PARKING FUNCTIONS, DIAGONAL COINVARIANTS, AND THE DELTA CONJECTURE

Athina Avrantini

A THESIS

in

Mathematics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Master of Arts

2026

Supervisor of Thesis

James Haglund, Professor of Mathematics

Graduate Group Chairperson

Angela Gibney, Presidential Professor of Mathematics

ABSTRACT

PARKING FUNCTIONS, DIAGONAL COINVARIANTS, AND
THE DELTA CONJECTURE

Athina Avrantini

James Haglund

This thesis investigates the connection between the representation theory of the diagonal coinvariant ring DR_n and the combinatorics of parking functions. We provide a self-contained exposition of the Shuffle Theorem, which expresses the bigraded Frobenius series of DR_n as a sum over standard parking functions. We then detail the Schedules Formula, which arises from a recursive insertion algorithm to organize parking functions into trees indexed by permutations. Finally, we extend this framework to state the Delta Conjecture, a generalization of the Shuffle Theorem involving the Macdonald eigenoperators $\Delta'_{e_{n-k-1}}$. We introduce valley-marked parking functions and generalize the insertion algorithm to ordered set partitions, establishing the combinatorial structure necessary to study the Delta Conjecture.

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

Parking functions form a central meeting point of enumerative combinatorics, symmetric functions, and the representation theory of the symmetric group. Originally introduced in the context of computer science to resolve collisions in hash tables, they have since appeared throughout algebraic combinatorics and algebraic geometry, notably in connection with diagonal harmonics and the geometry of the Hilbert scheme. The guiding theme of this thesis is the network of results and conjectures that relate (i) statistics on parking functions, (ii) graded and bigraded Frobenius series, and (iii) explicit formulas in the algebra of symmetric functions. Our main goal is to explain, in a self-contained way, how these perspectives fit together in the story surrounding the Shuffle Theorem, the Schedules Formula, and the Delta Conjecture.

Chapter 2 introduces parking functions, first from the classical preference-sequence viewpoint and then through their equivalent description as labeled Dyck paths. In the Dyck-path model we define the statistics area, dinv, the reading word $\sigma(\pi)$, and ides(π), which provide the basic combinatorial data appearing in later generating functions.

Chapter 3 provides the representation-theoretic background needed to interpret these generating functions. We review representations and characters of finite groups, specialize to S_n, and recall partitions, tableaux, and Schur functions. We then introduce the Hilbert and Frobenius series for graded S_n-modules.

Chapter 4 introduces the diagonal coinvariant ring

$$DR_n = \mathbb{C}[\mathbf{x_n}, \mathbf{y_n}]/I_n(x, y),$$

a fundamental bigraded S_n-module whose dimension is $(n+1)^{n-1}$. We describe its bigraded Hilbert and Frobenius series, state Haiman's theorem identifying the Frobenius series with ∇e_n, and then state the Shuffle Theorem, which gives a purely combinatorial formula for the same Frobenius series as a weighted sum over parking functions. As an immediate corollary, one obtains a compact expression for the bigraded Hilbert series as a (q, t)-enumerator of parking functions.

Chapter 5 develops an alternative approach to this Hilbert series. We describe an insertion procedure that builds parking functions recursively and organizes them into families indexed by permutations. The branching data of this construction is encoded by a schedule statistic on permutations, and iterating the insertion step leads to the Schedules Formula, which expresses $\mathcal{H}(DR_n; q, t)$ as a sum over S_n with a simple product contribution from each permutation.

Chapter 6 turns to the Delta Conjecture, introduced by Haglund, Remmel, and Wilson as a broad generalization of the Shuffle Theorem. The original Delta Conjecture posits the equality of a symmetric function side and two distinct combinatorial sides: the "rise version" and the "valley version." While the rise version has recently been proven by two different groups in (DM22) and (BHM$^+$23), the valley version remains an active open problem. In this thesis, "Delta Conjecture" will refer exclusively on the latter version, in which ordinary parking functions are replaced by "valley-marked" parking functions with a modified dinv statistic, and the relevant indexing objects become ordered set partitions. We extend the schedules perspective to this marked setting, obtaining a parallel product formula that efficiently packages the (q, t)-weights arising from the marked insertion process.

CHAPTER 2

PARKING FUNCTIONS

Definition 2.1. A parking function of length n is a sequence of integers between 1 and n with the property that for each $1 \leq i \leq n$, the sequence contains at least i values that are at most i.

Parking functions can be intuitively described as follows: suppose n cars are attempting to park along a one-way street with n parking spaces, numbered sequentially from 1 to n. Each car i has a preferred spot a_i. Cars park in their preferred spot if it is available; otherwise, they park in the first available spot after their preferred position. A sequence of preferences $(a_1, \ldots, a_n)$ is then a parking function if every car can find a parking spot.

Theorem 2.1. *The number of parking functions on n cars is $(n+1)^{n-1}$.*

Proof. Consider an extended scenario where the street is circular, containing $n+1$ parking spots. Each car selects a preferred spot from these, giving $(n+1)^n$ total choices. All cars will inevitably park due to the circular structure, leaving exactly one empty spot. A configuration corresponds to a valid parking function precisely when the empty spot is position $n+1$. By symmetry, each spot is equally likely to remain empty; hence, dividing by $n+1$, the number of valid parking functions is $(n+1)^{n-1}$. □

Another standard description of parking functions involves labeled Dyck paths.

Definition 2.2. A Dyck path of length n is a lattice path from $(0,0)$ to (n,n) only using steps of the form $(0,1)$ and $(1,0)$, never going below the diagonal $x=y$.

Starting from a Dyck path $\mathcal{D}$, we can obtain a parking function by labeling the squares immediately to the right of the vertical steps of $\mathcal{D}$ with the numbers 1 through n (representing the cars), subject to the constraint that if car i is placed immediately on top of car j, then $i<j$.

It is convenient to organize the squares of the $n \times n$ grid into diagonals: the squares intersected by the line $y=x$ form the 0-diagonal (main diagonal); those immediately above form the 1-diagonal;

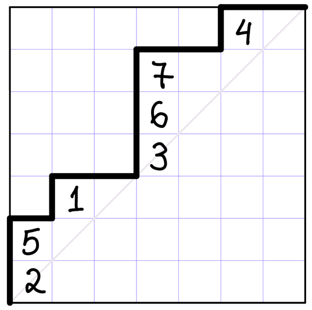

Figure 2.1: The parking function $(2, 1, 4, 6, 1, 4, 4)$ as a Dyck path

then the 2-diagonal, and so on.

We associate to each parking function π four statistics: $\text{area}(\pi)$, $\text{dinv}(\pi)$, the "word" $\sigma(\pi)$, and the "inverse descent set" $\text{ides}(\pi)$.

The area of π is simply the number of cars that fall strictly between the Dyck path and the main diagonal $x = y$. In Figure 2.1 we have area $= 6$.

To obtain the word of π, denoted by $\sigma(\pi)$, we write down the cars of π from highest to lowest diagonal, and from right to left within each diagonal. In Figure 2.1 we have $\sigma(\pi) = 7461532$. Then, we define $\text{ides}(\pi) = \{i : i + 1 \text{ occurs to the left of } i \text{ in } \sigma(\pi)\}$. In Figure 2.1, ides $= \{2, 3, 5, 6\}$.

Finally, dinv couts the number of "diagonal inversions". A pair of cars (a, b) with $a < b$ constitute a diagonal inversion if either i) a and b are on the same diagonal and b is further to the right or ii) b is one diagonal above and further to the left of a. In Figure 2.1, the diagonal inversions are $(1, 4)$, $(1, 6)$, $(3, 5)$, $(4, 7)$, and $(5, 6)$, so dinv $= 5$.

These statistics are required to state the Shuffle Theorem. Before doing so, we need to also cover some background material on representation theory.

CHAPTER 3

REPRESENTATION THEORY OF S_n

3.1. Preliminaries

Definition 3.1. Let G be a finite group. A representation of G is a vector space V together with a group homomorphism $\rho : G \to GL(V)$. The degree of a representation is the dimension of the vector space V.

It is often convenient to think of representations of G as G-modules. The trace of the matrix corresponding to a representation ρ does not depend on the choice of basis. It is denoted by χ_ρ and called the *character* of ρ. Characters are always class functions, meaning that they are constant on conjugacy classes of G.

Let (V, ρ) be a representation of G and W a linear subspace of V that is G-stable in the sense that $\rho(g)w \in W$ for all $g \in G, w \in W$. Then the restriction of ρ to W gives rise to a *subsrepresentation* $(W, \rho_{|W})$.

Definition 3.2. A representation (V, ρ) is called irreducible if it has exactly two subrepresentations, namely the trivial subspace $\{0\}$ and V itself. If V has proper nontrivial subrepresentations then it is called reducible.

In light of the following theorem, the question of classifying all representations of a finite group reduces to the more tractable task of classifying just the irreducible ones.

Theorem 3.1 (Maschke's Theorem)**.** *Let G be a finite group and V a representation over $\mathbb{C}$. Then*

$$V = V_1 \oplus V_2 \oplus \cdots \oplus V_k$$

where each V_i is an irreducible subrepresentation of V.

It then follows from the Jordan–Hölder theorem that, while the decomposition of V into a direct sum of irreducible subrepresentations may not be unique, the irreducible pieces have well-defined multiplicities. Given an irreducible character χ, we will denote its multiplicity in V by $\mathrm{Mult}(\chi, V)$.

Theorem 3.2. *The number of non-isomorphic irreducible representations of G is equal to the number of conjugacy classes of G. Furthermore, the irreducible characters of G form an orthonormal basis for the space of class functions on G.*

In the case of $G = S_n$, the conjugacy classes are determined by cycle type, and so they are in bijection with the partitions of n.

3.2. Partitions and tableaux

Definition 3.3. A partition of n is a weakly decreasing sequence $\lambda = (\lambda_1, \ldots, \lambda_l)$ of nonnegative integers with $\sum_{i=1}^{l} \lambda_i = n$. We write $\lambda \vdash n$.

We visualize a partition λ by means of Young diagrams: a Young diagram is a collection of boxes arranged in left-aligned rows such that the number of boxes in each row decreases weakly from top to bottom. If $\lambda = (\lambda_1, \ldots, \lambda_l)$ is a partition, the corresponding Young diagram has λ_i boxes on row i for $1 \leq i \leq l$. We denote that diagram by $dg(\lambda)$.

Definition 3.4. Given a partition λ, its conjugate partition λ' is the one obtain by flipping $dg(\lambda)$ along its main diagonal.

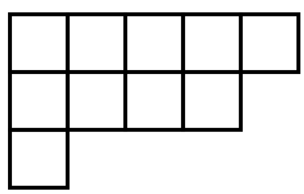

Figure 3.1: Young diagram corresponding to $\lambda = (5, 4, 1)$

Definition 3.5. A semi-standard Young tableau (SSYT) of shape λ is a filling in of the boxes of $dg(\lambda)$ with nonnegative integers such that the entries in each row are weakly increasing and the entries in each column are strictly increasing. If we require that the entries in each row are also strictly increasing, we get a standard Young tableau (SYT).

1	2	4	7	8
3	5	6	9	
10				

Figure 3.2: Young tableaux of shape $\lambda = (5, 4, 1)$. The numbers in the boxes increase in every row and every column

Let T be a SSYT whose entries are in the set $\{1, \ldots, n\}$. For each $1 \leq i \leq n$, let m_i be the number of boxes of T that are labeled i. We then associate to T the monomial

$$x^T = x_1^{m_1} \cdots x_n^{m_n}.$$

Given a partition λ, its *Schur function* s_λ is given by

$$s_\lambda(x_1, \ldots, x_n) = \sum_{T \text{ has shape } \lambda} x^T,$$

where the sum runs only over tableaux with entries in $\{1, \ldots, n\}$.

As discussed earlier, there is a bijective correspondence between partitions of n and irreducible representations of S_n; the latter are also referred to as Specht modules. Denote by S^λ the Specht module corresponding to λ. We can define an inner product on $\mathcal{R}_n$, the ring of all irreducible representations of S_n, via

$$\langle S^\lambda, S^\mu \rangle = \delta_{\lambda\mu} = \begin{cases} 1 & \text{if } \lambda = \mu \\ 0 & \text{if } \lambda \neq \mu \end{cases}.$$

The multiplicity of an irreducible compontent S^λ in a representation V is just $\langle V, S^\lambda \rangle$.

3.3. Hilbert and Frobenius series

So far, we have considered representations as single, finite-dimensional vector spaces. However, in algebraic combinatorics, we frequently encounter infinite-dimensional spaces that are naturally graded by degree, such as polynomial rings. To understand the structure of a graded S_n-module $V = \bigoplus_i V^i$, we need to track both the dimension and the representation type of each homogeneous component V^i.

To do this, we encode this information into generating functions. We define the Hilbert series and Frobenius series of V as:

$$\mathcal{H}(V;q) = \sum_{i \geq 0} q^i \dim(V^i)$$

and

$$\mathcal{F}(V;q) = \sum_{i \geq 0} q^i \mathrm{Frob}(V^i) = \sum_{i \geq 0} q^i \left(\sum_{\lambda} \langle \chi_{V^i}, \chi^\lambda \rangle s_\lambda \right).$$

The power of the Frobenius series lies in the Frobenius characteristic map, which provides an isometric isomorphism between the ring of class functions of S_n and the ring of symmetric functions. By mapping the irreducible character χ^λ to the Schur function s_λ, the Frobenius series translates the representation theory of the symmetric group into the computable algebra of symmetric functions. In the next chapter, we will use these tools to analyze the structure of the diagonal coinvariant ring.

CHAPTER 4

THE SHUFFLE THEOREM

4.1. Coinvariants

Let $\mathbb{C}[\mathbf{x_n}, \mathbf{y_n}]$ denote the polynomial ring $\mathbb{C}[x_1, \ldots, x_n, y_1, \ldots, y_n]$. Given a polynomial $f \in \mathbb{C}[\mathbf{x_n}, \mathbf{y_n}]$, the symmetric group S_n acts diagonally by permuting the x- and y- variables in the same way, that is, for all $\sigma \in S_n$, we have

$$\sigma f(x_1, \cdots, x_n, y_1, \cdots, y_n) = f(x_{\sigma(1)}, \ldots, x_{\sigma(n)}, y_{\sigma(1)}, \ldots, y_{\sigma(n)}).$$

Definition 4.1. The ring of diagonal coinvariants is the quotient

$$DR_n = \mathbb{C}[\mathbf{x_n}, \mathbf{y_n}]/I_n(x, y)$$

where $I_n(x, y)$ is the ideal generated by the homogeneous polynomials in $\mathbb{C}[\mathbf{x_n}, \mathbf{y_n}]$ without constant term that are invariant under the S_n-action above.

DR_n was first studied by Haiman, who proved it has dimension $(n+1)^{n-1}$; as we saw earlier, this equals the number of parking functions on n cars.

Definition 4.2. Given a subspace $W \subseteq \mathbb{C}[\mathbf{x_n}, \mathbf{y_n}]$, we define its bigraded Hilbert series as

$$\mathcal{H}(W; q, t) = \sum_{i,j \geq 0} t^i q^j \dim(W^{(i,j)})$$

where $W^{(i,j)}$ is the subspace consisting of bi-homogeneous elements of W with degree i in the x variables and degree j in the y variables.

Note that $W = \bigoplus_{i,j \geq 0} W^{(i,j)}$. The Hilbert series encodes information about the dimensions of bigraded subspaces. Next, we will define the bigraded Frobenius series, which will analogously give us information about the decomposition of these subspaces as S_n-representations.

Definition 4.3. Given a subspace $W \subseteq \mathbb{C}[\mathbf{x_n}, \mathbf{y_n}]$, we define its bigraded Frobenius series as

$$\mathcal{F}(W; q, t) = \sum_{i,j \geq 0} t^i q^j \mathrm{Frob}(W^{(i,j)}) = \sum_{i,j \geq 0} t^i q^j \sum_{\lambda \vdash n} s_\lambda \mathrm{Mult}(\chi^\lambda, W^{(i,j)}).$$

This is a symmetric function in infinitely many variables with coefficients in $\mathbb{Z}_{\geq 0}[q, t]$.

Next, we show that we can recover the bigraded Hilbert series from the bigraded Frobenius series:

Theorem 4.1. *We have that*

$$\mathcal{H}(W; q, t) = \langle \mathcal{F}(W; q, t), h_1^n \rangle$$

where $h_1^n = \sum_{\lambda \vdash n} s_\lambda f^\lambda$ and f^λ is the number of SYT of shape λ.

Proof. Recall that Schur functions are orthonormal with respect to the Hall inner product, so $\langle s_\lambda, s_\mu \rangle = \delta_{\lambda,\mu}$. Then, we calculate

$$\begin{aligned} \langle \mathcal{F}(W; q, t), h_1^n \rangle &= \left\langle \sum_{i,j \geq 0} t^i q^j \sum_{\lambda \vdash n} s_\lambda \mathrm{Mult}(\chi^\lambda, W^{(i,j)}), \sum_{\mu \vdash n} s_\mu f^\mu \right\rangle \\ &= \sum_{i,j \geq 0} t^i q^j \sum_{\lambda, \mu \vdash n} \mathrm{Mult}(\chi^\lambda, W^{(i,j)}) f^\mu \langle s_\lambda, s_\mu \rangle \\ &= \sum_{i,j \geq 0} t^i q^j \sum_{\lambda \vdash n} \mathrm{Mult}(\chi^\lambda, W^{(i,j)}) f^\lambda. \end{aligned}$$

By Maschke's Theorem, we have that

$$W^{(i,j)} \cong \bigoplus_{\lambda \vdash n} \mathrm{Mult}(\chi^\lambda, W^{(i,j)}) W_\lambda^{(i,j)}$$

where $W_\lambda^{(i,j)}$ is the irreducible representation corresponding to λ. Hence,

$$\dim(W^{(i,j)}) = \mathrm{Mult}(\chi^\lambda, W^{(i,j)}) f^\lambda,$$

so we obtain

$$\langle \mathcal{F}(W;q,t), h_1^n \rangle = \sum_{i,j\geq 0} t^i q^j \dim(W^{(i,j)}) = \mathcal{H}(W;q,t).$$

□

Example 4.1.1. For $n = 3$, we have the partitions $(3), (2,1)$, and $(1,1,1)$. Let $W^{(0,0)} = S^{(3)}$, $W^{(1,0)} = S^{(2,1)} \oplus S^{(2,1)}$ and $W^{(0,1)} = S^{(1,1,1)}$, where S^λ denotes the Specht module corresponding to λ. If

$$W = W^{(0,0)} \oplus W^{(1,0)} \oplus W^{(0,1)}$$

then the Hilbert and Frobenius series are given by:

$$\mathcal{H}(W;q,t) = 1 + 4t + q,$$

$$\mathcal{F}(W;q,t) = s_{(3)} + 2ts_{(2,1)} + qs_{(1,1,1)}.$$

Moreover, we have $h_1^3 = s_{(3)} + 2s_{(2,1)} + s_{(1,1,1)}$ and therefore

$$\langle \mathcal{F}(W;q,t), h_1^3 \rangle = 1 + 4t + q$$

which indeed recovers the Hilbert series.

4.2. The Nabla operator

This operator plays a crucial role in the developments relating Macdonald polynomials to symmetric group representation theory.

Let $\Lambda = \bigoplus_{n\geq 0} \Lambda^{(n)}$ be the algebra of symmetric functions in a formal infinite alphabet $X = x_1, x_2, \ldots$ with coefficients in the field of rational functions $\mathbb{Q}(q,t)$, graded by degree. For each partition $\mu \vdash n$ there is a *modified Macdonald polynomial* $\widetilde{H}_\mu(X;q,t) \in \Lambda^{(n)}$. The family $\{\widetilde{H}_\mu\}_{\mu\vdash n}$ forms a basis for $\Lambda^{(n)}$ over $\mathbb{Q}(q,t)$. For this thesis we will not need their explicit construction, only that they are an orthogonal basis. For a more detailed exposition, the reader should consult

(Mac95). The linear operator ∇ is defined in terms of this orthogonal basis by setting

$$\nabla \widetilde{H}_\mu = t^{\eta(\mu)} q^{\eta(\mu')} \widetilde{H}_\mu$$

where

$$\eta(\mu) = \sum_{i \geq 1} (i-1)\mu_i$$

for a partition $\mu = (m_1, \ldots, \mu_l) \vdash n$ and μ' its conjugate partition.

Theorem 4.2 (Haiman)**.** *Let e_n denote the nth elementary symmetric polynomial. Then, we have that*

$$\mathcal{F}(DR_n; q, t) = \nabla e_n[X].$$

The proof of Theorem 4.2 is highly non-trivial. It relies on deep connections with algebraic geometry, specifically the geometry of the isospectral Hilbert scheme of n points in the complex plane, $Hilb^n(\mathbb{C}^2)$. Haiman showed that DR_n can be realized as the ring of global sections of a specific vector bundle on this scheme (Hai02). The celebrated $(n+1)^{n-1}$ dimension formula is a corollary of this geometric equivalence.

4.3. The Shuffle Theorem

The Shuffle Theorem was conjectured by Haglund, Haiman, Loehr, Remmel, and Ulyanov (HHL$^+$04) and proven by Carlson and Mellit (CM18). It predicts a combinatorial formula involving parking functions for the Frobenius series of the diagonal coinvariant ring. The theorem supplies a three-way bridge among parking-function enumeration, symmetric-function theory, and the bigraded representation theory of DR_n.

Theorem 4.3 (The Shuffle Theorem)**.**

$$\mathcal{F}(DR_n; q, t) = \sum_{\pi \in PF_n} q^{dinv(\pi)} t^{area(\pi)} F_{ides(\pi)},$$

where

$$F_{ides(\pi)}(x_1, x_2, \dots) = \sum_{\substack{i_1 \le i_2 \le \cdots \le i_n \\ j \in ides(\pi) \implies i_j < i_{j+1}}} x_{i_1} x_{i_2} \cdots x_{i_n}$$

is Gessel's fundamental basis for the ring of quasi-symmetric functions.

While the full theory of quasi-symmetric functions is beyond the scope of this thesis, the role of Gessel's basis (Ges84) here is simply to translate the combinatorial ides(π) statistic into an algebraic object. By packaging the descent information of the parking function's reading word into a polynomial, we can use the inner product on symmetric functions to extract combinatorial data.

Now using the result of Theorem 4.1 we obtain the following

Corollary 4.3.1.

$$\mathcal{H}(DR_n; q, t) = \sum_{\pi \in PF_n} q^{\mathrm{dinv}(\pi)} t^{\mathrm{area}(\pi)}.$$

The Shuffle Theorem gets its name due to its connection to shuffles:

Definition 4.4. Given an ordered sequence of distinct positive integers $r_1, \dots, r_m \le n$, we say $\sigma \in S_n$ is an $(r_1, \dots, r_m)$-shuffle if r_i occurs before r_{i+1} in $\sigma_1\sigma_2\cdots\sigma_n$ for all i. Given a composition μ of n, we say σ is a μ-shuffle if it is a shuffle of each of the sequences $(1, 2, \dots, \mu_1), (\mu_1+1, \cdots, \mu_1+\mu_2), \dots$ of lengths μ_1, μ_2, and so on.

Example 4.3.1. Let $\mu = (2, 2)$. A permutation $\sigma \in S_4$ is an μ-shuffle it is a shuffle of both $(1, 2)$ and $(3, 4)$. That means that 1 occurs before 2 and 3 before 4 in $\sigma_1\sigma_2\sigma_3\sigma_4$. These are exactly:

$$1234 \quad 1324 \quad 1342 \quad 3124 \quad 3142 \quad 3412.$$

Given a composition $\mu = (\mu_1, \dots, \mu_k)$ of n, we consider the homogeneous symmetric function

$$h_\mu = h_{\mu_1} \cdots h_{\mu_k}$$

where for a nonnegative integer r, h_r is defined by

$$h_r(x_1, x_2, \dots) = \sum_{i_1 \leq \cdots \leq i_r} x_{i_1} \cdots x_{i_r}.$$

Then, given a parking function π,

$$\langle F_{ides(\pi)}, h_\mu \rangle = \begin{cases} 1, & \text{if the reading word } \sigma(\pi) \text{ is a } \mu\text{-shuffle} \\ 0, & \text{otherwise.} \end{cases}$$

Hence, by Theorems 4.2 and 4.3, we have another version of the Shuffle Theorem involving the graded dimension of the invariants under the Young subgroup

$$\langle \nabla e_n, h_\mu \rangle = \dim_{q,t}((DR_n)^{S_\mu}) = \sum_{\substack{\pi \in PF_n, \\ \sigma(\pi) \text{ is a } \mu\text{-shuffle}}} q^{\mathrm{dinv}(\pi)} t^{\mathrm{area}(\pi)}$$

where

$$S_\mu = S_{\mu_1} \times \cdots \times S_{\mu_k} \subset S_n$$

is the Young subgroup of S_n corresponding to μ. Combinatorially, S_μ is the subgroup of permutations that independently permute the elements within the disjoint sets $\{1, \dots, \mu_1\}$, $\{\mu_1 + 1, \dots, \mu_1 + \mu_2\}$ and so on. Taking the invariants under this subgroup precisely isolates the parking functions whose reading words are μ-shuffles.

CHAPTER 5

THE SCHEDULES FORMULA

The Shuffle Theorem gave us a formula for the Hilbert series of the diagonal coinvariant ring DR_n by summing over all parking functions in PF_n. However, this sum can be difficult to compute directly for large n. In this chapter, we introduce a method to group parking functions into families generated by permutations. This approach leads to the Schedules Formula, which provides an alternative way to compute $\mathcal{H}(DR_n; q, t)$.

5.1. Schedules

Definition 5.1. Given $\sigma \in S_n$, its runs $(r_1(\sigma), \ldots, r_k(\sigma))$ are its maximal consecutive increasing subsequences in order from left to right. If σ has k runs, we set $r_{k+1}(\sigma) = (0)$. If σ_i is in the jth run of σ, we define $sch_i(\sigma)$ as the number of elements in $r_j(\sigma)$ that are greater than σ_i together with the number of elements in $r_{j+1}(\sigma)$ that are less than σ_i. Then, the schedule of σ is the sequence $sch(\sigma) = (sch_1(\sigma), \ldots, sch_n(\sigma))$.

Example 5.1.1. For $\sigma = 3512746$, we have

$$r_1(\sigma) = 35 \quad r_2(\sigma) = 127 \quad r_3(\sigma) = 46 \quad r_4(\sigma) = 0$$

and

$$sch(\sigma) = (3, 2, 2, 1, 2, 2, 1).$$

This definition is designed so that each $sch_i(\sigma)$ measures a local "room for choice" across the unique descent that separates the run of σ_i from the next run: looking within the run, larger entries can be swapped past without breaking increasing order, while looking into the next run, smaller entries can be pulled left across the descent. These two freedoms are independent and add.

Definition 5.2. We also define the major index of a permutation

$$maj(\sigma) = \sum_{\sigma_i > \sigma_{i+1}} i = \sum_{i=1}^{k-1} \sum_{j=1}^{i} |r_j(\sigma)|,$$

where k is the number of runs in σ.

There is a connection between schedules and parking functions. The idea behind this name is that the schedule tells you the order (or schedule) of cars being inserted to build up a parking function. We make this more precise in the next section.

5.2. Insertion

We describe an algorithm for creating new parking functions from old ones.

Insertion Algorithm 5.1. *Let $\pi \in PF_n$ a parking function, c a car not already present in π, and k an integer such that the $(k+1)$-diagonal of π is empty and the k-diagonal contains no car smaller than c. We present three different ways of inserting c into π to obtain a new parking function in PF_{n+1}:*

(a) Let $a < c$ be a car in the $(k-1)$-diagonal of π. Move all cars in a higher row than a one step up and one step right. Place c directly above a.

(b) Let $b > c$ be a car in the k-diagonal of π. Move all cars in a higher row than b one step up and one step right. Place c one step above and one step to the right of b.

(c) If $k = 0$, move all cars one step up and one step right. Place c in the bottom-left corner.

Let the set of all these new parking functions be denoted by $Insert(\pi, c, k)$.

For any parking function π and any car c in the highest non-empty diagonal of π, let $\rho_c(\pi)$ be the parking function obtained by deleting c from π and shifting all cars in higher rows one step down and one step left. Then, we have the following:

Proposition 5.1. *Let π be any nonempty parking function. There is a unique choice of π', c, and*

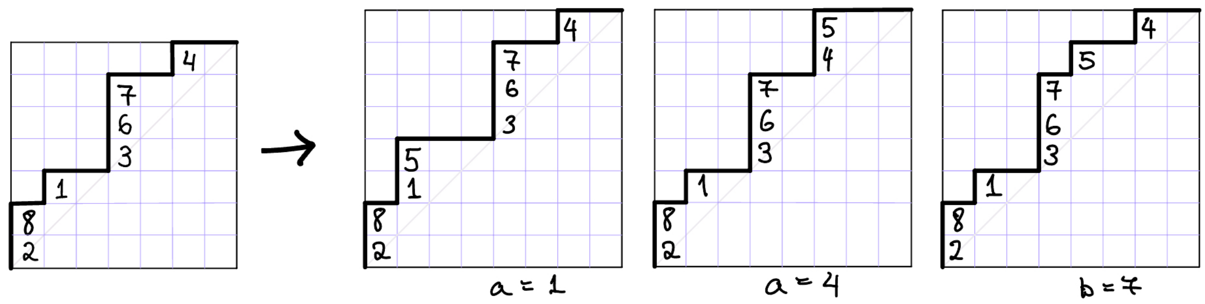

Figure 5.1: A parking function π together with the three parking functions in Insert$(\pi, 5, 2)$.

k such that $\pi \in Insert(\pi', c, k)$. In particular, k is the number of the highest non-empty diagonal of π, c is the smallest car in the k-diagonal, and $\pi' = \rho_c(\pi)$.

Thus one may build all parking functions by repeated insertions. One motivating factor behind this is that the area and *dinv* statistics change predictably under a single insertion. For fixed π, c, k, every $\pi' \in \text{Insert}(\pi, c, k)$ satisfies $\text{area}(\pi') = \text{area}(\pi) + k$. Moreover, the only new diagonal inversions involve the inserted car c; as one moves the insertion site for c from the rightmost admissible position leftwards across successive admissible sites, *dinv* increases by exactly 1 at each step. Consequently, we have

Lemma 5.1. *For any parking function π , car c, and diagonal k for which the algorithm 5.1 applies,*

$$\sum_{\pi' \in Insert(\pi,c,k)} t^{area(\pi')} q^{dinv(\pi')} = t^{area(\pi)+k} q^{dinv(\pi)} [|Insert(\pi, c, k)|]_q,$$

where given an integer m, $[m]_q$ denotes its q-analog:

$$[m]_q = \frac{1 - q^m}{1 - q} = 1 + q + \cdots + q^{m-1}.$$

To connect with schedules, we will fix a permutation $\sigma \in S_n$ and build a rooted tree of parking functions by inserting cars one by one. Start with the 1-car parking function containing just σ_n as the root at level 0. At level i, let $c = \sigma_{n-i}$ and k such that c lies in the $(k+1)$st from last run of σ. Then, for each parking function π from the previous level, we add as children the elements of $\text{Insert}(\pi, c, k)$. We denote the set of parking functions in PF_n generated from σ in this way by $PF(\sigma)$.

Example 5.2.1. If $\sigma = 3124$, the runs of σ are $r_1 = 3$ and $r_2 = 124$. Hence, $k = 0$ in the first three levels, and $k = 1$ in the last one. Repeatedly applying our algorithm, we obtain the tree shown in Figure 5.2.

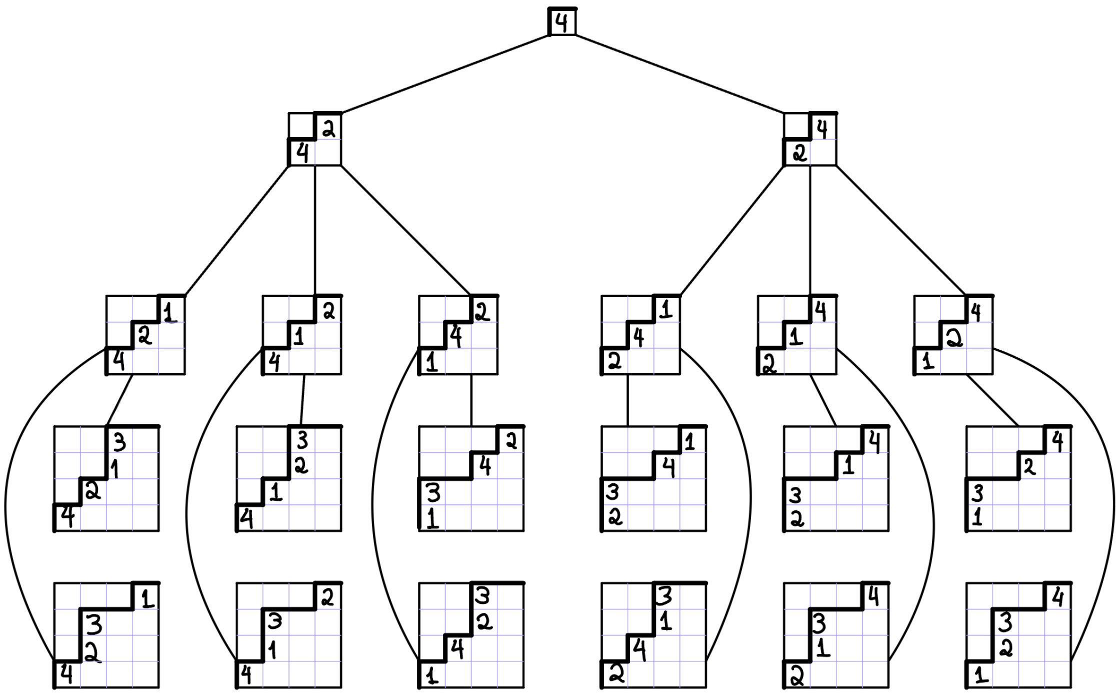

Figure 5.2: The tree of parking functions built from $\sigma = 3124$

Observe that in the above example, we have $sch(\sigma) = (2, 3, 2, 1)$. In particular, $sch_i(\sigma)$ is the number of children each parking function on level $n - 1 - i$ has. By construction, this is just $|\text{Insert}(\pi, c, k)|$. This turns out to be true in general.

Lemma 5.2. *Let $\sigma = \sigma_1\sigma_2 \cdots \sigma_n$ have maximal increasing runs $r_1, \ldots, r_m$, and set $r_{m+1} = \{0\}$. When inserting $c = \sigma_i$ into the diagonal $k = m - j$ according to the above procedure,the number of admissible insertion sites is exactly $sch_i(\sigma)$.*

Proof. By Algorithm 5.1, admissible sites for c in diagonal $k \neq 0$ are in bijection with

(a) $a \in r_{j+1}$ with $a < \sigma_i$ (place c above a),

(b) $b \in r_j$ with $b > \sigma_i$ (place c above–right of a)

These two sets are disjoint and exhaustive, hence the count equals

$$\#\{b \in r_j : b > \sigma_i\} + \#\{a \in r_{j+1} : a < \sigma_i\} = sch_i(\sigma)$$

If $j = m$, then $r_{m+1} = \{0\}$ contributes one insertion site at the bottom-left corner. □

Now using Lemma 5.2 and repeated applications of Lemma 5.1 we obtain:

Theorem 5.1. *For any permutation $\sigma \in S_n$, we have*

$$\sum_{\pi \in PF(\sigma)} q^{dinv(\pi)} t^{area(\pi)} = t^{maj(\sigma)} \prod_{i=1}^{n} [sch_i(\sigma)]_q.$$

Now since for every parking function $\pi \in PF_n$ there is exactly one $\sigma \in S_n$ such that $\pi \in PF(\sigma)$, Theorem 5.1 together with Corollary 4.3.1 give us another formula for the Hilbert series of the diagonal coinvariant ring:

Theorem 5.2 (The Schedules Formula)**.**

$$\mathcal{H}(DR_n; q, t) = \sum_{\sigma \in S_n} t^{maj(\sigma)} \prod_{i=1}^{n} [sch_i(\sigma)]_q.$$

This formula was first conjectured by Haglund and Loehr (HL05) and proven in 2018 by Carlson and Mellit as a consequence of their proof of the Shuffle Conjecture.

Example 5.2.2. We use the schedules formula to compute $\mathcal{H}(DR_3; q, t)$.
Summing the contributions for each $\sigma \in S_3$, we get

$$\mathcal{H}(DR_3; q, t) = 1 + 2q + 2q^2 + q^3 + 2t + 3tq + tq^2 + 2t^2 + t^2q + t^3.$$

σ	runs of σ	$sch(\sigma)$	$t^{maj(\sigma)} \prod_i [sch_i(\sigma)]_q$
123	123	$(3, 2, 1)$	$1 + 2q + 2q^2 + q^3$
132	$13 \mid 2$	$(1, 1, 1)$	t^2
213	$2 \mid 13$	$(1, 2, 1)$	$t(1 + q)$
231	$23 \mid 1$	$(2, 1, 1)$	$t^2(1 + q)$
312	$3 \mid 12$	$(2, 2, 1)$	$t(1 + 2q + q^2)$
321	$3 \mid 2 \mid 1$	$(1, 1, 1)$	t^3

Table 5.1: Schedules–formula contributions for $n = 3$

As a sanity check, we can set $q = t = 1$ to verify that $\mathcal{H}(DR_3; 1, 1) = 16 = (3 + 1)^{3-1}$, the number of parking functions on 3 cars, as predicted by the Shuffle Theorem.

CHAPTER 6

THE DELTA CONJECTURE

Shortly before the proof of the Shuffle Theorem, Haglund, Remmel and Wilson introduced a substantial generalization of the first two sides of the Shuffle Theorem, known as the Delta Conjecture. On the symmetric functions side, the operator ∇ is replaced by Δ, while the combinatorial side has two distinct formulations, the "rise version" and the "valley version". The rise version been proven independently by two different groups of researchers, while the valley version remains open. In this chapter, "Delta Conjecture" will refer to the valley version, in which parking functions are generalize by "valley-marked" parking functions.

6.1. The Delta Operator

Let $\widetilde{H}_\mu[X; q, t]$ denote the modified Macdonald polynomials, which form a basis for the ring of symmetric functions Λ. We define the Garsia-Haiman Delta operators Δ_f and Δ'_f acting on this basis. For any symmetric function $f \in \Lambda$, the action is defined by:

$$\Delta_f \widetilde{H}_\mu = f[B_\mu(q,t)]\widetilde{H}_\mu$$

and

$$\Delta'_f \widetilde{H}_\mu = f[B_\mu(q,t) - 1]\widetilde{H}_\mu$$

where $B_\mu(q,t) = \sum_{(i,j)\in\mu} q^{a'(i,j)} t^{l'(i,j)}$ is the sum of q,t-monomials over the cells of the Ferrers diagram of μ.

These operators are natural generalizations of ∇. In particular, ∇ acts on the Macdonald basis by $\nabla \widetilde{H}_\mu = e_n[B_\mu(q,t) - 1]\widetilde{H}_\mu$. Thus, we have the identity:

$$\nabla = \Delta'_{e_{n-1}}.$$

The Delta Conjecture, formulated by Haglund, Remmel, and Wilson, concerns the family of operators Δ'_{e_k} for $0 \leq k \leq n-1$. Specifically, it predicts a combinatorial expansion for

$$\Delta'_{e_{n-k-1}} e_n.$$

When $k = 0$, we recover $\Delta'_{e_{n-1}} e_n = \nabla e_n$, which is the Shuffle Theorem. For $k > 0$, the operator "deletes" generic dinv contributions, which combinatorially corresponds to "marking" certain valleys in the Dyck paths to exclude them from the dinv statistic. This motivates the definition of our primary combinatorial object: the valley-marked parking function.

6.2. Valley-marked parking functions

Definition 6.1. Given a parking function represented as a Dyck path, we say that a valley occurs whenever an East step is followed by a North step and either (1) there is no car under the East step or (2) the car under the East step is smaller than the car adjacent to the North step. A valley-marked parking function is a parking function with markings on some of its valleys. We also say that the car adjacent to a marked valley's North step is marked. The set of marked parking functions on n cars with m marked valleys will be denoted by $MPF_{n,m}$.

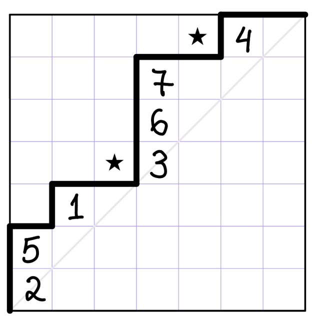

Figure 6.1: A valley-marked parking function. Here, both valleys are marked and dinv$= 5 - 2 = 3$.

Given a marked parking function $\pi^* \in MPF_{n,m}$, we define its area, word, and ides statistics to be those of the underlying (unmarked) parking function π. The dinv statistic is slightly different:

Definition 6.2. A pair of cars (a, b) with $a < b$ constitute a diagonal inversion if either i) a and b are on the same diagonal, b is further to the right, and a is unmarked or ii) b is one diagonal above and further to the left of a, and b is unmarked. Then, $\text{dinv}(\pi^*)$ is the total number of diagonal inversions minus the number of marked valleys.

Note that while the positions that could create a diagonal inversion are the same as for unmarked parking functions, we also require the car further to the left to be unmarked, so we could have

$$\operatorname{dinv}(\pi^*) < \operatorname{dinv}(\pi) - m.$$

Despite this, $\operatorname{dinv}(\pi^*)$ is always nonnegative. To see why this is true, for any marked car c, consider the closest unmarked car to the left of c in the same diagonal. Either it is smaller than c, creating a diagonal inversion of type i), or it is larger and must have some even larger, unmarked car on top of it, creating a diagonal inversion of type ii).

We now describe an analog to the insertion algorithm 5.1 for valley-marked parking functions.

Insertion Algorithm 6.1. *Let $\pi^* \in MPF_{n,m}$ a marked parking function, c a car not already present in π^*, and k an integer such that the k-diagonal of π^* contains no marked car smaller than c. Moreover, let $a < c$ be an unmarked car in the k-diagonal of π^* or $b > c$ any car in the $(k+1)$-diagonal. We obtain a marked parking function in $MPF_{n+1,m+1}$ as follows: starting from the a or b, look to the right for the first time that either (a) an unmarked car $b' > c$ appears in the $(k+1)$-diagonal or (b) the Dyck path returns to the line $y = x + k$.*

- *If (a) occurs before (b), move b' and all the cars in a higher row one step up and one step right. Place the marked car c directly below b'.*

- *If (b) occurs before (a), move all cars above that point one step up and one step right. Place the marked car c one step above and one step to the right of that point, so it lies in the k-diagonal.*

Let the set of all these new parking functions be denoted by $\operatorname{Insert}^*(\pi^*, c, k)$.

Similarly to the unmarked version, we obtain the following

Lemma 6.1. *For any marked parking function π^* , car c, and diagonal k for which Algorithm 6.1*

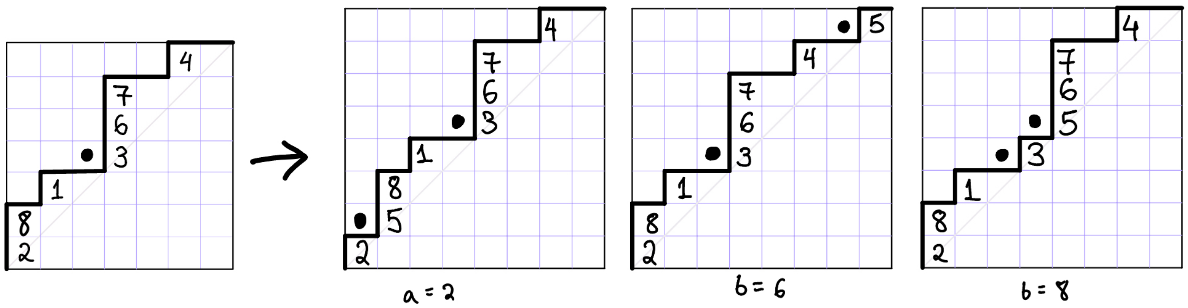

Figure 6.2: A valley marked parking function π^* together with the three parking functions in Insert$^*(\pi^*, 5, 0)$.

applies,

$$\sum_{\widetilde{\pi^*} \in Insert^*(\pi^*,c,k)} t^{area(\widetilde{\pi^*})} q^{dinv(\widetilde{\pi^*})} = t^{area(\pi^*)+k} q^{dinv(\pi^*)} [|Insert^*(\pi^*, c, k)|]_q.$$

We have to be a little more careful with stating the marked analogue to Proposition 5.1. Given $\pi^* \in MPF_{n,m}$ and a marked car c in the highest non-empty diagonal of π^*, we define $\rho_c(\pi^*) \in MPF_{n-1,m-1}$ to be the marked parking function obtained by deleting c (and its mark) from π^* and shifting all cars in higher rows one step down and one step left. Note that, unlike the unmarked case, there might be more than one choice of c and k for which $\pi^* \in$ Insert$^*(\rho_c(\pi^*), c, k)$. This is because we can insert marked cars from different diagonals in any order. To remedy this, we add one extra condition to our proposition:

Proposition 6.1. *Let π^* be any nonempty marked parking function. There is a unique choice of $\widetilde{\pi^*}, c,$ and k such that $\pi \in Insert(\widetilde{\pi^*}, c, k)$ and for which no diagonal above the k-diagonal contains any marked cars. In particular, k is the number of the highest diagonal of π^*, containing any marked cars, c is the largest marked car in the k-diagonal, and $\widetilde{\pi^*} = \rho_c(\pi^*)$.*

We will now build a tree of marked parking functions starting from a permutation σ^* with some of its elements marked. We will again insert the cars one by one, starting from σ_n^* to σ_1^*. We will first use Algorithm 5.1 to insert all unmarked cars and then Algorithm 6.1 for marked cars. Like before, if the car c lies in the $(k+1)$-st from last run of σ^*, it will be inserted into the k-diagonal.

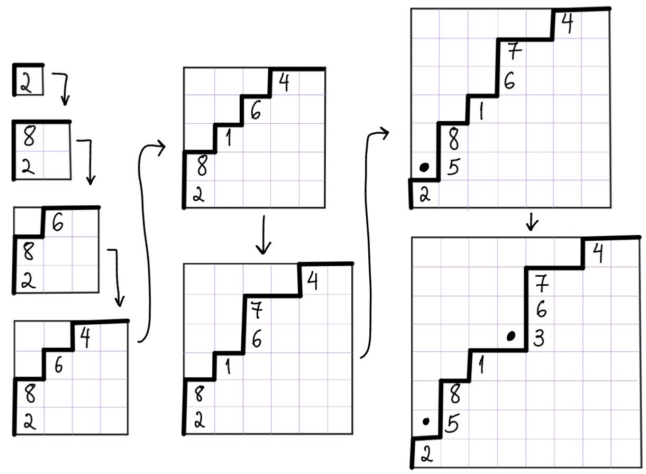

Figure 6.3: A path in the insertion tree of the marked permutation $\sigma^* = 714682\overline{35}$.

An example of a path in such a tree is shown in Figure 6.3.

To any $\pi^* \in MPF_{n,m}$ we can associate a marked permutation of size n so that π^* is in its insertion tree; this is obtained by the reading word $\sigma(\pi^*)$. By Proposition 6.1, π^* will not appear in any other permutation's insertion tree.

However, there are marked permutations of size n that do not have any marked parking functions on n cars in their insertion tree. We call the permutations that do *insertable.*

Example 6.2.1. $21\overline{3}$ is insertable; take two North steps with cars 1 and 2 next to them, followed by two East steps, a North step with the marked 3, and finally an East step . However, $1\overline{3}2$ is not insertable, since the marked 3 would have to be inserted in the 1-diagonal, but there is no way to do that following Algorithm 6.1.

It turns out that insertable permutations correspond to ordered set partitions.

Definition 6.3. An ordered set partition of a set S is a list of pairwise disjoint nonempty subsets of S such that their union is S . These subsets are called the blocks of the partition.

Example 6.2.2. The ordered set partitions of $S = \{1, 2, 3\}$ are

- $[\{1\},\{2\},\{3\}]$,
- $[\{1\},\{3\},\{2\}]$,
- $[\{2\},\{1\},\{3\}]$,
- $[\{3\},\{1\},\{2\}]$,
- $[\{2\},\{3\},\{1\}]$,
- $[\{3\},\{2\},\{1\}]$,
- $[\{1\},\{2,3\}]$,
- $[\{2\},\{1,3\}]$,
- $[\{3\},\{1,2\}]$,
- $[\{1,2\},\{3\}]$,
- $[\{1,3\},\{2\}]$,
- $[\{2,3\},\{1\}]$,
- $[\{1,2,3\}]$

For brevity, we will also represent a partition as a word with bars between the blocks, so for instance $23|1$ corresponds to $[\{2,3\},\{1\}]$.

Definition 6.4. Given an ordered set partition Π, we can inductively construct a permutation $\sigma(\Pi)$. If Π consists of only one block, we define $\sigma(\Pi)$ to be the elements of Π in increasing order. Otherwise, let Π' be the ordered set partition obtained by removing the first block Π_1 of Π and let r be the first element of $\sigma(\Pi')$. Then, $\sigma(\Pi)$ is formed by first writing the elements of Π_1 that are greater than r in increasing order, followed by the elements of Π_1 that are less than r in increasing order, followed by $\sigma(\Pi')$. We also define $\sigma^*(\Pi)$ by adding markings to every number that is not the left-most element, or "*leader*" of its block in Π.

Example 6.2.3. Let $\Pi = 9|57|1368|24$. We construct $\sigma(\Pi)$ one block at a time:

Π	$\sigma(\Pi)$
24	24
1368|24	368124
57|1368|24	57368124
9|57|1368|24	957368124

Table 6.1

Then, $\sigma^*(\Pi) = 95\overline{7}3\overline{681}2\overline{4}$.

Observe that we can recover Π from $\sigma^*(\Pi)$ by introducing a block break before each unmarked element, so there is at most one ordered set partition giving rise to each marked permutation.

Theorem 6.1. *A marked permutation is insertable if and only if it equals* $\sigma^*(\Pi)$ *for some ordered set partition* Π.

Proof. By the definition of the insertion algorithm, a marked permutation σ^* is insertable if and only if at every step of the insertion process, there is at least one valid diagonal k to insert the next car. In terms of schedules, this is equivalent to the condition that $sch_{\sigma^*}(c) \geq 1$ for every car c in σ^*. We prove the theorem by analyzing these schedule numbers.

First, suppose $\sigma^* = \sigma^*(\Pi)$ for some ordered set partition Π.

- If c is a marked car, it must be in the same block as some unmarked leader u to its left in $\sigma^*(\Pi)$. By the construction of $\sigma^*(\Pi)$, c is placed in the run immediately following the run containing u (or in the same run if $c > u$). The definition of the schedule for marked cars ensures that the presence of the unmarked leader u contributes at least 1 to $sch_{\sigma^*}(c)$, corresponding to the specific relative order enforced by the block structure.

- If c is an unmarked car, let $c = u_i$ be the leader of block B_i. We analyze its relationship with u_{i+1}, the leader of the next block B_{i+1} (or 0 if $i = k$). The recursive construction of $\sigma^*(\Pi)$ places B_i around the sequence generated by $B_{i+1} \dots B_k$, with the pivot element being $r = u_{i+1}$.

  - If $u_i < u_{i+1}$: The construction places u_i in the set of elements smaller than the pivot, appearing immediately before the sequence starting with u_{i+1}. Since $u_i < u_{i+1}$, they form an increasing subsequence and thus lie in the same run. Therefore, u_{i+1} contributes to the first term of $sch_{\sigma^*}(u_i)$ (unmarked cars in r_j greater than c).

  - If $u_i > u_{i+1}$: The construction places u_i in the set of elements larger than the pivot. This creates a descent between the run containing u_i and the run starting with u_{i+1}. Thus, u_{i+1} lies in the next run r_{j+1}. Since $u_{i+1} < u_i$, it contributes to the second term of $sch_{\sigma^*}(u_i)$ (unmarked cars in r_{j+1} smaller than c).

Finally, for the last leader u_k, the virtual element 0 acts as u_{k+1}, ensuring $sch_{\sigma^*}(u_k) \geq 1$.

In both cases, $sch_{\sigma^*}(c) \geq 1$, so $\sigma^*(\Pi)$ is insertable.

Conversely, suppose σ^* is an insertable marked permutation. We define Π by cutting σ^* before every unmarked element. We must show that $\sigma^* = \sigma^*(\Pi)$. If σ^* deviated from the canonical form $\sigma^*(\Pi)$ (for example, if a marked car c appeared in a run too far from its unmarked leader, or if the relative order of marked cars violated the block construction), the schedule definition would force $sch_{\sigma^*}(c) = 0$. Since σ^* is insertable, no such zero schedules exist. Therefore, σ^* must strictly adhere to the relative ordering prescribed by Π, implying $\sigma^* = \sigma^*(\Pi)$. □

Definition 6.5. Let π^* be a marked parking function. We say that the type of π^* is the marked permutation whose ith from last run contains the cars in the i-diagonal of π^* with markings above the marked cars of π^*.

Example 6.2.4. The marked parking function in Figure 6.1 is of type $71\overline{4}562\overline{3}$.

Given an ordered set partition Π of $\{1, \ldots, n\}$, let $MPF(\Pi)$ denote the set of all marked parking functions of type $\sigma^*(\Pi)$. So for instance, the parking function in Figure 6.1 is an element of $MPF(7|14|5|6|23)$. By Proposition 6.1, each marked parking function belongs to exactly one insertion tree. Therefore, by Theorem 6.1, for each marked parking function π^*, there is a unique ordered set partition Π such that $\pi^* \in MPF(\Pi)$. Thus,

$$MPF_{n,m} = \bigsqcup_{\Pi} MPF(\Pi)$$

where Π runs over all ordered set partitions of $\{1, \ldots, n\}$ withe exactly $n - m$ blocks.

The observations in this section lead us to seek a generalization of Theorem 5.1. With that goal in mind, we define schedules for marked parking functions.

6.3. Schedules for marked parking functions

Definition 6.6. Let Π be an ordered set partition of $\{1, \ldots, n\}$ and $\sigma^*(\Pi)$ its marked permutation. We define the runs $r_1, \ldots, r_k$ of $\sigma^*(\Pi)$ in the same way as in the unmarked case, with $\{0\}$ as the last run.

- If c is an unmarked car in the r_j, we define $sch_\Pi(c)$ as the number of unmarked cars in r_j that are greater than c together with the number of unmarked cars in r_{j+1} that are less than c.
- If c is a marked car in r_j, we define $sch_\Pi(c)$ as the number of unmarked cars in r_j that are smaller than c together with the number of unmarked cars in r_{j-1} that are larger than c.

Then, $sch_\Pi = (sch_\Pi(1), \ldots, sch_\Pi(n))$.

Example 6.3.1. If $\Pi = 9|57|1368|24$, as in the previous example, we have $\sigma^*(\Pi) = 95\overline{7}3\overline{681}2\overline{4}$. The runs of $\sigma^*(\Pi)$ are

$$r_1 = 9, r_2 = \overline{5}7, r_3 = 3\overline{68}, r_4 = \overline{1}2\overline{4}, r_5 = 0.$$

Hence, $sch_\Pi = (1, 1, 1, 1, 2, 1, 1, 1)$.

Similarly to the unmarked case, the schedule numbers $sch_\Pi(c)$ count how many cars can be used to insert c according to Algorithm 6.1 and thus correspond to the number of descendants in each level of the insertion tree built from $\sigma^*(\Pi)$. Once again, the appended 0 at the end accounts for the extra point of insertion in the special case $k = 0$ and does not affect the schedule numbers of marked cars. We now arrive at the following generalization of Theorem 5.1 (we obtain the latter as a special case where all the blocks of Π have size 1):

Theorem 6.2. *For any order set partition Π we have*

$$\sum_{\pi^* \in MPF(\Pi)} q^{dinv(\pi^*)} t^{area(\pi^*)} = t^{maj(\sigma(\Pi))} \prod_{c \in \Pi} [sch_\Pi(c)]_q.$$

Theorem 6.2 provides the q, t-enumerator for the set of marked parking functions associated with a fixed ordered set partition Π. If we sum this result over all ordered set partitions of $\{1, \ldots, n\}$ with $n - k$ blocks, we obtain the Hilbert series of the generalized module of diagonal coinvariants. However, the Delta Conjecture is a stronger statement: it predicts the full Frobenius series. To capture the action of the symmetric group, we must weight the sum by the quasisymmetric function $F_{ides(\pi^*)}$, just as we did in the Shuffle Theorem.

Conjecture 6.1 (The Delta Conjecture)**.** *For all* $k \leq n$,

$$\Delta'_{e_{n-k-1}} e_n = \sum_{\pi^* \in MPF_{n,k}} t^{area(\pi^*)} q^{dinv(\pi^*)} F_{ides(\pi^*)}.$$

The Delta Conjecture generalizes the Shuffle Theorem by replacing the standard diagonal coinvariant ring with a generalized module, and standard parking functions with valley-marked parking functions. Theorem 6.2 provides an explicit formula for the combinatorial side of this conjecture by extending the schedules method to the marked case. While the original Shuffle Theorem was proven in 2018, the Delta Conjecture remains open.

BIBLIOGRAPHY

[BHM+23] Jonah Blasiak, Mark Haiman, Jennifer Morse, Anna Pun, and George H. Seelinger. A proof of the extended delta conjecture. *Forum Math. Pi*, 11:Paper No. e6, 28, 2023.

[CM18] Erik Carlsson and Anton Mellit. A Proof of the Shuffle Conjecture, 2018.

[DM22] Michele D'Adderio and Anton Mellit. A proof of the compositional delta conjecture. *Adv. Math.*, 402:Paper No. 108342, 17, 2022.

[Ges84] Ira Gessel. Multipartite P-Partitions and Inner Products of Skew Schur Functions. *Contemporary Mathematics*, 34:289–301, 1984.

[Hag08] James Haglund. *The q,t-Catalan Numbers and the Space of Diagonal Harmonics*. American Mathematical Society, 2008.

[Hai02] Mark Haiman. Vanishing Theorems and Character Formulas for the Hilbert Scheme of Points in the Plane. *Inventiones Mathematicae*, 149(2):371–407, August 2002.

[HHL+04] J. Haglund, M. Haiman, N. Loehr, J. B. Remmel, and A. Ulyanov. A Combinatorial Formula for the Character of the Diagonal Coinvariants, 2004.

[HL05] J. Haglund and Nicholas Loehr. A Conjectured Combinatorial Formula for the Hilbert Series for Diagonal Harmonics. *Discrete Mathematics*, 298:189–204, 08 2005.

[HRW17] James Haglund, Jeffrey Remmel, and Andrew Timothy Wilson. The Delta Conjecture, 2017.

[HS20] James Haglund and Emily Sergel. Schedules and the Delta Conjecture, 2020.

[KW66] Alan G. Konheim and Benjamin Weiss. An Occupancy Discipline and Applications. *SIAM Journal on Applied Mathematics*, 14(6):1266–1274, 1966.

[Mac95] I.G. Macdonald. *Symmetric Functions and Hall Polynomials*. Oxford University Press, Oxford, UK, 2nd edition, 1995.