppose it involves nodes A, B, C and D where the edge \overrightarrow{AB} is diagonal. Let $A := (i_1, 2)$ and $B := (i_2, 1)$ where $i_1 > i_2$. Figures 18, 19, and 20 depict all the possibilities for such a reduced 4-cycle. We analyze each case separately.




FIGURE
18. Case 1




FIGURE
19. Case 2




FIGURE
20. Case 3

In Case 1 shown in Figure 18, let E and F be the nodes $(i_2, 2)$ and $(i_1, 1)$ respectively. Let C and D be the nodes $(i_3, 1)$ and $(i_3, 2)$ respectively where $i_3 < i_2$. As \overrightarrow{AB} is diagonal, following the argument used in establishing the non-existence of a reduced 3-cycle, we infer the existence of vertical edges \overrightarrow{AE} and \overrightarrow{FB} respectively. Observe that the vertical edges \overrightarrow{FB} and \overrightarrow{BC} together imply that $\sigma_1(i_3) < \sigma_1(i_2) < \sigma_1(i_1)$. Similarly, the vertical edges \overrightarrow{DA} and \overrightarrow{AE} together imply that $\sigma_2(i_2) < \sigma_2(i_1) < \sigma_2(i_3)$. These two sequences of inequalities contradict the fact that (σ_1, σ_2) is $(123, 312)$ -avoiding. Therefore, the 4-cycle shown in Figure 18 cannot occur in \mathbf{G} .

We will treat Cases 2 and 3 in Figures 19 and 20 uniformly. Let C be the node $(i_3, 1)$ where $i_3 > i_2$. Note that $i_3 \neq i_1$ as that would give rise to a 3-cycle, which is impossible. Let E be the node $(i_2, 2)$. As (σ_1, σ_2) is $(21, 12)$ -avoiding, the edge \overrightarrow{BC} implies the existence of the vertical edge \overrightarrow{ED} . The diagonal edge \overrightarrow{AB} implies the existence of the vertical edge \overrightarrow{AE} . Now note that the nodes A, D, E are involved in a 3-cycle in \mathbf{G}' , an impossibility.

At this point we have established that \mathbf{G} does not contain any reduced 3-cycle or 4-cycle. Hence it does not contain a 3-cycle or a 4-cycle. We proceed by induction to show that \mathbf{G} does not contain a reduced k -cycle for $k \geq 5$. Consider a reduced k -cycle in \mathbf{G} whose nodes read by following the directed edges clockwise are A_1, A_2, \dots, A_k where we further assume that the edge $\overrightarrow{A_1A_2}$ is diagonal. There exist two possibilities for the edge $\overrightarrow{A_2A_3}$. If it is horizontal, then the non-existence of a 3-cycle in \mathbf{G} implies that there is a vertical edge $\overrightarrow{A_1A_3}$. This in turn yields a shorter cycle of length $k - 1$ on the nodes A_1, A_3, \dots, A_k . If, on the other hand, the edge from $\overrightarrow{A_2A_3}$ is vertical, then we are guaranteed that the edge $\overrightarrow{A_3A_4}$ is horizontal. As \mathbf{G} contains no reduced 4-cycles, we infer the existence of the vertical edge $\overrightarrow{A_1A_4}$. Thus, from the original cycle we obtain a shorter cycle of length $k - 2$ on the nodes A_1, A_4, \dots, A_k .

It follows that if \mathbf{G} contains a reduced k -cycle, then it contains a strictly shorter cycle on a subset of the original set of nodes participating in the cycle. As \mathbf{G} contains no 3-cycles or 4-cycles, it does not contain any reduced k -cycles for $k \geq 5$. This finishes the proof. \square

As $\mathbf{G}(\sigma_1, \sigma_2)$ is acyclic for an allowable pair (σ_1, σ_2) where $\sigma_1, \sigma_2 \in \mathfrak{S}_n$, we infer that there exists $\tau \in \text{SPCT}((2^n))$ such that $(\text{st}_1(\tau), \text{st}_2(\tau)) = (\sigma_1, \sigma_2)$. Using Theorem 4.9, we conclude that the number of allowable pairs (σ_1, σ_2) with $\sigma_1, \sigma_2 \in \mathfrak{S}_n$ is $(n+1)^{n-1}$. Of course, this fact was proved before [3, 4, 22] but our proof is different from the aforementioned ones.

Armed with the above insight, we will show that allowable pairs are the same as compatible pairs. Prior to that, we need the following generalization of Theorem 5.4.

Theorem 5.5. *Consider the allowable sequence $(\sigma_1, \dots, \sigma_k)$ for $k \geq 2$. The directed graph $\mathbf{G}(\sigma_1, \dots, \sigma_k)$ is acyclic.*

Proof. Throughout this proof, let $\mathbf{G} := \mathbf{G}(\sigma_1, \dots, \sigma_k)$. Let \mathbf{G}_h (respectively \mathbf{G}_v) be the subgraph of \mathbf{G} induced by the horizontal edges (respectively vertical edges). All edges in \mathbf{G}_h are oriented eastwards, whereas all edges in \mathbf{G}_v go from larger values to smaller values in terms of the permutations in our allowable sequence. Thus, \mathbf{G}_h and \mathbf{G}_v are both acyclic.

We will establish the assertion in the theorem by induction on k . The base case $k = 2$ is Theorem 5.4. Assume that the claim holds for all allowable sequences of length i where $2 \leq i < k$. Suppose for the sake of contradiction that \mathbf{G} is not acyclic.

Let \mathbf{C} be a cycle in \mathbf{G} . Observe that \mathbf{C} must contain at least one node each from the first column and the k -th column. If it did not, then \mathbf{C} would in fact be a cycle in either $\mathbf{G}(\sigma_1, \dots, \sigma_{k-1})$ or $\mathbf{G}(\sigma_2, \dots, \sigma_k)$, considered as subgraphs of \mathbf{G} in the natural manner. Our inductive hypothesis forbids either possibility.

Since diagonal edges connect nodes in consecutive columns, \mathbf{C} has nodes in all columns from 1 through k . Consider the induced subgraph of \mathbf{C} formed from nodes that belong to the first and second columns. This induced subgraph is a disjoint union of directed paths, say $\mathbf{P}_1, \dots, \mathbf{P}_m$. Suppose that, amongst the aforementioned paths, the paths $\mathbf{P}_{i_1}, \dots, \mathbf{P}_{i_r}$ are those that contain at least one node from the first column. As our cycle is simple, the initial and terminal nodes of each such path are distinct and furthermore, belong to the second column. Thus, each \mathbf{P}_{i_j} for $1 \leq j \leq r$ contains at least one diagonal edge and one horizontal edge. As $\mathbf{G}(\sigma_1, \sigma_2)$ considered as a subgraph of \mathbf{G} is acyclic, the vertical edge that connects the initial and terminal node of each \mathbf{P}_{i_j} for $1 \leq j \leq r$ acquires a unique orientation. Replacing each \mathbf{P}_{i_j} with the vertical edge connecting its initial and terminal nodes gives another cycle \mathbf{C}' in \mathbf{G} . Our construction implies that \mathbf{C}' does not involve any node in the first column. Thus, $\mathbf{G}(\sigma_2, \dots, \sigma_k)$ considered as a subgraph of \mathbf{G} contains \mathbf{C}' , which is a contradiction. This finishes the proof. \square

The acyclicity of $\mathbf{G}(\sigma_1, \dots, \sigma_k)$ implies that we can label its nodes with distinct positive integers from 1 through kn such that all edges in \mathbf{G} point from nodes with larger labels to those with smaller ones. Figure 21 demonstrates the resulting SPCT for the directed graph from Figure 16.

5	4	3
8	7	2
9	6	1

 FIGURE 21. An SPCT corresponding to $\mathbf{G}((123, 132, 321))$.

Theorem 5.6. *Given an allowable sequence $(\sigma_1, \dots, \sigma_k)$ where $\sigma_1, \dots, \sigma_k \in \mathfrak{S}_n$, there exists an SPCT τ of shape (k^n) such that $\text{st}_i(\tau) = \sigma_i$ for $1 \leq i \leq k$.*

To obtain a similar statement for SCTs, we require that σ_1 be the identity permutation in \mathfrak{S}_n . First, we use the weak Bruhat order to extend allowable pairs to allowable sequences whose first coordinate is the identity.

Lemma 5.7. *Given an allowable pair (δ_1, δ_2) where $\delta_1, \delta_2 \in \mathfrak{S}_n$, there exists an allowable sequence $(\sigma_1, \dots, \sigma_k)$ where $k \geq 3$ such that $\sigma_{k-1} = \delta_1$, $\sigma_k = \delta_2$ and $\sigma_1 = \epsilon_n$, and $\delta_1 = \epsilon_n$ if $k = 2$.*

Proof. Consider a maximal chain $\sigma_1 \prec_L \sigma_2 \prec_L \dots \prec_L \sigma_{k-1} = \delta_1$ in the weak Bruhat order, where $\sigma_1 = \epsilon_n$. We claim that $(\sigma_1, \dots, \sigma_{k-1})$ is an allowable sequence. It suffices to show that if $\gamma_1, \gamma_2 \in \mathfrak{S}_n$ satisfy $\gamma_1 \prec_L \gamma_2$, then (γ_1, γ_2) is an allowable pair.

Since $\gamma_1 \prec_L \gamma_2$, we know that $\gamma_2 = s_p \gamma_1$ for some p where p comes before $p+1$ in γ_1 . As inversions in γ_1 are also inversions in γ_2 , we infer that (γ_1, γ_2) is $(21, 12)$ -avoiding. We proceed to show that (γ_1, γ_2) is $(123, 312)$ -avoiding as well.

Suppose there are $1 \leq i < j < k \leq n$ such that $\gamma_1(i) < \gamma_1(j) < \gamma_1(k)$. We claim that we cannot have $\gamma_2(j) < \gamma_2(k) < \gamma_2(i)$. Suppose to the contrary. As γ_2 is obtained from γ_1 by swapping p and $p+1$, we have that γ_1 and γ_2 differ in exactly two positions. If neither p nor $p+1$ belong to $\{\gamma_1(i), \gamma_1(j), \gamma_1(k)\}$, then $\gamma_2(i) < \gamma_2(j) < \gamma_2(k)$ as well. If both p and $p+1$ belong to $\{\gamma_1(i), \gamma_1(j), \gamma_1(k)\}$, then $\text{stan}(\gamma_2(i)\gamma_2(j)\gamma_2(k))$ is a permutation that is distance 1 from the identity in the Hasse diagram of the weak Bruhat order on \mathfrak{S}_3 . In particular, $\text{stan}(\gamma_2(i)\gamma_2(j)\gamma_2(k)) \neq 312$. If exactly one of p or $p+1$ belongs to $\{\gamma_1(i), \gamma_1(j), \gamma_1(k)\}$, then $\text{stan}(\gamma_1(i)\gamma_1(j)\gamma_1(k)) = \text{stan}(\gamma_2(i)\gamma_2(j)\gamma_2(k))$. Thus, (γ_1, γ_2) is $(123, 312)$ -avoiding and we conclude that (γ_1, γ_2) is an allowable pair.

The preceding argument implies that $(\sigma_1, \dots, \sigma_{k-1} = \delta_1)$ is an allowable sequence and therefore, $(\sigma_1, \dots, \sigma_{k-1}, \delta_2)$ is an allowable sequence as well. \square

We have the following corollary of Lemma 5.7 which establishes that allowable pairs and compatible pairs are the same notion.

Corollary 5.8. *Given an allowable pair (σ, γ) where $(\sigma, \gamma) \in \mathfrak{S}_n$, there exists an SCT τ of shape (k^n) where $k \geq 2$ and such that $\text{st}_{k-1}(\tau) = \sigma$ and $\text{st}_k(\tau) = \gamma$.*

Proof. Construct an allowable sequence $(\sigma_1, \dots, \sigma_{k-1}, \sigma_k)$ where $\sigma_1 = \epsilon_n$, $\sigma_{k-1} = \sigma$ and $\sigma_k = \gamma$. Then $\mathbf{G}(\sigma_1, \dots, \sigma_k)$ is acyclic. The conclusion now follows. \square

6. FINAL REMARKS

- (1) As mentioned earlier, the representation theory of $H_n(0)$ in type A is related to the algebra of quasisymmetric functions. Indeed, there is a map ch called the *quasisymmetric characteristic* [13] that associates a quasisymmetric function to an $H_n(0)$ -module. Mimicking the proof of [45, Theorem 5.2], one can establish that

$$ch(\mathbf{S}_\alpha) = \sum_{\tau \in \text{SPCT}(\alpha)} F_{\text{comp}(\tau)}.$$

Here $\text{comp}(\tau)$ is the composition of n naturally associated with the subset $\text{Des}(\tau)$ of $[n-1]$ and F_γ is the fundamental quasisymmetric function indexed by a composition γ [43, Chapter 7]. It would be interesting to investigate the algebraic/combinatorial properties of $ch(\mathbf{S}_\alpha)$ in addition to the representation-theoretic properties of the \mathbf{S}_α .

- (2) In the introduction we mentioned that the study of ascents and descents on labeled binary trees is tied to the enumeration of chambers in various Coxeter deformations. The number of tableaux in $\text{SPCT}((2^n))$ is $n! \text{Cat}_n$, which equals the number of regions in the Catalan arrangement in \mathbb{R}^n defined by the hyperplanes $x_i - x_j = 0, \pm 1$ for $1 \leq i < j \leq n$. The sink tableaux of shape (2^n) are distinguished representatives of the equivalence classes under $\sim_{(2^n)}$ and by Theorem 4.9, there are $(n+1)^{n-1}$ many of them. This is the number of regions in the Shi arrangement in \mathbb{R}^n defined by the hyperplanes $x_i - x_j = 0, 1$ for $1 \leq i < j \leq n$. In view of this, it would be interesting to understand how the coarsening of regions of the Catalan arrangement into the regions of the Shi arrangement relates to grouping of tableaux in $\text{SPCT}((2^n))$ into equivalence classes under $\sim_{(2^n)}$. Our 0-Hecke operators translate into operators on regions of the Catalan arrangement and this viewpoint merits further study.

ACKNOWLEDGEMENTS

We would like to thank Sara Billey, Ira Gessel and Sean Griffin for extremely helpful discussions. We would also like to thank the anonymous referees for their valuable feedback and suggestions.

REFERENCES

- [1] E. ALLEN, J. HALLAM AND S. MASON, *Dual immaculate quasisymmetric functions expand positively into Young quasisymmetric Schur functions*, J. Combin. Theory Ser. A 157 (2018) 70–108.
- [2] S. ASSAF AND D. SEARLES, *Kohnert tableaux and a lifting of quasi-Schur functions*, J. Combin. Theory Ser. A 156 (2018) 85–118.
- [3] M. ATKINSON, S. LINTON AND L. WALKER, *Priority queues and multisets*, Electron. J. Combin. 2 (1995).
- [4] M. ATKINSON AND M. THIYAGARAJAH, *The permutational power of a priority queue*, BIT 33 (1993) 2–6.
- [5] C. BERG, N. BERGERON, F. SALIOLA, L. SERRANO AND M. ZABROCKI, *A lift of the Schur and Hall-Littlewood bases to non-commutative symmetric functions*, Canad. J. Math. 66 (2014) 525–565.

- [6] C. BERG, N. BERGERON, F. SALIOLA, L. SERRANO AND M. ZABROCKI, *Indecomposable modules for the dual immaculate basis of quasi-symmetric functions*, Proc. Amer. Math. Soc. 143 (2015) 991–1000.
- [7] O. BERNARDI, *Deformations of the braid arrangement and trees*, Adv. Math. 335 (2018) 466–518.
- [8] C. BESSENRODT, V. TEWARI AND S. VAN WILLIGENBURG, *Littlewood-Richardson rules for symmetric skew quasisymmetric Schur functions*, J. Combin. Theory Ser. A 137 (2016) 179–206.
- [9] A. BJÖRNER AND F. BRENTI, *Combinatorics of Coxeter Groups*, Springer, 2005.
- [10] R. CARTER, *Representation theory of the 0-Hecke algebra*, J. Algebra 104 (1986) 89–103.
- [11] S. CORTEEL, D. FORGE AND V. VENTOS, *Bijections between affine hyperplane arrangements and valued graphs*, European J. Combin. 50 (2015) 30–37.
- [12] B. DRAKE, *An inversion theorem for labeled trees and some limits of areas under lattice paths*, Phd. Dissertation, Brandeis University (2008).
- [13] G. DUCHAMP, D. KROB, B. LECLERC AND J-Y. THIBON, *Fonctions quasi-symétriques, fonctions symétriques noncommutatives, et algèbres de Hecke à $q = 0$* , C. R. Math. Acad. Sci. Paris 322 (1996) 107–112.
- [14] D. FORGE, *Linial arrangements and local binary search trees*, arxiv: <http://arxiv.org/abs/1411.7834>.
- [15] J. FRAME, G. ROBINSON AND R. THRALL, *The hook graphs of the symmetric groups*, Canad. J. Math. 6 (1954) 316–324.
- [16] I. GESSEL, *Oberwolfach Reports (Enumerative Combinatorics)*, https://www.mfo.de/document/1410/OWR_2014_12.pdf, 2014, Page 709.
- [17] I. GESSEL, S. GRIFFIN AND V. TEWARI, *Labeled plane binary trees and Schur-positivity*, arxiv: <https://arxiv.org/abs/1706.03055>.
- [18] J. GILBEY AND L. KALIKOW, *Parking functions, valet functions and priority queues*, Discrete Math. 197/198 (1999) 351–373.
- [19] J. HAGLUND, K. LUOTO, S. MASON AND S. VAN WILLIGENBURG, *Quasisymmetric Schur functions*, J. Combin. Theory Ser. A 118 (2011) 463–490.
- [20] J. HAGLUND, K. LUOTO, S. MASON AND S. VAN WILLIGENBURG, *Refinements of the Littlewood-Richardson rule*, Trans. Amer. Math. Soc. 363 (2011) 1665–1686.
- [21] J. HAGLUND, S. MASON AND J. REMMEL, *Properties of the nonsymmetric Robinson-Schensted-Knuth algorithm*, J. Algebraic Combin. 38 (2013) 285–327.
- [22] A. HAMEL, *Priority queue sorting and labeled trees*, Ann. Comb. 7 (2003) 49–54.
- [23] J. HUANG, *0-Hecke algebra actions on coinvariants and flags*, J. Algebraic Combin. 40 (2014) 245–278.
- [24] J. HUANG, *0-Hecke algebra action on the Stanley-Reisner ring of the Boolean algebra*, Ann. Comb. 19 (2015) 293–323.
- [25] J. HUANG, *A tableau approach to the representation theory of 0-Hecke algebras*, Ann. Comb. 20 (2016) 831–868.
- [26] J. HUANG AND B. RHOADES, *Ordered set partitions and the 0-Hecke algebra*, Algebraic Combinatorics 1 (2018) 47–80.
- [27] N. JING AND Y. LI, *A lift of Schur’s Q -functions to the peak algebra*, J. Combin. Theory Ser. A 135 (2015) 268–290.
- [28] L. KALIKOW, *Symmetries in trees and parking functions*, Discrete Math. 256 (2002) 719–741.
- [29] S. KÖNIG, *The decomposition of 0-Hecke modules associated to quasisymmetric Schur functions*, arxiv: <https://arxiv.org/abs/1711.08737>.
- [30] A. LAUVE AND S. MASON, *$QSym$ over Sym has a stable basis*, J. Combin. Theory Ser. A 118 (2011) 1661–1673.
- [31] K. LUOTO, S. MYKYTIUK AND S. VAN WILLIGENBURG, *An introduction to quasisymmetric Schur functions. Hopf algebras, quasisymmetric functions, and Young composition tableaux*, SpringerBriefs in Mathematics, Springer, 2013.

- [32] S. MASON, *An explicit construction of type A Demazure atoms*, J. Algebraic Combin. 29 (2009) 295–313.
- [33] S. MASON, *A decomposition of Schur functions and an analogue of the Robinson-Schensted-Knuth algorithm*, Sémin. Lothar. Combin. 57 (2006).
- [34] A. MATHAS, *Iwahori-Hecke algebras and Schur algebras of the symmetric group*, American Mathematical Society, 1999.
- [35] C. MONICAL, *Set-valued skyline fillings*, arxiv: <https://arxiv.org/abs/1611.08777>.
- [36] P. NORTON, *0-Hecke algebras*, J. Aust. Math. Soc. 27 (1979) 337–357.
- [37] E. OĞUZ, *A note on Jing and Li's type B quasischur functions*, Ann. Comb. (2018), to appear.
- [38] I. PAK, *Tree bijections*, slides: <http://www.math.ucla.edu/~pak/lectures/bij-trees-talk1.pdf>.
- [39] I. PAK AND A. POSTNIKOV, *Enumeration of trees and one amazing representation of S_n* , Proc. FPSAC'96 Conf., Minneapolis, MN.
- [40] B. SAGAN, *The symmetric group. Representations, combinatorial algorithms, and symmetric functions. 2nd ed.*, Springer, 2001.
- [41] I. SCHUR, *Über eine Klasse von Matrizen, die sich einer gegebenen Matrix zuordnen lassen*. Dissertation, Berlin, 1901. In: I. Schur, Gesammelte Abhandlungen I, Springer, Berlin, (1973) 1–70.
- [42] D. SEARLES, *Polynomial bases: positivity and Schur multiplication*, Trans. Amer. Math. Soc. (2018), to appear.
- [43] R. STANLEY, *Enumerative Combinatorics vol. 2*, Cambridge University Press, 1999.
- [44] V. TEWARI, *Gessel polynomials, rooks, and extended Linial arrangements*, arxiv: <https://arxiv.org/abs/1604.06894>.
- [45] V. TEWARI AND S. VAN WILLIGENBURG, *Modules of the 0-Hecke algebra and quasisymmetric Schur functions*, Adv. Math. 285 (2015) 1025–1065.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195, USA
E-mail address: vasut@math.washington.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF BRITISH COLUMBIA, VANCOUVER, BC V6T 1Z2,
 CANADA
E-mail address: steph@math.ubc.ca