



HAL
open science

The Steep-Bounce zeta map in Parabolic Cataland

Cesar Ceballos, Wenjie Fang, Henri Mühle

► **To cite this version:**

Cesar Ceballos, Wenjie Fang, Henri Mühle. The Steep-Bounce zeta map in Parabolic Cataland. Journal of Combinatorial Theory, Series A, 2020, 172, pp.105210. 10.1016/j.jcta.2020.105210 . hal-02453173

HAL Id: hal-02453173

<https://hal.science/hal-02453173>

Submitted on 21 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

THE STEEP-BOUNCE ZETA MAP IN PARABOLIC CATALAND

CESAR CEBALLOS, WENJIE FANG, AND HENRI MÜHLE

ABSTRACT. As a classical object, the Tamari lattice has many generalizations, including ν -Tamari lattices and parabolic Tamari lattices. In this article, we unify these generalizations in a bijective fashion. We first prove that parabolic Tamari lattices are isomorphic to ν -Tamari lattices for bounce paths ν . We then introduce a new combinatorial object called “left-aligned colorable tree”, and show that it provides a bijective bridge between various parabolic Catalan objects and certain nested pairs of Dyck paths. As a consequence, we prove the Steep-Bounce Conjecture using a generalization of the famous zeta map in q, t -Catalan combinatorics. A generalization of the zeta map on parking functions, which arises in the theory of diagonal harmonics, is also obtained as a labeled version of our bijection.

CONTENTS

Introduction	1
Part 1. Parabolic Cataland	3
1.1. Members of Parabolic Cataland	4
1.1.1. Parabolic 231-Avoiding Permutations	4
1.1.2. Parabolic Noncrossing Partitions	4
1.1.3. Parabolic Dyck Paths	5
1.1.4. Left-Aligned Colorable Trees	6
1.2. Bijections in Parabolic Cataland	7
1.2.1. Noncrossing Partitions and Dyck Paths	7
1.2.2. 231-Avoiding Permutations and Noncrossing Partitions	8
1.2.3. 231-Avoiding Permutations and Left-Aligned Colorable Trees	9
1.2.4. Noncrossing Partitions and Left-Aligned Colorable Trees	12
1.2.5. Dyck Paths and Left-Aligned Colorable Trees	16
Part 2. Tamari Lattices in Parabolic Cataland	21
2.1. Parabolic Tamari Lattices	21
2.1.1. On the Duality of the Parabolic Tamari Lattices	22
2.2. ν_α-Tamari Lattices	23
2.2.1. ν_α -Bracket Vectors	24
2.2.2. The Galois Graph of \mathcal{T}_{ν_α}	25

Key words and phrases. parabolic Tamari lattice, ν -Tamari lattice, bijection, left-aligned colorable tree, zeta map.

Cesar Ceballos was supported by the Austrian Science Foundation FWF, grant F 5008-N15, in the framework of the Special Research Program Algorithmic and Enumerative Combinatorics.

Wenjie Fang was supported by the Austrian Science Foundation FWF, grant P27290 and I2309.

2.3. \mathcal{T}_α and \mathcal{T}_{ν_α} are Isomorphic	28
Part 3. The Zeta Map in Parabolic Cataland	31
3.1. Level-Marked Dyck Paths	31
3.2. The Steep-Bounce Theorem	35
3.3. A New Interpretation of the Zeta Map	37
3.3.1. The Zeta Map	37
3.3.2. Some Statistics on Dyck Paths	38
3.3.3. The Steep-Bounce Zeta Map	39
3.3.4. The Zeta Map on Parking Functions	42
Discussion	45
Acknowledgements	46
Appendix A. Lattice-Theoretic Background	46
A.1. Basics	46
A.2. Extremal Lattices	46
A.3. Congruence-Uniform Lattices	48
References	49

INTRODUCTION

The Tamari lattice can be realized as a partial order on various Catalan objects. It has been studied widely from various perspectives, leading to numerous generalizations. Two of its recent variants are the parabolic Tamari lattices of Williams and the third author [26] and the ν -Tamari lattices of Préville-Ratelle and Viennot [28]. The parabolic Tamari lattices are defined as certain lattice quotients of the weak order in parabolic quotients of the symmetric group, which generalizes a construction by Björner and Wachs [6]. The ν -Tamari lattices are, in contrast, partial orders defined by manipulating lattice paths. They generalize the m -Tamari lattices which were initially motivated by connections to trivariate diagonal harmonics [4], and now have remarkable applications to the theory of multivariate diagonal harmonics [5] and bijective links to other objects in combinatorics [12].

In this article, we reunite the objects in these two variants of the Tamari lattice. More precisely, we introduce a new family of parabolic Catalan objects which we call *left-aligned colorable trees* (or simply *LAC trees*). These trees connect the parabolic generalizations of 231-avoiding permutations, noncrossing set partitions, and Dyck paths introduced in [26], all of which can be recovered easily from the tree. We have the following result, with precise definitions delayed to subsequent sections.

Theorem I. *For every integer composition α , there is an explicit bijection from the set \mathbb{T}_α of α -trees to each of the following: the set $\mathfrak{S}_\alpha(231)$ of $(\alpha, 231)$ -avoiding permutations, the set \mathcal{NC}_α of noncrossing α -partitions, and the set \mathcal{D}_α of α -Dyck paths.*

We continue with a structural study of the two variants of the Tamari lattice. With each composition α one can associate a bounce path ν_α , and we show that

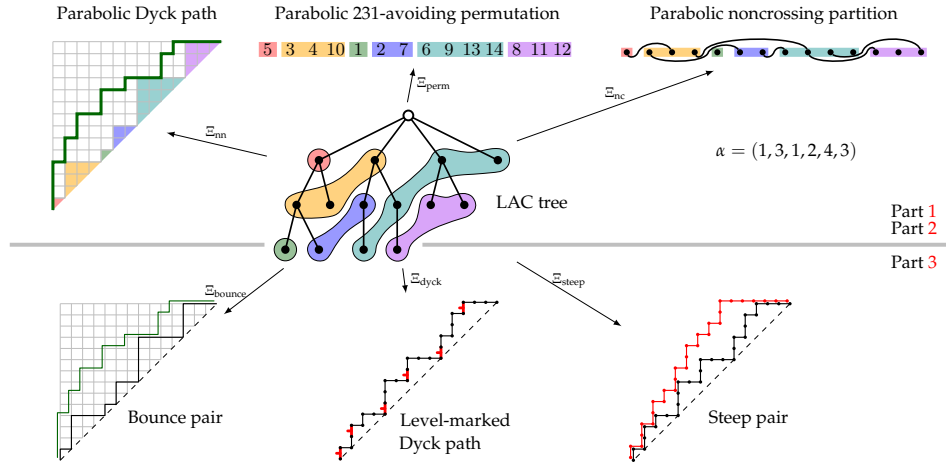


Figure 1. An overview of the bijections presented in this article.

the corresponding ν_α -Tamari lattice \mathcal{T}_{ν_α} and the parabolic Tamari lattice \mathcal{T}_α are isomorphic. This result is particularly interesting because these two lattices provide completely different perspectives. By definition of \mathcal{T}_{ν_α} , it is simple to check when two elements form a cover relation, but in general it is not easy to check whether they are comparable. In \mathcal{T}_α , the situation is quite the opposite: by definition it is easy to check when two elements are comparable, but to determine whether they form a cover relation is a nontrivial task.

We exploit the fact that \mathcal{T}_α and \mathcal{T}_{ν_α} are extremal lattices, and can therefore be uniquely represented by certain directed graphs. We show that the bijection between parabolic 231-avoiding permutations and parabolic Dyck paths from [26] extends to an isomorphism of these graphs, and we may therefore conclude the following result.

Theorem II. *For every integer composition α , the parabolic Tamari lattice \mathcal{T}_α is isomorphic to the ν_α -Tamari lattice \mathcal{T}_{ν_α} .*

In Part 3, we use LAC trees to prove the Steep-Bounce Conjecture of Bergeron, Ceballos, and Pilaud [5, Conjecture 2.2.8], which connects the graded dimensions of a certain Hopf algebra of pipe dreams to the enumeration of certain family of lattice walks in the positive quarter plane. Our proof is based on a bijection Γ between two families of nested Dyck paths via LAC trees.

Theorem III (The Steep-Bounce Theorem). *For $n > 0$ and every $r \in [n]$, the map Γ is a bijection from*

- *the set of nested pairs (μ_1, μ_2) of Dyck paths with $2n$ steps, where μ_2 is a steep path ending with exactly r east-steps with y -coordinate equal to n , to*
- *the set of nested pairs (μ'_1, μ'_2) of Dyck paths with $2n$ steps, where μ'_1 is a bounce path that touches the main diagonal $r + 1$ times.*

Interestingly, we show that the bijection Γ generalizes the famous zeta map in q, t -Catalan combinatorics. Given a Dyck path μ , we use two different greedy algorithms to define its associated *steep path* μ_{steep} and its *bounce path* μ_{bounce} . Details are delayed to related sections. With these notions, we have the following result.

Theorem IV. For $n > 0$, the map Γ restricts to a bijection from

- the set of pairs $(\mu, \mu_{\text{steep}})$, where μ is a Dyck path with $2n$ steps, to
- the set of pairs (v_{bounce}, v) , where v is a Dyck path with $2n$ steps.

Moreover, if $(v_{\text{bounce}}, v) = \Gamma(\mu, \mu_{\text{steep}})$, then $v = \zeta(\mu)$, where ζ is the zeta map from q, t -Catalan combinatorics.

The study of q, t -Catalan combinatorics originated over twenty years ago in connection to Macdonald polynomials and Garsia–Haiman’s theory of diagonal harmonics [13, 14, 16, 17, 20]. The q, t -Catalan polynomial can be obtained as the bi-graded Hilbert series of the alternating component of certain module of diagonal harmonics, whose dimension is equal to the Catalan number [20]. The zeta map is a bijection on Dyck paths that explains the equivalence of two combinatorial interpretations of these polynomials [17], which were discovered almost simultaneously by Haglund and Haiman in terms of two pairs of statistics on Dyck paths.

The zeta map was further generalized by Haglund and Loehr [19], who introduced an extension from the set of parking functions to certain diagonally labeled Dyck paths. As in the q, t -Catalan case, their extension explains the equivalence of two combinatorial interpretations of the Hilbert series of the full diagonal harmonics module [17, Chapter 5] [7], whose dimension is equal to the number of parking functions [20]. In Corollary 3.22, we show that this generalized zeta map can be obtained as a labeled version of our bijection Γ .

The organization of this article is as follows. In the first part, we introduce several families of parabolic Catalan objects, including left-aligned colorable trees, and then describe bijections between these families. In the second part, we dive into the order structure of parabolic Tamari lattices to explain how they are isomorphic to certain v -Tamari lattices. In the third part, we demonstrate how our bijection Γ settles the Steep-Bounce Conjecture while generalizing the zeta map in q, t -Catalan combinatorics, and present a labeled version which gives rise to the extension of the zeta map from parking functions to diagonally labeled Dyck paths.

Part 1. Parabolic Cataland

In the first part of this article, we present an overview of the currently explored parts of Parabolic Cataland. The name “Cataland” was coined in [36], and describes the interaction of the various tribes of Catalan families associated with finite Coxeter groups¹.

Recently, the frontiers of Cataland were pushed back in two different directions. Firstly, in [29], a positive integral parameter m was introduced, which yields Fuß-Catalan families associated with finite Coxeter groups. These families live in “Fuß-Cataland”. Secondly, in [26] (building on [36]), the various Catalan families were generalized to parabolic quotients of finite Coxeter groups. These generalizations constitute “Parabolic Cataland”, of which only the type A case is reasonably well understood. In this section, we explore these parabolic Catalan families and their web of connections.

¹By a “Catalan family” we mean a family of combinatorial objects enumerated by some Coxeter-Catalan numbers.

1.1. MEMBERS OF PARABOLIC CATALAND

We now introduce the main protagonists of this story, but we first fix some notation.

For a natural number $n > 0$, let $[n] \stackrel{\text{def}}{=} \{1, 2, \dots, n\}$, and let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$ be a composition of n . Moreover, let $\bar{\alpha} \stackrel{\text{def}}{=} (\alpha_r, \alpha_{r-1}, \dots, \alpha_1)$, and $|\alpha| = \alpha_1 + \alpha_2 + \dots + \alpha_r$.

For $i \in \{0, 1, \dots, r\}$, we define $s_i \stackrel{\text{def}}{=} \alpha_1 + \alpha_2 + \dots + \alpha_i$ and $t_i \stackrel{\text{def}}{=} n - s_{r-i}$, where we consider $s_0 = 0$. For $i \in [r]$ the set $\{s_{i-1} + 1, s_{i-1} + 2, \dots, s_i\}$ is the i -th α -region.

1.1.1. Parabolic 231-Avoiding Permutations. Let \mathfrak{S}_n be the symmetric group of degree n . The *parabolic quotient* of \mathfrak{S}_n with respect to α is defined by

$$\mathfrak{S}_\alpha \stackrel{\text{def}}{=} \{w \in \mathfrak{S}_{|\alpha|} \mid w(k) < w(k+1) \text{ for } k \notin \{s_1, s_2, \dots, s_{r-1}\}\}.$$

Definition 1.1 ([26, Definition 3.2]). In a permutation $w \in \mathfrak{S}_\alpha$, an $(\alpha, 231)$ -pattern is a triple of indices $i < j < k$ each in different α -regions such that $w(k) < w(i) < w(j)$ and $w(i) = w(k) + 1$. A permutation in \mathfrak{S}_α without $(\alpha, 231)$ -patterns is $(\alpha, 231)$ -avoiding.

A pair of indices (i, k) such that $w(i) = w(k) + 1$ is a *descent* of w . We denote the set of all $(\alpha, 231)$ -avoiding permutations of \mathfrak{S}_α by $\mathfrak{S}_\alpha(231)$. An example is shown in the top-left corner of Figure 1. In this article, a permutation w is always presented in its one-line notation $w(1)w(2) \dots w(n)$.

If $\alpha = (1, 1, \dots, 1)$, then we obtain precisely the stack-sortable permutations introduced in [21, Exercise 2.2.1.4].

1.1.2. Parabolic Noncrossing Partitions. For $|\alpha| = n$, an α -partition is a set partition of $[n]$ where every block intersects an α -region in at most one element. Let Π_α denote the set of all α -partitions. A *bump* of $\mathbf{P} \in \Pi_\alpha$ is a pair of consecutive elements in some block of \mathbf{P} , or more precisely, a pair (a, b) such that a and b belong to the same block of \mathbf{P} , and there exists no $c \in [n]$ with $a < c < b$ such that a and c belong to the same block of \mathbf{P} .

Definition 1.2 ([26, Definition 4.1]). An α -partition $\mathbf{P} \in \Pi_\alpha$ is *noncrossing* if it satisfies the following two conditions.

- (NC1): If two distinct bumps (a_1, b_1) and (a_2, b_2) of \mathbf{P} satisfy $a_1 < a_2 < b_1 < b_2$, then either a_1 and a_2 belong to the same α -region, or b_1 and a_2 belong to the same α -region.
- (NC2): If two distinct bumps (a_1, b_1) and (a_2, b_2) of \mathbf{P} satisfy $a_1 < a_2 < b_2 < b_1$, then a_1 and a_2 belong to different α -regions.

Let us denote the set of all noncrossing α -partitions by NC_α . We represent α -partitions graphically by its *diagram* as follows. First of all, we draw n vertices on a horizontal line, label them from 1 through n , and group them by α -regions. For $\mathbf{P} \in \Pi_\alpha$ we represent a bump (a, b) of \mathbf{P} by a curve that leaves the a -th vertex to the bottom, stays below all vertices in the same α -region as a , then moves up and continues above every vertex in the subsequent α -regions, until it reaches the b -th vertex, which it enters from the top. Throughout this article we say “elements of \mathbf{P} ” instead of “vertices of the diagram of \mathbf{P} ”.

We can verify that an α -partition is noncrossing if and only if it admits a diagram in which no two curves cross.

An example for $\alpha = (1, 3, 1, 2, 4, 3)$ is the noncrossing α -partition

$$\{\{1, 3\}, \{2, 6\}, \{4, 9, 12\}, \{5\}, \{7, 8\}, \{10, 14\}, \{11\}, \{13\}\},$$

whose diagram is shown in the top-central part of Figure 1.

If $\alpha = (1, 1, \dots, 1)$, then we obtain precisely the noncrossing set partitions studied in [22].

Let $\mathbf{P} \in \text{NC}_\alpha$. The unique block of \mathbf{P} containing 1 is the *first block* of \mathbf{P} , and we usually denote it by P_\circ .

If i is an element of \mathbf{P} , and (a, b) a bump of \mathbf{P} with $a < i \leq b$ such that a and i belong to different α -regions, then we say that i *lies below* (a, b) . Moreover, another bump (a', b') of \mathbf{P} *separates* i from (a, b) if i lies below (a', b') and either $a' < a$ with a, a' in the same α -region, or $a < a'$ for a, a' in different α -regions. In the diagram, it means that the bump (a', b') is between i and (a, b) . Finally, i *lies directly below* (a, b) if either $b = i$ or i is not separated from (a, b) by any other bump of \mathbf{P} .

For $P, P' \in \mathbf{P}$ we say that P' *starts below* P if $\min P'$ lies below some bump of P .

1.1.3. Parabolic Dyck Paths. A *Dyck path* is a lattice path in \mathbb{N}^2 starting from the origin, composed of steps $E \stackrel{\text{def}}{=} (1, 0)$ (so-called *east steps*) and $N \stackrel{\text{def}}{=} (0, 1)$ (so-called *north steps*), ending on the main diagonal while staying always weakly above it. A Dyck path can be regarded as a word in the alphabet $\{N, E\}$. For a Dyck path ν , we derive the path $\bar{\nu}$ by reversing ν and exchanging east and north steps.

A *valley* on a Dyck path is a lattice point \vec{p} that is preceded by an east-step and followed by a north-step. Similarly, a *peak* is a lattice point \vec{p} that is preceded by a north-step and followed by an east-step.

A Dyck path is *steep* if it does not have consecutive east-steps before the last north step. A Dyck path is *bounce* if it is of the form $N^{i_1}E^{i_1}N^{i_2}E^{i_2}\dots N^{i_s}E^{i_s}$, for some strictly positive integers i_1, i_2, \dots, i_s (because it ‘‘bounces off’’ the main diagonal $s - 1$ times).

The α -*bounce path* is $\nu_\alpha \stackrel{\text{def}}{=} N^{\alpha_1}E^{\alpha_1}N^{\alpha_2}E^{\alpha_2}\dots N^{\alpha_r}E^{\alpha_r}$. In particular, we have $\bar{\nu}_\alpha = \nu_{\bar{\alpha}}$.

Definition 1.3. Any Dyck path with $2|\alpha|$ steps that stays weakly above ν_α is an α -*Dyck path*.

We denote the set of all α -Dyck paths by \mathcal{D}_α . An example is shown on the left of Figure 1.

If $\alpha = (1, 1, \dots, 1)$, then we obtain the well-studied Dyck paths. See for instance [11] for an enumerative study of these paths.

1.1.4. Left-Aligned Colorable Trees. In this section we introduce a new family of combinatorial objects that are naturally parametrized by an integer n and some composition α of n . In Section 1.2 we show that these objects are in bijection with each of the previously mentioned parabolic Catalan objects.

Let T be a plane rooted tree with vertex set V , whose root is denoted by v_\circ . Throughout this article we refer to the vertices of a tree as *nodes*.

A *partial coloring* of T is a map $C: V' \rightarrow \mathbb{N}$ for some $V' \subseteq V \setminus \{v_\circ\}$. If $V' = V \setminus \{v_\circ\}$, then C is a *full coloring*. A *colored tree* is a pair (T, C) , where T is a plane rooted tree, and C is a full coloring of T . We refer to the elements in the image of C as C -colors.

A node $v \in V \setminus \{v_\circ\}$ is *active* with respect to a partial coloring C on $V' \subseteq V$ if $v \notin V'$ and its parent belongs to $V' \cup \{v_\circ\}$.

Recall that the *left-to-right traversal* of a plane rooted tree T is a depth-first search in T starting from the root, where children of the same node are visited *from left to right*. The *LR-prefix order* (resp. *LR-postfix order*) of T is the linearization of the nodes of T , where we record the nodes in order of first (resp. last) visits in left-to-right traversal.

Construction 1.4. Let $n > 0$ and let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$ be a composition of n . Let T be a plane rooted tree with n non-root nodes.

We initialize our algorithm by $\mathcal{T}_0 = (T, C_0)$, with C_0 the empty coloring. In the k -th step, we denote by A_k the set of active nodes of T with respect to the partial coloring C_{k-1} . If $|A_k| < \alpha_k$, then the algorithm fails. Otherwise, we let F_k be the set of the first α_k elements of A_k in the LR-prefix order of T , and we set $V_k = V_{k-1} \uplus F_k$. (Here \uplus denotes disjoint set union.) We define a partial coloring $C_k: V_k \rightarrow \mathbb{N}$ by

$$C_k(v) = \begin{cases} C_{k-1}(v), & \text{if } v \in V_{k-1}, \\ k, & \text{if } v \in F_k, \end{cases}$$

and we set $\mathcal{T}_k = (T, C_k)$.

If the algorithm has not failed after r steps, then we return the colored tree $\mathcal{T}_r = (T, C_r)$.

Every plane rooted tree for which Construction 1.4 does not fail is *compatible* with α ; the corresponding coloring is a *left-aligned coloring*. It is clear that every plane rooted tree has at most one left-aligned coloring with respect to α .

Definition 1.5. Any plane rooted tree with $|\alpha|$ non-root nodes that is compatible with α is an *α -tree*.

We denote the set of all α -trees by \mathbb{T}_α . An example is shown in the middle of Figure 1. We remark that every plane rooted tree is compatible with $\alpha = (1, 1, \dots, 1)$. Moreover, if we fix a plane rooted tree T , and consider $\alpha_T = (\alpha_1, \alpha_2, \dots)$, where α_i equals the number of nodes of T whose distance to the root is exactly i , then T is compatible with α_T . (In every step of Construction 1.4 we simply color all active nodes.) This composition will play an important role in the last part of this article.

We say that two nodes in a rooted tree are *comparable* if one is an ancestor of the other. Otherwise, the two nodes are *incomparable*.

Lemma 1.6. *Every α -tree has the following properties.*

- (i) *The color of every node is smaller than the color of its children, and nodes of the same color are incomparable.*
- (ii) *The active nodes of each step of the construction will eventually receive colors that are weakly increasing in LR-prefix order.*

Proof. The first point follows from the fact that a node becomes active before its children do. The second point follows from the construction of coloring where active nodes are selected for coloring in LR-prefix order, starting always from the first active node. \square

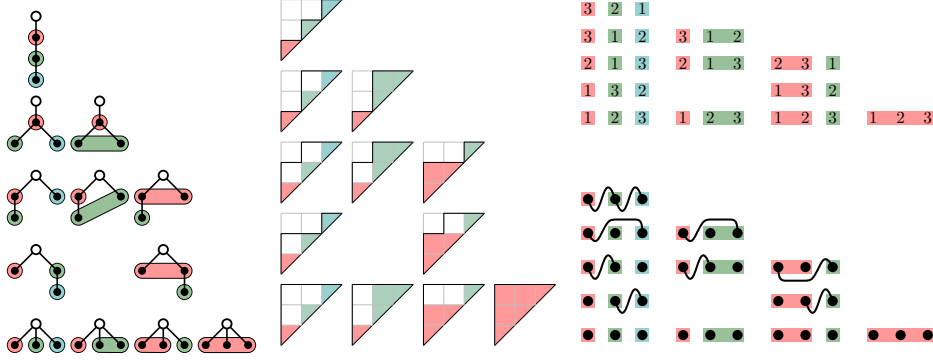


Figure 2. Examples of α -trees, α -Dyck paths, $(\alpha, 231)$ -avoiding permutations and noncrossing α -partitions of size $n = 3$. Different columns correspond to different α 's.

1.2. BIJECTIONS IN PARABOLIC CATALAND

In this section we describe explicit bijections between the parabolic Catalan objects introduced in Section 1.1. The bijections in the two following subsections 1.2.1 and 1.2.2 are reproduced from [26], the bijections of the remaining subsections 1.2.3, 1.2.4 and 1.2.5 are new. Figure 2 shows the various parabolic Catalan families for the compositions of 3.

1.2.1. Noncrossing Partitions and Dyck Paths. The following construction associates a noncrossing α -partition with each $\bar{\alpha}$ -Dyck path.

Construction 1.7. Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$ be an integer composition. Let $\mu \in \mathcal{D}_{\bar{\alpha}}$, whose valleys are $(p_1, q_1), (p_2, q_2), \dots, (p_m, q_m)$ with $p_1 < p_2 < \dots < p_m$. We set $p_0 = 0$.

We initialize our algorithm with the α -partition \mathbf{P}_0 without any bumps. In the i -th step, we construct the α -partition \mathbf{P}_i from \mathbf{P}_{i-1} by adding a bump corresponding to the valley (p_i, q_i) . To do so, we pick the unique index $0 \leq k < r$ such that $t_k \leq q_i \leq t_{k+1}$, so that we can write $q_i = t_k + \ell$ for some $\ell \in \{0, 1, \dots, t_{k+1} - t_k\}$. We then obtain \mathbf{P}_i by adding the bump (a, b) to \mathbf{P}_{i-1} , where a is the $(\ell + 1)$ -st element in the $(r - k)$ -th α -region and b is the $(p_i - p_{i-1})$ -th element after the $(r - k)$ -th α -region that is not already below some bump.

After m steps we return the α -partition $\mathbf{P}_\mu = \mathbf{P}_m$.

Two examples of this construction are illustrated in the top part of Figure 11.

Theorem 1.8 ([26, Theorem 5.2]). *For $\mu \in \mathcal{D}_{\bar{\alpha}}$ the α -partition \mathbf{P}_μ is noncrossing. Moreover, the map*

$$(1) \quad \Theta_1: \mathcal{D}_{\bar{\alpha}} \rightarrow \text{NC}_\alpha, \quad \mu \mapsto \mathbf{P}_\mu.$$

is a bijection that sends valleys to bumps.

1.2.2. 231-Avoiding Permutations and Noncrossing Partitions. Recall that a binary relation $R \subseteq M \times M$, for M a finite set, is *acyclic* if there does not exist a sequence $(a_1, a_2), (a_2, a_3), \dots, (a_k, a_1) \in R$ for any $k > 0$. It is well known that

an acyclic relation can be extended to a partial order by taking the reflexive and transitive closure.

We now define a binary relation $\vec{R}_{\mathbf{P}}$ on \mathbf{P} by setting $(P, P') \in \vec{R}_{\mathbf{P}}$ if and only if P' starts below P .

Lemma 1.9. *For every $\mathbf{P} \in NC_{\alpha}$, the relation $\vec{R}_{\mathbf{P}}$ is acyclic.*


Proof. This follows from the fact that $(P, P') \in \vec{R}_{\mathbf{P}}$ implies $\min P < \min P'$. \square

As a consequence, we may define a partial order $\vec{O}_{\mathbf{P}}$ on the blocks of \mathbf{P} by taking the reflexive and transitive closure of $\vec{R}_{\mathbf{P}}$. By construction the first block of \mathbf{P} is a minimal element of the poset $(\mathbf{P}, \vec{O}_{\mathbf{P}})$.

Example 1.10. Let $\alpha = (1, 3, 1, 2, 4, 3, 1)$, and consider the noncrossing α -partition

$$\mathbf{P} = \left\{ \begin{array}{cccccccc} \{1, 3, 10\}, & \{2, 6\}, & \{4\}, & \{5\}, & \{7, 8\}, & \{9, 13\}, & \{11, 14\}, & \{12, 15\} \\ \parallel & \parallel & \parallel & \parallel & \parallel & \parallel & \parallel & \parallel \\ P_1 & P_2 & P_3 & P_4 & P_5 & P_6 & P_7 & P_8 \end{array} \right\}.$$


The diagram of \mathbf{P} is shown below, with its first block highlighted.



The block P_2 , for instance, starts below P_1 because $(1, 3)$ is a bump and $\min P_2 = 2$. The relation $\vec{R}_{\mathbf{P}}$ is the following:

$$\vec{R}_{\mathbf{P}} = \{(P_1, P_2), (P_1, P_4), (P_1, P_5), (P_1, P_6), (P_2, P_4), (P_6, P_8), (P_7, P_8)\}.$$

The poset diagram of $(\mathbf{P}, \vec{O}_{\mathbf{P}})$ is shown below.



Now let $|\alpha| = n$, and $\mathbf{P} \in NC_{\alpha}$ whose first block is P_{\circ} . Further, let $\mathbf{P}' = \mathbf{P} \setminus P_{\circ}$. We may view \mathbf{P}' as a noncrossing α' -partition of $[n] \setminus P_{\circ}$ for some appropriate composition α' . By construction, we obtain $\vec{O}_{\mathbf{P}'}$ by removing every pair of the form (P_{\circ}, \cdot) from $\vec{O}_{\mathbf{P}}$.

Construction 1.11. Let α be an integer composition, and let $|\alpha| = n$. Let $\mathbf{P} \in NC_{\alpha}$ whose first block is $P_{\circ} = \{i_1, i_2, \dots, i_s\}$ with $i_1 < i_2 < \dots < i_s$. (Thus, $i_1 = 1$.)

We construct a permutation $w_{\mathbf{P}}$ recursively such that

$$w_{\mathbf{P}}(i_1) = w_{\mathbf{P}}(i_2) + 1 = \dots = w_{\mathbf{P}}(i_s) + s - 1,$$

and $w_{\mathbf{P}}(i_1)$ is as small as possible. To achieve that, we determine

$$D = \bigcup \{P \in \mathbf{P} \mid (P_{\circ}, P) \in \vec{O}_{\mathbf{P}}\},$$

and we set $w_{\mathbf{P}}(i_1) = |D|$. (In other words, D consists of all elements of $[n]$ that belong to a block of \mathbf{P} that lies weakly above P_{\circ} in $(\mathbf{P}, \vec{O}_{\mathbf{P}})$.)


Now we break \mathbf{P} into two pieces, namely

$$\begin{aligned}\mathbf{P}_1 &= \{P \cap D \mid P \in \mathbf{P} \setminus P_o\} \setminus \{\emptyset\}, \\ \mathbf{P}_2 &= \{P \cap ([n] \setminus D) \mid P \in \mathbf{P} \setminus P_o\} \setminus \{\emptyset\}.\end{aligned}$$

We may view \mathbf{P}_1 and \mathbf{P}_2 as noncrossing partitions with respect to some integers n_1 and n_2 and compositions α_1 of n_1 and α_2 of n_2 , respectively. In particular, $n_1, n_2 < n$, so that we may determine the missing values of $w_{\mathbf{P}}$ via this recursion. We will have $w_{\mathbf{P}}(i) > |D|$ if $i \in [n] \setminus D$, and $w_{\mathbf{P}}(i) < |D| + s - 1$ if $i \in D \setminus P_o$.

The initial condition for this recursion assigns the identity permutation to the α -partition without bumps.

Example 1.12. Let us continue Example 1.10. The first block of \mathbf{P} is $P_o = \{1, 3, 10\}$, and we see from the poset diagram above that $D = \{1, 2, 3, 5, 6, 7, 8, 9, 10, 12, 13, 15\}$. Therefore we set $w_{\mathbf{P}}(1) = 12$, $w_{\mathbf{P}}(3) = 11$, and $w_{\mathbf{P}}(10) = 10$. The two smaller partitions are $\mathbf{P}_1 = \{\{2, 6\}, \{5\}, \{7, 8\}, \{9, 13\}, \{12, 15\}\}$ and $\mathbf{P}_2 = \{\{4\}, \{11, 14\}\}$ whose diagrams are shown below.



Recursively, we obtain two permutations $w_1 = w_{\mathbf{P}_1} = 3 \mid 1 \mid 2 \ 5 \mid 4 \ 9 \mid 7 \ 8 \mid 6$ and $w_2 = w_{\mathbf{P}_2} = 1 \mid 3 \mid 2$. We now embed these permutations into $w_{\mathbf{P}}$, where we have to increase the values of w_2 by $|D| = 12$. The resulting permutation is

$$w_{\mathbf{P}} = 12 \mid 3 \ 11 \ 13 \mid 1 \mid 2 \ 5 \mid 4 \ 9 \ 10 \ 15 \mid 7 \ 8 \ 14 \mid 6.$$

Two more examples of this construction are illustrated in the bottom part of Figure 11.

Theorem 1.13 ([26, Theorem 4.2]). *For $\mathbf{P} \in NC_\alpha$ the permutation $w_{\mathbf{P}}$ is $(\alpha, 231)$ -avoiding. Moreover, the map*

$$(2) \quad \Theta_2: NC_\alpha \rightarrow \mathfrak{S}_\alpha(231), \quad \mathbf{P} \mapsto w_{\mathbf{P}}$$

is a bijection that sends bumps to descents.

1.2.3. 231-Avoiding Permutations and Left-Aligned Colorable Trees. The following construction associates a permutation in \mathfrak{S}_α with an α -tree.

Construction 1.14. Let α be an integer composition. Let $T \in \mathbb{T}_\alpha$ whose left-aligned coloring produced by Construction 1.4 is C .

We label the non-root nodes of T in LR-postfix order. We then group together the labels that correspond to nodes with the same color under C , and order these labels increasingly. Finally, we order the blocks according to their color, and obtain a permutation $w_T \in \mathfrak{S}_\alpha$.

An example of Construction 1.14 is illustrated in Figure 3.

Lemma 1.15. *For a non-root node u in a plane rooted tree T , the next node v in the LR-postfix order is either its parent, of which u is the last child, or the left-most leaf of the subtree induced by its sibling immediately to the right.*

Proof. This follows from the definition of the LR-postfix order. \square




Figure 3. An illustration of Construction 1.14 for $\alpha = (1, 3, 1, 2, 4, 3)$.

Construction 1.14 suggests the definition of the following map, which we next show is well-defined:

$$(3) \quad \Xi_{\text{perm}}: \mathbb{T}_\alpha \rightarrow \mathfrak{S}_\alpha(231), \quad T \mapsto w_T.$$

Proposition 1.16. *For an α -tree T , the permutation w_T obtained by Construction 1.14 is $(\alpha, 231)$ -avoiding.*

Proof. Let T be an α -tree with corresponding left-aligned coloring C . Suppose that (i, j, k) is an $(\alpha, 231)$ -pattern in w_T . Let u_i, u_j, u_k denote the nodes of T labeled by i, j, k , respectively, in Construction 1.14. It follows that $C(u_i) < C(u_j) < C(u_k)$.

By Lemma 1.15, since $w_T(k) + 1 = w_T(i)$, the node u_i is either the parent of u_k , or the left-most leaf of the subtree induced by the sibling of u_k on the right. However, the second case cannot happen, since the node u_k would then become active before u_i (because u_k precedes u_i in LR-prefix order). By Lemma 1.6(ii), we would have $C(u_i) \geq C(u_k)$, which is a contradiction.

Therefore, u_i is the parent of u_k , which means that u_k is active during steps $i + 1, i + 2, \dots, k$ in Construction 1.4, in particular in step j . Since $C(u_j) = j < C(u_k)$, the node u_j must come before u_k in LR-prefix order. However, since $w(i) < w(j)$, we know that u_j comes after u_i in LR-postfix order, and $C(u_i) < C(u_j)$ means that u_j is not an ancestor of u_i . Therefore, u_j must come after u_i in LR-prefix order, and thus also after u_k . This is a contradiction, and we conclude that w_T is $(\alpha, 231)$ -avoiding. \square

Let us now describe how to obtain a plane rooted tree from an $(\alpha, 231)$ -avoiding permutation.

Construction 1.17. Let α be an integer composition, and let $w \in \mathfrak{S}_\alpha(231)$.

We initialize our algorithm with the tree T_0 that has a single node labeled by $|\alpha| + 1$. In the i -th step we create the tree T_i by inserting a node labeled by $w(i)$ into T_{i-1} . This insertion proceeds as follows. We start at the root and walk around T_{i-1} . Suppose that we have reached a node v with label a . If $w(i) < a$, then we move to the left-most child of v , otherwise we move to the first sibling of v on the right. If the destination does not exist, then we create it and label it by $w(i)$.

After $|\alpha|$ steps we return the plane rooted tree $T_w = T_{|\alpha|}$.

Readers familiar with binary search trees may notice that this construction is essentially the insertion algorithm for binary search trees composed with the classical bijection between binary and plane trees.

It turns out that the output of Construction 1.17 is a plane rooted tree compatible with α , which enables us to define the following map:

$$(4) \quad \Lambda_{\text{perm}} : \mathfrak{S}_\alpha(231) \rightarrow \mathbb{T}_\alpha, \quad w \mapsto T_w.$$

Proposition 1.18. *For an $(\alpha, 231)$ -avoiding permutation w , the plane rooted tree T_w obtained by Construction 1.17 is compatible with α .*

Proof. Let us define a full coloring f of T_w as follows. For every node v of T_w with label a we set $f(v) = i$ if and only if $w^{-1}(a)$ belongs to the i -th α -region. It remains to prove that f is the left-aligned coloring of T_w produced in Construction 1.4. (Because in that case, Construction 1.4 does not fail, and T_w is indeed an α -tree.)

Construction 1.17 implies that the nodes of T_w are labeled in LR-postfix order. Now we show that for every non-root node u of T_w whose parent v is not the root, we have $f(u) > f(v)$. Since v is the parent of u , it is inserted before u , which yields $f(u) \geq f(v)$. If $f(u) = f(v)$, then the labels of u and v must come from indices in the same α -region. However, by construction, the label of u is smaller than that of v , which would produce an inversion in their common α -region, contradicting the assumption that $w \in \mathfrak{S}_\alpha$. We conclude that nodes of T_w labeled by elements coming from the same α -region are incomparable.

Now let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$. We now prove that T_w is an α -tree, and f is its left-aligned coloring. We consider the partial colorings C_s obtained after step s of Construction 1.4, and proceed by induction to prove that, for all $0 \leq s \leq r$, the algorithm does not fail in step s , and C_s is the same as the restriction f_s of f to values from 1 to s . The base case $s = 0$ is clearly correct. Now suppose that the induction hypothesis holds for $s - 1$.

We first show that there are always enough active nodes for step s . Since $w \in \mathfrak{S}_\alpha(231)$, labels $w(i)$ for i in the s -th α -region are inserted in increasing order, and by Construction 1.17, these newly inserted nodes cannot be comparable. Therefore, they are all children of nodes in the domain of C_{s-1} , thus active at step s . We therefore have at least α_s active nodes, and C_s is well defined.

We now show that C_s coincides with f_s . Suppose otherwise, and let u_* be the first node of T_w in LR-prefix order that is in the domain of C_s but not in that of f_s . We have $f(u_*) > s$. Let u_i be the parent of u_* , which is in the domain of C_{s-1} , and therefore in that of f_{s-1} by induction hypothesis. We thus have $f(u_i) \leq s - 1$.

By minimality of u_* , the node u_i is in the domain of C_s . Let u_k be the rightmost child of u_i (which exists as u_i has a least one child u_*). Since u_k is inserted after u_* , we conclude that $f(u_k) \geq f(u_*)$. It is clear that the domains of C_s and of f_s are of the same size. Hence, there is a node u_j in the domain of f_s but not in that of C_s , and we take it the minimal in LR-prefix order.

By induction hypothesis, we have $f(u_j) \geq s$, thus $f(u_j) = s$. Since nodes are colored in LR-prefix order in Construction 1.17, u_* is also the first node that differs in the domains of C_s and f_s , meaning that u_* precedes u_j in LR-prefix order. If u_j is a child of u_i , since $f(u_*) > s = f(u_j)$, it would lead to a contradiction to Construction 1.17. Therefore, u_j comes after u_i in LR-postfix order.

Let i, j, k be the indices such that u_i, u_j, u_k are labeled by $w(i), w(j), w(k)$ respectively. Since $f(u_i) < f(u_j) = s < f(u_k)$, we have $i < j < k$. Moreover, by construction, since u_k is the rightmost child of u_i , we have $w(k) + 1 = w(i)$, and

since u_j comes after u_i in LR-postfix order, we have $w(j) > w(i)$. Therefore, (i, j, k) is an $(\alpha, 231)$ -pattern in w , which is a contradiction.

We thus conclude the induction, meaning that f is the left-aligned coloring of T_w , which is therefore an α -tree. \square

We conclude this section with the proof that Ξ_{perm} is a bijection with inverse Λ_{perm} . In a plane rooted tree, an *internal* node is a node that is neither the root nor a leaf.

Theorem 1.19. *For every integer composition α , the map Ξ_{perm} is a bijection whose inverse is Λ_{perm} . Moreover, Ξ_{perm} sends internal nodes to descents.*

Proof. Given Propositions 1.16 and 1.18, it remains to prove that $\Xi_{\text{perm}} \circ \Lambda_{\text{perm}} = \text{id}$ and $\Lambda_{\text{perm}} \circ \Xi_{\text{perm}} = \text{id}$. Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$.

To show that $\Xi_{\text{perm}} \circ \Lambda_{\text{perm}} = \text{id}$, for $w \in \mathfrak{S}_\alpha(231)$, we consider $T = \Lambda_{\text{perm}}(w)$ and $w' = \Xi_{\text{perm}}(T)$. Let C be the full coloring of T constructed in the proof of Proposition 1.16. We have seen in that proof that C is precisely the left-aligned coloring of T from Construction 1.4. Therefore, entries in each α -region of w and w' agree, and since there is a unique way to order a set of distinct integers increasingly, we conclude that $w = w'$.

Now for $\Lambda_{\text{perm}} \circ \Xi_{\text{perm}} = \text{id}$, given $T \in \mathbb{T}_\alpha$, we consider $w = \Xi_{\text{perm}}(T)$ and $T' = \Lambda_{\text{perm}}(w)$. Let C, C' be the left-aligned colorings of T and T' , respectively, and for $s \in [r]$ let T_s and T'_s denote the induced subtrees consisting of the nodes of color at most s in T and T' respectively. We prove by induction on s that $T_s = T'_s$ holds for all $0 \leq s \leq r$.

The base case $s = 0$ is clear, as T_0 and T'_0 consist only of the root. Now suppose that $T_{s-1} = T'_{s-1}$. Let u be a node of T with $C(u) = s$ and label $w(i)$, and let u' be the node of T' with the same label. Since elements in any α -region are ordered increasingly, we know that every node of color s inserted in T' before u' has a smaller label than u' , and by construction no such node can be the parent of u' . It follows that the parent v' of u' belongs to T'_{s-1} . By our induction hypothesis, it follows that the parent v of u in T has the same label as v' . Now, for the order of newly inserted children of a node, by Construction 1.17, in T' the children of every node at each step are ordered by increasing labels, which is the same as in T due to the LR-postfix order. We thus have $T_s = T'_s$, which completes the induction step. We conclude that $T = T_r = T'_r = T'$. \square

1.2.4. Noncrossing Partitions and Left-Aligned Colorable Trees. The next construction associates an α -partition with every α -tree.

Construction 1.20. Let α be an integer composition. Let $T \in \mathbb{T}_\alpha$ whose left-aligned coloring produced by Construction 1.4 is C . We write the $|\alpha|$ non-root nodes of T on a horizontal line, where we group them by their color under C . Each group is then sorted by LR-prefix order, while groups are sorted by increasing color. We then connect two nodes u and v by a bump if and only if v is the rightmost child of u . The diagram thus obtained is the *flattening* of T .

By Lemma 1.6(i) there are no bumps between elements in the same color group, which implies that the flattening of T is the diagram of some α -partition \mathbf{P}_T .

See Figure 4 for an illustration of Construction 1.20. We now prove that every α -partition obtained by Construction 1.20 is in fact noncrossing, which allows us




Figure 4. An illustration of Construction 1.20 for $\alpha = (1, 3, 1, 2, 4, 3)$.

to define the following map:

$$(5) \quad \mathbb{E}_{\text{nc}} : \mathbb{T}_\alpha \rightarrow \text{NC}_\alpha, \quad T \mapsto \mathbf{P}_T.$$

Proposition 1.21. *For an α -tree T , the α -partition \mathbf{P}_T obtained by Construction 1.20 is noncrossing.*

Proof. We recall that by Lemma 1.6(i), nodes with the same color in T are incomparable. Thus, \mathbf{P}_T is indeed an α -partition. Now suppose that \mathbf{P}_T is not noncrossing. There are thus two bumps $(a_1, b_1), (a_2, b_2)$ that violate either (NC1) or (NC2). Let u_1, v_1, u_2, v_2 be the nodes of T corresponding to the elements a_1, b_1, a_2, b_2 , whose colors in the left-aligned coloring of T are c_1, d_1, c_2, d_2 .

If the two bumps violate (NC1), then $a_1 < a_2 < b_1 < b_2$, but a_1, a_2, b_1 belong to three different α -regions. It follows that $c_1 < c_2 < d_1$. We conclude that u_2 precedes v_1 in LR-prefix order, since v_1 becomes active immediately after step c_1 of Construction 1.4 in which its parent u_1 receives its color. Moreover, since $c_1 < c_2$, the node u_2 is not an ancestor of u_1 . Consequently, the right-most child of u_2 , which is v_2 , precedes v_1 in LR-prefix order, too. However, since $b_1 < b_2$, we have $d_1 \leq d_2$, which implies that v_1 precedes v_2 in LR-prefix order. This is a contradiction, and we conclude that (NC1) is not violated.

If the two bumps violate (NC2), then $a_1 < a_2 < b_2 < b_1$, and a_1, a_2 are in the same α -region. It follows that $c_1 = c_2$, and from Lemma 1.6(i) we know that u_1 and u_2 are incomparable in T . By Construction 1.20, the node u_1 thus precedes u_2 in LR-prefix order. Since v_1 is the rightmost child of u_1 , it also precedes u_2 in LR-prefix order. Since v_2 is the rightmost child of u_2 , it is also preceded by v_1 in the LR-prefix order. However, since $b_2 < b_1$, we have $d_2 \leq d_1$, which implies that v_2 precedes v_1 in LR-prefix order. This is a contradiction, and we conclude that (NC2) is not violated, either.

Hence, \mathbf{P}_T must be noncrossing. \square

Let us now describe how to obtain a plane rooted tree from a noncrossing α -partition.

Construction 1.22. For α a composition of n , let \mathbf{P} be a noncrossing α -partition, and we label the elements of \mathbf{P} by $1, 2, \dots, n$ from left to right.

We start with a collection of $n + 1$ nodes v_0, v_1, \dots, v_n . If the element i does not lie below any bump of \mathbf{P} , then we add an edge from v_0 to v_i , making v_i a child of v_0 . If i lies directly below a bump (a, b) , then we add an edge from v_a to v_i in the same way.

Since we associate with each element of \mathbf{P} a parent with strictly smaller index, the resulting graph does not have cycles. Moreover, it has $n + 1$ nodes and n edges,

and therefore is a tree. Finally, we order the children of each node by their label in increasing order, and root the tree at v_0 . We thus obtain a plane rooted tree $T_{\mathbf{P}}$.

We now prove that the plane rooted tree $T_{\mathbf{P}}$ obtained from Construction 1.22 is compatible with α so that we obtain the following map:

$$(6) \quad \Lambda_{\text{nc}}: \text{NC}_{\alpha} \rightarrow \mathbb{T}_{\alpha}, \quad \mathbf{P} \mapsto T_{\mathbf{P}}.$$

Proposition 1.23. *For a noncrossing α -partition \mathbf{P} , the plane rooted tree $T_{\mathbf{P}}$ is compatible with α .*

Proof. Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$ with $|\alpha| = n$, and suppose that the nodes of $T_{\mathbf{P}}$ are v_0, v_1, \dots, v_n as in Construction 1.22. We define a full coloring f of $T_{\mathbf{P}}$, where $f(v_a) = i$ if and only if a belongs to the i -th α -region in \mathbf{P} . We now prove that Construction 1.4 does not fail on $T_{\mathbf{P}}$ for α , and f agrees with the coloring of $T_{\mathbf{P}}$ thus obtained. We proceed by induction on the number s of steps.

The elements in the first α -region do not lie below any bump of \mathbf{P} , which means that Construction 1.4 does not fail in the first step. Moreover, by Construction 1.22 the nodes $v_1, v_2, \dots, v_{\alpha_1}$ are the leftmost children of the root of $T_{\mathbf{P}}$, and the partial labeling of Construction 1.4 therefore assigns the color 1 to all these nodes. Hence, this partial labeling agrees with f on the nodes colored by 1. This establishes the base case $s = 1$ of our induction.

Now suppose that the induction hypothesis holds of $s - 1 < r$, and we now consider step s . For $a \in \{\alpha_{s-1} + 1, \alpha_{s-1} + 2, \dots, \alpha_s\}$, let w be the parent of v_a . By Construction 1.22, w is either the root, or $w = v_{a'}$ for some $a' \leq \alpha_{s-1}$. This implies $f(w) \leq s - 1$. Thus, v_a is active in step s of Construction 1.4. Since a was chosen arbitrarily, we conclude that all vertices of $T_{\mathbf{P}}$ that come from the s -th α -region are active in step s of Construction 1.4. Consequently, Construction 1.4 does not fail in step s .

We prove now that the nodes corresponding to elements in the s -th α -region are the first active nodes in LR-prefix order at step s . Suppose that at step s there are two active nodes v_{b_1}, v_{b_2} such that v_{b_1} precedes v_{b_2} in LR-prefix order, and $f(v_{b_1}) > s = f(v_{b_2})$. Consequently, $b_2 < b_1$, and both nodes belong to different α -regions. Let v_{a_1}, v_{a_2} denote the parents of v_{b_1}, v_{b_2} , respectively. If $v_{a_1} = v_{a_2}$, then by construction we must have $b_1 < b_2$, which is a contradiction. If v_{a_1} is the root, then v_{b_1} also precedes v_{a_2} in LR-prefix order, and it is either colored or active in the step in which v_{a_2} receives its color. By our induction hypothesis we have $s < f(v_{b_1}) \leq f(v_{a_2}) < s$, which is a contradiction. If v_{a_2} is the root, then since b_2 is not below any bump, the region of b_2 must not come later than that of a_1 (they can be in the same region). In combination with the induction hypothesis, we have $s = f(v_{b_2}) \leq f(v_{a_1}) < s$, which is a contradiction. We conclude that b_1 and b_2 are directly below bumps (a_1, b'_1) and (a_2, b'_2) respectively.

Since v_{b_1}, v_{b_2} are active in the s -th step, it follows that a_1, a_2 are in α -regions before the s -th α -region, with $b_1 < b'_1$ and $b_2 < b'_2$. There are four possibilities.

(i) If $a_1 < a_2 < b'_1 < b'_2$, then either a_1 and a_2 are in the same α -region, or a_2 and b'_1 are in the same α -region by (NC1). In both cases, we would have $b_1 < b_2$, which is a contradiction.

(ii) If $a_1 < a_2 < b'_2 < b'_1$, then a_1 and a_2 belong to different α -regions by (NC2). Thus, v_{a_1} is colored before v_{a_2} is, which implies that v_{b_1} is active before v_{a_2} receives its color. Since v_{b_1} precedes v_{b_2} in LR-prefix order and v_{b_1} is not a descendant of v_{a_2} , we conclude that v_{b_1} precedes v_{a_2} in LR-prefix order, too. It follows that v_{b_1}

receives its color not later than v_{a_2} , and by our induction hypothesis, we have $f(v_{b_1}) \leq f(v_{a_2}) < s$. This contradicts the assumption $f(v_{b_1}) > s$.

(iii) If $a_2 < a_1 < b'_1 < b'_2$, then a_1 and a_2 belong to different α -regions by (NC2). This means that (a_1, b'_1) separates b_1 from (a_2, b'_2) . Since $f(v_{a_1}) < s = f(v_{b_2})$, we have $a_1 < b_2$, thus $b_1 < b'_1 < b_2$ by the separation, which is a contradiction.

(iv) If $a_2 < a_1 < b'_2 < b'_1$, then either a_1 and a_2 belong to the same α -region, or a_1 and b'_2 belong to the same α -region by (NC1). If a_1 and a_2 belong to the same α -region, then v_{a_2} precedes v_{a_1} in LR-prefix order (since $a_2 < a_1$). Moreover, since $f(v_{a_1}) = f(v_{a_2}) \leq s - 1$, by our induction hypothesis these two nodes receive the same color in Construction 1.4, and by Lemma 1.6(i) they are thus incomparable in $T_{\mathbf{P}}$. Since v_{b_2} is a child of v_{a_2} , it follows that v_{b_2} precedes v_{b_1} in LR-prefix order, which is a contradiction. If a_1 and b'_2 belong to the same α -region, then necessarily b_2 belongs to this α -region, too. By induction hypothesis we have $s > f(v_{a_1}) = f(v_{b_2}) = s$, which is a contradiction.

We therefore conclude that all the nodes of $T_{\mathbf{P}}$ from the s -th α -region receive color s in step s of Construction 1.4, and we finish our induction step. We conclude that $T_{\mathbf{P}}$ is an α -tree, and f is its left-aligned coloring. \square

We conclude this section with the proof that Ξ_{nc} is a bijection with inverse Λ_{nc} .

Theorem 1.24. *For every integer composition α , the map Ξ_{nc} is a bijection whose inverse is Λ_{nc} . Moreover, Ξ_{nc} sends internal nodes to bumps.*

Proof. In view of Propositions 1.21 and 1.23, it remains to prove that $\Xi_{\text{nc}} \circ \Lambda_{\text{nc}} = \text{id}$ and $\Lambda_{\text{nc}} \circ \Xi_{\text{nc}} = \text{id}$.

To show that $\Xi_{\text{nc}} \circ \Lambda_{\text{nc}} = \text{id}$, for $\mathbf{P} \in \text{NC}_{\alpha}$, let $T = \Lambda_{\text{nc}}(\mathbf{P})$ be the α -tree obtained from \mathbf{P} by Construction 1.22, and $\mathbf{P}' = \Xi_{\text{nc}}(T)$ the noncrossing α -partition obtained from T by Construction 1.20.

Since every bump of \mathbf{P} corresponds to an internal node of T , and every internal node of T corresponds to a bump of \mathbf{P}' , we conclude that \mathbf{P} and \mathbf{P}' have the same number of bumps. Let (a, b) be a bump of \mathbf{P} . The node v_a is thus the parent of v_b . Let i be an element of \mathbf{P} such that v_i is a child of v_a in T . By Construction 1.22, we have $a < i \leq b$. Hence, v_i weakly precedes v_b in LR-prefix order. This is valid for any child of v_a in T , implying that v_b is the rightmost child of v_a in T , leading to a bump (a, b) in \mathbf{P}' by Construction 1.20. We thus conclude that $\mathbf{P} = \mathbf{P}'$.

To show that $\Lambda_{\text{nc}} \circ \Xi_{\text{nc}} = \text{id}$, for $T \in \mathbb{T}_{\alpha}$, let C be the left-aligned coloring of T , $\mathbf{P} = \Xi_{\text{nc}}(T)$ and $T' = \Lambda_{\text{nc}}(\mathbf{P})$. We now prove $T = T'$. For u a non-leaf node of T and v a child of u , let a, b be the elements of \mathbf{P} corresponding to u, v , and let u', v' be the nodes of T' corresponding to a, b . We first show that v' is a child of u' in T' .

If u is the root, then b is not below any bump in \mathbf{P} , meaning that u' is the root of T' and v' one of its children by Construction 1.22. We now suppose that u is not the root. If v is the rightmost child of u , then (a, b) is a bump of \mathbf{P} , and it follows from Construction 1.22 that v' is the rightmost child of u' , too.

It remains to consider the case that v is not the rightmost child of u . Let v_r be the rightmost child of u , and let b_r be the element of \mathbf{P} corresponding to v_r . It follows that (a, b_r) is a bump of \mathbf{P} . Moreover, by Construction 1.4 we have $C(u) < C(v) \leq C(v_r)$, and Construction 1.20 then implies that $a < b < b_r$. Hence, b lies below some bump, which implies that u' is not the root of T' . Moreover, b lies directly below a certain bump (a_*, b_*) .

Suppose that $a_* \neq a$. Let u_* and v_* denote the nodes of T that correspond to a_* and b_* . In particular, a_* and b belong to different α -regions, which implies $C(u_*) < C(v)$. Since $b < b_*$, we have $C(v) \leq C(v_*)$. By Construction 1.20, the node v_* is the rightmost child of u_* , and we conclude that $b \leq b_* < b_r$. There are two cases.

(i) If $a < a_* < b_* < b_r$, then a and a_* are in different α -regions by (NC2). Thus, $C(u) < C(u_*)$, which means that v is active in the step in which u_* receives its color in Construction 1.4. Since $C(u_*) < C(v)$, it follows that u_* precedes v in LR-prefix order. Consequently, v_* precedes v in LR-prefix order, too, as u_* and v_* are consecutive in LR-prefix order. But this implies $b_* < b$, which is a contradiction.

(ii) If $a_* < a < b_* < b_r$, then either a and a_* belong to the same α -region, or a and b_* belong to the same α -region by (NC1). If a and a_* belong to the same α -region, then it follows that $C(u) = C(u_*)$ and that u_* precedes u in LR-prefix order. Hence, v and v_* become active in the same iteration of Construction 1.4 and v_* precedes v in LR-prefix order, too. Thus, we have $C(v_*) \leq C(v)$, which by assumption yields $C(v_*) = C(v)$. It follows that b and b_* belong to the same α -region, and by Construction 1.20 we obtain that v precedes v_* in LR-prefix order, which is a contradiction. If a and b_* belong to the same α -region, then combining $b < b_*$ with the case condition, we have $a < b < b_*$, and it follows that a and b also belong to the same α -region. This implies $C(u) = C(v)$, which contradicts Lemma 1.6(i), since v is a child of u .

We conclude that $a_* = a$, which implies that b lies directly below the bump (a, b_r) . We observe that the order of the children a node in T and T' depends only on the order of the elements of \mathbf{P} by Constructions 1.20 and 1.22. We thus conclude that $T = T'$. \square

1.2.5. Dyck Paths and Left-Aligned Colorable Trees. The following construction associates an α -Dyck path with every α -tree.

Construction 1.25. Given $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$ an integer composition with $|\alpha| = n$, let $T \in \mathbb{T}_\alpha$. For $k \in [r]$, let $s_k = \alpha_1 + \alpha_2 + \dots + \alpha_k$ as defined at the beginning of Section 1.1, which is precisely the number of nodes of color $\leq k$. Moreover, let a_k denote the number of active nodes in step $(k+1)$ of Construction 1.4 before coloring, which is precisely the number of nodes of color $> k$ whose parents have color $\leq k$. It is immediate that $s_k + a_k \leq n$ for all possible k .

We now construct a Dyck path which will have exactly one valley for each internal node of T . Suppose that the rightmost child of the i -th internal node of color k is the j -th active node in the $(k+1)$ -st step of Construction 1.4 before coloring. We mark the coordinate (p, q) in the plane, where $p = s_k - i + 1$ and $q = s_k + a_k - j$. We then return the unique Dyck path μ_T from $(0, 0)$ to (n, n) that has valleys precisely at the marked coordinates.

Remark 1.26. There is a simpler way to informally describe Construction 1.25. We first draw the bounce path ν_α , which encloses triangular regions with the main diagonal. Each such region corresponds to one color. For each internal node u of color k , with v its rightmost child, suppose that u is the i -th node from right to left of color k , and v the j -th active node from right to left at step $(k+1)$ of Construction 1.4. We then add a dot on the i -th column from left to right intersecting the triangular region of color k , in the j -th cell above the region. These dots indicate the valleys of μ_T . The description here is equivalent to Construction 1.25.




Figure 5. An illustration of Construction 1.25 for $\alpha = (1, 3, 1, 2, 4, 3)$.

See Figure 5 for an illustration of Construction 1.25. We now prove that every Dyck path obtained by Construction 1.25 is in fact an α -Dyck path, so that we obtain the following map:

$$(7) \quad \Xi_{\text{nn}}: \mathbb{T}_\alpha \rightarrow \mathcal{D}_\alpha, \quad T \mapsto \mu_T.$$

Proposition 1.27. *For an α -tree T , the Dyck path μ_T obtained by Construction 1.25 is an α -Dyck path.*

Proof. Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$ be a composition of n . Recall from Section 1.1.3 that the α -bounce path is $v_\alpha = N^{\alpha_1} E^{\alpha_1} N^{\alpha_2} E^{\alpha_2} \dots N^{\alpha_r} E^{\alpha_r}$. We need to show that μ_T is weakly above v_α , or equivalently, that all valleys of μ_T lie weakly above v_α .

Let u be the i -th node of color k in T with respect to LR-prefix order. If u is a leaf, then it does not contribute a valley of μ_T , and thus need not be considered. Let v be the rightmost child of u . Suppose that v is the j -th active node in the k -th step of Construction 1.4.

By construction, the corresponding valley of μ_T has coordinate (p, q) with $p = s_k - i + 1$ and $q = s_k + a_k - j$, where a_k is the number of active nodes in the $(k + 1)$ -st step of Construction 1.4 before coloring. By construction we have $i \leq \alpha_k$, therefore $s_{k-1} < p \leq s_k$. Moreover, since $j \leq a_k$, we have $s_k \leq q < s_k + a_k$. Therefore, (p, q) is weakly above v_α , with which we conclude the proof. \square

Let us now describe how to obtain a colored tree from an α -Dyck path.

Construction 1.28. Given $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$ an integer composition, let $\mu \in \mathcal{D}_\alpha$. We initialize our algorithm with the tree T_0 which consists of a single root node v_0 (implicitly colored by color 0), and the empty coloring C_0 .

At the k -th step, we construct a (partially) colored tree (T_k, C_k) from (T_{k-1}, C_{k-1}) by adding as many children to the i -th node of color $k - 1$ as there are north-steps in μ with x -coordinate equal to $s_{k-1} - i + 1$. After all children have been added, we extend C_{k-1} to C_k by assigning the value k to the first α_k uncolored nodes in LR-prefix order. The algorithm fails if there are not enough uncolored nodes. After r steps, we obtain a colored tree (T_r, C_r) , and we return the plane rooted tree $T_\mu = T_r$.

We now prove that the plane rooted tree obtained in Construction 1.28 is indeed compatible with α , which allows for the definition of a map

$$(8) \quad \Lambda_{\text{nn}}: \mathcal{D}_\alpha \rightarrow \mathbb{T}_\alpha, \quad \mu \mapsto T_\mu.$$

Proposition 1.29. *For an α -Dyck path μ , the plane rooted tree T_μ obtained from Construction 1.28 is compatible with α .*

Proof. Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$. We first prove that we never run out of uncolored nodes in Construction 1.28. To that end, let $s_0 = 0$. For $i \in \{0, 1, \dots, r-1\}$ we denote by m_i the maximal y -coordinate of a lattice point on μ that has x -coordinate s_i . Since μ is weakly above v_α , we have $m_i \geq s_{i+1}$.

We now prove by induction on $k \in [r]$ that the number of uncolored non-root nodes at the end of step k in Construction 1.28 is precisely $m_{k-1} - s_k$. The base case $k = 1$ holds, because in the first step, we add exactly m_0 children to the root, and we color $\alpha_1 = s_1$ of them. Now suppose that the claim holds up to the k -th step of Construction 1.28. In particular, at the beginning of the $(k+1)$ -st step we already have $m_{k-1} - s_k$ uncolored nodes. In the $(k+1)$ -st step, we add $m_k - m_{k-1}$ new nodes, and color α_{k+1} of the uncolored nodes. This leaves us with $m_{k-1} - s_k + m_k - m_{k-1} - \alpha_{k+1} = m_k - s_{k+1}$ uncolored nodes, and we conclude the induction.

Now we prove that T_μ is an α -tree. Let f be the full coloring of T_μ obtained in Construction 1.28. We show that f is the left-aligned coloring of T_μ . Let f_k denote the restriction of f to the nodes of T_μ that have f -color at most k . We now prove that Construction 1.4 does not failed up to step k and that f_k agrees with C_k defined in Construction 1.4, once again by induction on $k \in [r]$.

For the base case $k = 1$, by construction, the root of T_μ has $m_0 \geq \alpha_1$ children, and we have $f(v) = 1$ if and only if v belongs to the α_1 leftmost children of the root. We see that Construction 1.4 does not fail in the first step, and the same nodes receive color 1 in both constructions.

Suppose that the induction hypothesis holds for k . The active nodes with respect to C_k are thus precisely those whose parents have C_k -color at most k . In step $k+1$ of Construction 1.28, we add precisely $m_k - m_{k-1}$ new nodes, and after this addition, the parents of the uncolored nodes have f_k -color at most k . By induction hypothesis, the uncolored nodes with respect to f_k are precisely the active nodes with respect to C_k . In view of the first part of this proof, Construction 1.4 does not fail in step $k+1$, and since we color the α_{k+1} leftmost nodes in LR-prefix order by $k+1$ in both constructions, the induction step is established.

We conclude that T_μ is compatible with α and the coloring from Construction 1.28 is precisely the left-aligned coloring of T_μ . \square

Remark 1.30. It follows from the proof of Proposition 1.29 that after we have created all the nodes in the $(k+1)$ -st step of Construction 1.4, there are $m_k - s_k$ active nodes before coloring, and $m_k - s_{k+1}$ active nodes after coloring.

Lemma 1.31. *Given $T \in \mathbb{T}_\alpha$, let u be an internal node in T and $\vec{p} = (p, q)$ the valley of $\mu = \Xi_{\text{nn}}(T)$ corresponding to u . The number of children of u is equal to that of consecutive north-steps in $\Xi_{\text{nn}}(T)$ immediately after \vec{p} .*

Proof. Let c_1 be the number of children of u in T , and c_2 the number of consecutive north-steps in $\Xi_{\text{nn}}(T)$ immediately after \vec{p} . Suppose that u has color k , and that the rightmost child of u is the j -th active node in LR-prefix order in the $(k+1)$ -st step of Construction 1.4 before coloring among a total of a_k active nodes. Moreover, let m_k be the maximal y -coordinate of a lattice point on μ with x -coordinate s_k . It follows from Remark 1.30 that $a_k = m_k - s_k$. By Construction 1.25 we have $q = s_k + a_k - j = m_k - j$.

If u is the first internal node in its color group, then $c_1 = j$. By construction, \vec{p} is the last valley whose x -coordinate is at most s_k , which implies that $c_2 = m_k - q = j$.

Otherwise let u_* be the internal node of color k that immediately precedes u in LR-prefix order. Suppose that the rightmost child of u_* is the j_* -th active node in LR-prefix order at step $k + 1$ of Construction 1.4 before coloring. It follows that $c_1 = j - j_*$. Moreover, let $\vec{p}_* = (p_*, q_*)$ be the valley of μ corresponding to u_* . We know that \vec{p}_* is the valley immediately after \vec{p} , and we have $c_2 = q_* - q$. Since $q_* = s_k + a_k - j_* = m_k - j_*$ by Construction 1.25, we conclude that $c_2 = j - j_*$. We thus have $c_1 = c_2$ in all cases. \square

We now prove that Ξ_{nn} is a bijection with inverse Λ_{nn} .

Theorem 1.32. *For every integer composition $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$, the map Ξ_{nn} is a bijection with inverse Λ_{nn} . Moreover, Ξ_{nn} sends internal nodes to valleys.*

Proof. In view of Propositions 1.27 and 1.29, it remains to prove that $\Xi_{\text{nn}} \circ \Lambda_{\text{nn}} = \text{id}$ and $\Lambda_{\text{nn}} \circ \Xi_{\text{nn}} = \text{id}$.

To show that $\Xi_{\text{nn}} \circ \Lambda_{\text{nn}} = \text{id}$, for $\mu \in \mathcal{D}_\alpha$, let $T = \Lambda_{\text{nn}}(\mu)$ be the α -tree obtained from μ by Construction 1.28, and $\mu' = \Xi_{\text{nn}}(T)$ the α -Dyck path obtained from T by Construction 1.25. Since every valley of μ corresponds to an internal node of T , which in turn corresponds to a valley of μ' , we conclude that μ and μ' have the same number of valleys.

Let (p, q) be a valley of μ , and u the corresponding internal node of T , which is the $(s_k - p + 1)$ -st node of color k in the LR-prefix order of T . Let m_k be the maximal y -coordinate of a lattice point on μ with x -coordinate s_k . By Remark 1.30, the number of active nodes in the $(k + 1)$ -st step of Construction 1.4 (before coloring) is precisely $a_k = m_k - s_k$. The rightmost child of u is thus the $(m_k - q)$ -th active node at step k in the LR-prefix order of T . By Construction 1.28, the internal node u and its rightmost child contribute a valley (p', q') to μ' , where $p' = s_k - (s_k - p + 1) + 1 = p$ and $q' = s_k + a_k - (m_k - q) = q$. Thus, the valleys of μ and μ' agree, and it follows that $\mu = \mu'$.

To show that $\Lambda_{\text{nn}} \circ \Xi_{\text{nn}} = \text{id}$, for $T \in \mathbb{T}_\alpha$, let $\mu = \Xi_{\text{nn}}(T)$ be the α -Dyck path obtained from T by Construction 1.25, and $T' = \Lambda_{\text{nn}}(\mu)$ the α -tree obtained from μ by Construction 1.28. For $s \in [r]$, let T_s, T'_s denote the induced subtrees of T, T' , respectively, that consist of the nodes of color at most s . We prove by induction on s that $T_s = T'_s$ holds for all $0 \leq s \leq r$.

The base case $s = 0$ is trivially true, since both trees T_0 and T'_0 consist only of a root node. Suppose that $T_{s-1} = T'_{s-1}$. Suppose that u is a node of T_{s-1} that has $k > 0$ children in T . By Lemma 1.31, either μ starts with k north-steps (if u is the root), or the valley corresponding to u is followed by k north-steps. By Construction 1.28, it follows that u' also has k children. We also notice that, by construction, the positions of u and u' among all active nodes in T_{s-1} and T'_{s-1} are the same. Therefore, we have $T_s = T'_s$. \square

We may now prove our first main result.

Theorem I. *For every integer composition α , there is an explicit bijection from the set \mathbb{T}_α of α -trees to each of the following: the set $\mathfrak{S}_\alpha(231)$ of $(\alpha, 231)$ -avoiding permutations, the set NC_α of noncrossing α -partitions, and the set \mathcal{D}_α of α -Dyck paths.*

Proof. The bijections are established in Theorems 1.19, 1.24, and 1.32. \square

We conclude this section by showing that the bijections Θ_1 and Θ_2 from Sections 1.2.1 and 1.2.2, respectively, can be recovered using the new bijections from Sections 1.2.3–1.2.5.

Proposition 1.33. *For every $\mu \in \mathcal{D}_\alpha$, we have $\Theta_1(\bar{\mu}) = \Xi_{\text{nc}} \circ \Lambda_{\text{nn}}(\mu)$.*

Proof. Let $T = \Lambda_{\text{nn}}(\mu)$, $\mathbf{P} = \Xi_{\text{nc}}(T)$, and $\mathbf{P}' = \Theta_1(\bar{\mu})$. By Theorem 1.32, we know that every valley of μ corresponds to an internal node u of T , with its coordinates dictating the exact position of u and its rightmost child v in the α -regions. By Theorem 1.24, each such pair (u, v) corresponds to a bump (a, b) of \mathbf{P} , where a and b are determined by the positions of u and v . If we now take the corresponding valley of $\bar{\mu}$, and apply Construction 1.7, then we see that (a, b) is also a bump of \mathbf{P}' . We conclude that $\mathbf{P} = \mathbf{P}'$. \square

Proposition 1.34. *For every $\mathbf{P} \in \text{NC}_\alpha$, we have $\Theta_2(\mathbf{P}) = \Xi_{\text{perm}} \circ \Lambda_{\text{nc}}(\mathbf{P})$.*

Proof. Let $T = \Lambda_{\text{nc}}(\mathbf{P})$, $w = \Xi_{\text{perm}}(T)$, and $w' = \Theta_2(\mathbf{P})$. We prove by induction on $n = |\alpha|$ that $w = w'$. The base case $n = 0$ clearly holds. Now suppose that our claim holds for all α' with $|\alpha'| < n$.

We take the block of \mathbf{P} containing 1, and denote it by $P_\circ = \{i_1, i_2, \dots, i_s\}$ with $i_1 < i_2 < \dots < i_s$ and $i_1 = 1$. By Construction 1.11, we have $w'(i_1) = w'_{i_2} + 1 = \dots = w'_{i_s} + s - 1$. The remaining values of w' are determined inductively from two parabolic noncrossing partitions \mathbf{P}_1 and \mathbf{P}_2 with strictly less than n elements each. By induction hypothesis we have $\Theta_2(\mathbf{P}_i) = \Xi_{\text{perm}} \circ \Lambda_{\text{nc}}(\mathbf{P}_i)$ for $i \in \{1, 2\}$.


On the other hand, by Construction 1.22, the elements in P_\circ correspond to nodes in T on the rightmost branch of the first child u of the root, that is, the set of nodes visited from u by always moving to the rightmost child when possible. The label of u in Construction 1.14 is always $w(1)$, and the label of the rightmost child of any node v is always one less than the label of v itself. Finally, by comparing Constructions 1.11 and 1.22 we see that the elements contributing to $w'(1)$ are precisely the nodes in the subtree of u , and we conclude that w and w' agree on P_\circ .

Now let T_2 be the induced subtree of T containing all the elements that are descendants of all but the first root children, and let $w_2 = \Xi_{\text{perm}}(T_2)$. Moreover, we construct a new tree T_1 from T by first deleting the root and all nodes of T_2 , and then contracting the rightmost branch of u into a new root node while keeping the LR-prefix order of the remaining nodes. Let $w_1 = \Xi_{\text{perm}}(T_1)$. The elements of T_1 and T_2 , respectively, are in the same relative order as in T with respect to LR-postfix order. Hence, w can be reconstructed from the values $w(i_1), w(i_2), \dots, w(i_s)$, the values in w_1 and the values in w_2 increased by $w(1)$. This is precisely the way that w' was created, so we conclude that $w = w'$. \square

Part 2. Tamari Lattices in Parabolic Cataland

In the middle part of this article, we equip each of the sets $\mathfrak{S}_\alpha(231)$ and \mathcal{D}_α with a partial order, and show that the resulting posets are isomorphic (Theorem II). Both posets are in fact lattices and generalize the well-known *Tamari lattice* \mathcal{T}_n introduced in [30, 31]. See [27] for a recent survey on topics related to the Tamari lattices.

For the proof of this isomorphism, we need some lattice-theoretic background, which is only relevant for the current section. We have thus decided to put it in Appendix A, which can be consulted whenever necessary.

Figure 6. The lattice $\mathcal{T}_{(2,2,1)}$ and its Galois graph.

2.1. PARABOLIC TAMARI LATTICES

For every permutation $w \in \mathfrak{S}_n$ we define its *inversion set* by

$$\text{inv}(w) \stackrel{\text{def}}{=} \{(i, j) \mid i < j \text{ and } w(i) > w(j)\}.$$

The elements of $\text{inv}(w)$ are the *inversions* of w . In particular, every descent of w is also an inversion. The *(left) weak order* on \mathfrak{S}_n is defined by setting $w_1 \leq_L w_2$ if and only if $\text{inv}(w_1) \subseteq \text{inv}(w_2)$ for all $w_1, w_2 \in \mathfrak{S}_n$.

Definition 2.1. The *parabolic Tamari lattice* $\mathcal{T}_\alpha \stackrel{\text{def}}{=} (\mathfrak{S}_\alpha(231), \leq_L)$ is defined to be the restriction of the weak order to $(\alpha, 231)$ -avoiding permutations.

Figure 6a shows the parabolic Tamari lattice $\mathcal{T}_{(2,2,1)}$. The elements of $\mathfrak{S}_{(2,2,1)}(231)$ are shown in one-line notation, where the $(2, 2, 1)$ -regions are separated by vertical bars. Theorem 1.1 in [26] states that \mathcal{T}_α is indeed a lattice. Moreover, the following structural result holds.

Theorem 2.2 ([25, Theorem 1.3]). *For every integer composition α , the parabolic Tamari lattice \mathcal{T}_α is extremal and congruence uniform.*

It follows from [25, Corollary 4.5] that the join-irreducible elements of \mathcal{T}_α are precisely the $(\alpha, 231)$ -avoiding permutations with a unique descent. As a consequence of Proposition A.4, the Galois graph of \mathcal{T}_α (defined in Section A.2) can be realized as a directed graph whose vertices are pairs of integers. Let us make this more precise.

Theorem 2.3 ([25, Theorem 1.8]). *Let α be an integer composition. The Galois graph of \mathcal{T}_α is isomorphic to the directed graph with vertex set*

$$\{(a, b) \mid 1 \leq a < b \leq |\alpha| \text{ and } a, b \text{ belong to different } \alpha\text{-regions}\},$$

where there is a directed edge $(a_1, b_1) \rightarrow (a_2, b_2)$ if and only if $(a_1, b_1) \neq (a_2, b_2)$ and

- either a_1 and a_2 belong to the same α -region and $a_1 \leq a_2 < b_2 \leq b_1$,

- or a_1 and a_2 belong to different α -regions and $a_2 < a_1 < b_2 \leq b_1$, where a_1 and b_2 belong to different α -regions, too.

Figure 6b illustrates this result, and shows the Galois graph of $\mathcal{T}_{(2,2,1)}$. If $\alpha = (1, 1, \dots, 1)$ and $n = |\alpha|$, then the classical Tamari lattice introduced in [30] is $\mathcal{T}_n \stackrel{\text{def}}{=} \mathcal{T}_{(1,1,\dots,1)}$ [6, Theorem 9.6(ii)].

2.1.1. On the Duality of the Parabolic Tamari Lattices. We now show that reversing the composition α , produces a parabolic Tamari lattice which is dual to \mathcal{T}_α .

Recall from Section 1.1 that, for an integer composition $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$, the reverse composition is $\bar{\alpha} = (\alpha_r, \alpha_{r-1}, \dots, \alpha_1)$.

Moreover, for every $n > 0$, and every permutation $w \in \mathfrak{S}_n$ we define the *reverse permutation* \bar{w} by setting $\bar{w}(i) = w(n+1-i)$ for $i \in [n]$. If $w \in \mathfrak{S}_\alpha$, and we additionally reverse the order of the entries of each $\bar{\alpha}$ -region of \bar{w} , then we obtain a permutation $\tilde{w} \in \mathfrak{S}_{\bar{\alpha}}$.

For example, let $w = 24|157|3|6 \in \mathfrak{S}_{(2,3,1,1)}$. Then we have $\bar{w} = 6375142$, and $\tilde{w} = 6|3|157|24 \in \mathfrak{S}_{(1,1,3,2)}$.

Lemma 2.4. *Let $w_1, w_2 \in \mathfrak{S}_\alpha$. We have $\text{inv}(w_1) \subseteq \text{inv}(w_2)$ if and only if $\text{inv}(\tilde{w}_1) \supseteq \text{inv}(\tilde{w}_2)$.*

Proof. Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$ be an integer composition. We only prove the implication “ $\text{inv}(w_1) \subseteq \text{inv}(w_2)$ implies $\text{inv}(\tilde{w}_1) \supseteq \text{inv}(\tilde{w}_2)$ ”, the other implication follows analogously.

Let $w_1, w_2 \in \mathfrak{S}_\alpha$ such that $\text{inv}(w_1) \subseteq \text{inv}(w_2)$, and choose integers a, b in different $\bar{\alpha}$ -regions such that $(a, b) \notin \text{inv}(\tilde{w}_1)$.

Let a', b' be the unique indices with $w_1(a') = \tilde{w}_1(a)$ and $w_1(b') = \tilde{w}_1(b)$. Since the entries in each $\bar{\alpha}$ -region are ordered linearly, and by going from \tilde{w}_1 to w_1 we revert this order, we conclude that $w_2(a') = \tilde{w}_2(a)$ and $w_2(b') = \tilde{w}_2(b)$.

Suppose that a belongs to the i -th $\bar{\alpha}$ -region, and suppose that b belongs to the j -th $\bar{\alpha}$ -region. Since $a < b$ we have $i < j$. (By construction, it follows from $\tilde{w}_1 \in \mathfrak{S}_{\bar{\alpha}}$ that there are no inversions within the same $\bar{\alpha}$ -region.) By construction we conclude that a and a' belong to the $(r+1-i)$ -th α -region and b and b' belong to the $(r+1-j)$ -th α -region. It follows that $b' < a'$.

Since $(a, b) \notin \text{inv}(\tilde{w}_1)$ we must have $w_1(a') = \tilde{w}_1(a) < \tilde{w}_1(b) = w_1(b')$, which implies $(b', a') \in \text{inv}(w_1) \subseteq \text{inv}(w_2)$. It follows that $\tilde{w}_2(b) = w_2(b') > w_2(a') = \tilde{w}_2(a)$, which yields $(a, b) \notin \text{inv}(\tilde{w}_2)$.


By contraposition we conclude that $\text{inv}(\tilde{w}_1) \supseteq \text{inv}(\tilde{w}_2)$. \square

Definition 2.5. In a permutation $w \in \mathfrak{S}_\alpha$, an $(\alpha, 132)$ -*pattern* is a triple of indices $i < j < k$ each in different α -regions such that $w(i) < w(k) < w(j)$ and $w(k) = w(i) + 1$. A permutation in \mathfrak{S}_α without $(\alpha, 132)$ -patterns is $(\alpha, 132)$ -*avoiding*.

We denote the set of all $(\alpha, 132)$ -avoiding permutations of \mathfrak{S}_α by $\mathfrak{S}_\alpha(132)$, and consider how reversal acts on $(\alpha, 231)$ -patterns.

Lemma 2.6. *A permutation $w \in \mathfrak{S}_\alpha$ has an $(\alpha, 231)$ -pattern if and only if $\tilde{w} \in \mathfrak{S}_{\bar{\alpha}}$ has an $(\bar{\alpha}, 132)$ -pattern.*

Proof. We only prove the implication “ w has an $(\alpha, 231)$ -pattern implies that \tilde{w} has an $(\bar{\alpha}, 132)$ -pattern”, the other implication follows analogously.

Figure 7. The parabolic Tamari lattice $\mathcal{T}_{(1,2,2)}$.

Let $w \in \mathfrak{S}_\alpha$ have an $(\alpha, 231)$ -pattern (i, j, k) . Let $a = w(i)$, $b = w(j)$ and $c = w(k)$. By Definition 1.1 we have $i < j < k$ in different α -regions, $a < b$ and $a = c + 1$. Let i', j', k' be such that $\tilde{w}(i') = a$, $\tilde{w}(j') = b$ and $\tilde{w}(k') = c$. As in the proof of Lemma 2.4 we find that $k' < j' < i'$ are in different α -regions. By Definition 2.5, (k', j', i') is an $(\alpha, 132)$ -pattern of \tilde{w} . \square

Theorem 2.7. *For every integer composition α , the lattice \mathcal{T}_α is isomorphic to the dual of $\overline{\mathcal{T}}_\alpha$.*

Proof. We use the fact that \mathcal{T}_α arises as a quotient lattice of $(\mathfrak{S}_\alpha, \leq_L)$ [26, Proposition 3.18]. It follows from [26, Section 3.4] that the set of greatest elements of the corresponding congruence-classes is precisely $\mathfrak{S}_\alpha(132)$, and the set of least elements of the corresponding congruence-classes is $\mathfrak{S}_\alpha(231)$.

Let us define a partial order on the symmetric group by setting $w_1 \leq_R w_2$ if and only if $\text{inv}(w_1) \supseteq \text{inv}(w_2)$. We conclude that the dual lattice of $\mathcal{T}_\alpha = (\mathfrak{S}_\alpha(231), \leq_L)$ is isomorphic to $(\mathfrak{S}_\alpha(132), \leq_R)$. Lemmas 2.4 and 2.6 now imply that $(\mathfrak{S}_\alpha(132), \leq_R) \cong (\mathfrak{S}_{\bar{\alpha}}(231), \leq_L) = \overline{\mathcal{T}}_\alpha$. \square

Compare Figures 6a and 7 for an illustration of Theorem 2.7.

2.2. ν_α -TAMARI LATTICES

In this section we define a partial order on the set \mathcal{D}_α of all α -Dyck paths. This construction can in fact be carried out for any set of Dyck paths that stay weakly above a fixed path ν formed by north and east steps. See [9, 28] for the general setup. All of the following structural properties of ν_α -Tamari lattices hold in this general case, too.

Let $\mu \in \mathcal{D}_\alpha$, and consider a lattice point \vec{p} in μ . We define $\text{horiz}_\alpha(\vec{p})$ to be the maximal number of east-steps that can be added after \vec{p} without crossing ν_α . For any valley \vec{p} of μ , let \vec{q} be the first lattice point on μ such that $\text{horiz}_\alpha(\vec{p}) = \text{horiz}_\alpha(\vec{q})$. We denote by $\mu[\vec{p}, \vec{q}]$ the subpath of μ from \vec{p} to \vec{q} . Let μ' be the unique α -Dyck path which arises from μ by swapping the subpath $\mu[\vec{p}, \vec{q}]$ with the east step preceding \vec{p} . Let us abbreviate this operation by $\mu <_\alpha \mu'$.

This operation induces an acyclic relation on \mathcal{D}_α , and we denote by \leq_α its reflexive and transitive closure.

Definition 2.8. The poset $\mathcal{T}_{v_\alpha} \stackrel{\text{def}}{=} (\mathcal{D}_\alpha, \leq_\alpha)$ is the v_α -*Tamari lattice*.

The $v_{(1,2,2)}$ -Tamari lattice is shown in Figure 8a.

Theorem 2.9. For every integer composition α , the v_α -lattice \mathcal{T}_{v_α} is extremal and congruence uniform.

Proof. First of all, [28, Theorem 1] states that \mathcal{T}_{v_α} is indeed a lattice. Moreover, [28, Theorem 3] states that \mathcal{T}_{v_α} is an interval in $\mathcal{T}_{2|\alpha|+1}$.

It was shown in [23, Theorem 22] that \mathcal{T}_N is extremal for every integer $N > 0$. However, it follows from [23, Theorem 14(ii)] that intervals of extremal lattices need not be extremal. Fortunately, [33, Theorem 9] implies that \mathcal{T}_N is *trim*. (This is a property that is somewhat stronger than extremality.) Theorem 1 in [33] states that every interval of a trim lattice is trim, too. We conclude that \mathcal{T}_{v_α} is trim, and therefore extremal.

Theorem 3.5 in [15] states that \mathcal{T}_N is congruence uniform for every integer $N > 0$, and [10, Theorem 4.3] implies that intervals of congruence-uniform lattices are congruence uniform again. We conclude that \mathcal{T}_{v_α} is congruence uniform. \square

We conclude this section with another useful result, which nicely parallels Theorem 2.7.

Theorem 2.10 ([28, Theorem 2]). For every integer composition α , the lattice \mathcal{T}_{v_α} is isomorphic to the dual of $\mathcal{T}_{\bar{v}_\alpha}$.

2.2.1. v_α -Bracket Vectors. Following [9], we may as well represent \mathcal{T}_{v_α} as a lattice of certain integer vectors under componentwise order.

Let $n = |\alpha|$. The *minimal v_α -bracket vector* \mathbf{b}^{\min} is the vector with $2n + 1$ entries that contains the y -coordinates of the lattice points of v_α read in order from $(0, 0)$ to (n, n) . The *fixed positions* are the elements of the set $F = \{f_0, f_1, \dots, f_n\}$, where f_i denotes the last occurrence of i in \mathbf{b}^{\min} .

Definition 2.11. A *v_α -bracket vector* is a vector $\mathbf{b} = (b_1, b_2, \dots, b_{2n+1})$ that satisfies

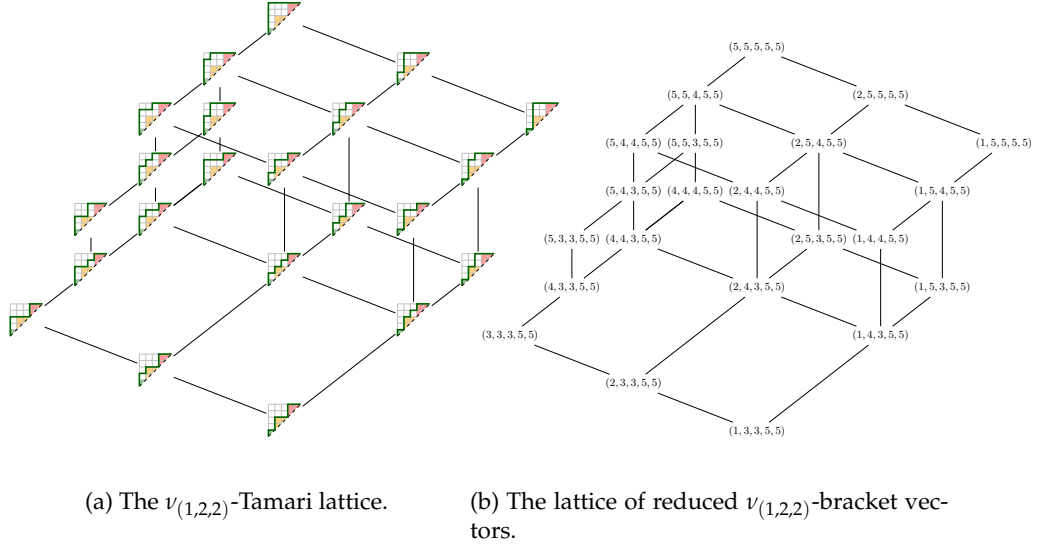
- $b_{f_i} = i$ for $0 \leq i \leq n$;
- $\mathbf{b}_i^{\min} \leq b_i \leq n$ for all i ;
- if $b_i = k$, then $b_j \leq k$ for all $i \leq j \leq f_k$.

Theorem 2.12 ([9, Theorem 4.2]). For every integer composition α , the lattice of v_α -bracket vectors under componentwise order is isomorphic to \mathcal{T}_{v_α} .

For $\mu \in \mathcal{D}_\alpha$ we may directly construct the corresponding v_α -bracket vector as follows; see [9, Section 4.3].

Construction 2.13. Let $\mu \in \mathcal{D}_\alpha$ for some integer composition α . Suppose that μ runs through k_i lattice points of y -coordinate i . We start with an empty vector \mathbf{b}_0 of length $2n + 1$. In the i -th step we construct \mathbf{b}_i from \mathbf{b}_{i-1} by filling the k_i rightmost available spaces before and including the fixed position f_i with the value i . After n steps we return the bracket vector $\mathbf{b}_\mu = \mathbf{b}_n$.

Given a v_α -bracket vector \mathbf{b} , we denote by $\underline{\mathbf{b}}$ the *reduced v_α -bracket vector*, i.e. the vector that contains all entries of \mathbf{b} in the same order, except for those at the fixed positions. It is clear that $\mathbf{b} \leq \mathbf{b}'$ if and only if $\underline{\mathbf{b}} \leq \underline{\mathbf{b}'}$, where \leq is componentwise order. Figure 8b shows the lattice of reduced $v_{(1,2,2)}$ -bracket vectors.

Figure 8. Two representations of a ν_α -Tamari lattice.

2.2.2. **The Galois Graph of \mathcal{T}_{ν_α} .** By Theorem 2.9, the lattice \mathcal{T}_{ν_α} is extremal for every integer composition α , and by Theorem A.1 we may represent \mathcal{T}_{ν_α} by its Galois graph. In this section, we characterize the Galois graph of \mathcal{T}_{ν_α} .

Since \mathcal{T}_{ν_α} is also congruence uniform, we may apply Proposition A.4 and define $\text{Gal}(\mathcal{T}_{\nu_\alpha})$ in terms of the join-irreducible elements of \mathcal{T}_{ν_α} . In general, however, it is much easier to describe meet-irreducible Dyck paths (because they have a unique valley). In view of Theorem 2.10, we may regard $\text{Gal}(\mathcal{T}_{\nu_\alpha})$ as a directed graph with vertex set $\mathcal{M}(\mathcal{T}_{\bar{\nu}_\alpha})$, where there is a directed edge $\mu \rightarrow \bar{\mu}$ if and only if $\mu \neq \bar{\mu}$ and $\bar{\mu} \geq \bar{\mu}^* \wedge \mu$.

Since meet-irreducible Dyck paths are precisely those with a unique valley, we may write $\mu_{p,q}$ for the unique meet-irreducible element of $\mathcal{T}_{\bar{\nu}_\alpha}$ that has its only valley at (p, q) . Moreover, we abbreviate $k^{(m)} \stackrel{\text{def}}{=} \underbrace{k, k, \dots, k}_m$ for any integer k , and any nonnegative integer m .


Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$ be a composition of n , and recall that we have defined $s_i = \alpha_1 + \alpha_2 + \dots + \alpha_i$ and $t_i = n - s_{r-i}$ for all $i \in \{0, 1, \dots, r\}$ in the beginning of Section 1.1.

Lemma 2.14. *Let $\mu_{p,q} \in \mathcal{M}(\mathcal{T}_{\bar{\nu}_\alpha})$, where $t_k \leq q < t_{k+1}$ for some $k \in [r]$. The reduced $\bar{\nu}_\alpha$ -bracket vector of $\mu_{p,q}$ is*

$$\underline{\mathbf{b}}(\mu_{p,q}) = \left(n^{(t_k-p)}, q^{(p)}, n^{(n-t_k)} \right).$$

Moreover, the reduced $\bar{\nu}_\alpha$ -bracket vector of the unique upper cover $\mu_{p,q}^*$ of $\mu_{p,q}$ is

$$\underline{\mathbf{b}}(\mu_{p,q}^*) = \begin{cases} \left(n^{(t_k-p)}, q+1, q^{(p-1)}, n^{(n-t_k)} \right), & \text{if } q < t_{k+1} - 1, \\ \left(n^{(t_k-p+1)}, q^{(p-1)}, n^{(n-t_k)} \right), & \text{if } q = t_{k+1} - 1. \end{cases}$$

Figure 9. The Galois graph of $\mathcal{T}_{v_{(2,2,1)}}$.

Proof. We will use Construction 2.13. Since $\mu_{p,q}$ has a unique valley, it contains precisely p lattice points of y -coordinate q , and $n - p$ lattice points of y -coordinate n . Every other y -coordinate is met exactly once. Therefore, the y -coordinates in $[n] \setminus \{q, n\}$ fill only fixed positions, and all the other positions in the bracket vector of $\mu_{p,q}$ are filled with either q or n . Now, since we reach y -coordinate q strictly before we reach y -coordinate n , we may first insert p -times the value q , and then fill the remaining available positions with values n (possibly before the first q and certainly after the last).

Let (p', q') be the next lattice point on $\mu_{p,q}$ with $\text{horiz}_{\bar{v}_\alpha}((p', q')) = \text{horiz}_{\bar{v}_\alpha}((p, q))$. If $q = t_{k+1} - 1$, then (p', q') is $(n - \text{horiz}_{\bar{v}_\alpha}((p, q)), n)$, and if $q < t_{k+1} - 1$, then $(p', q') = (p, q + 1)$. The description of the reduced bracket vector of $\mu_{p,q}^*$ follows immediately. \square

Lemma 2.15. Let $\mu_{p_1, q_1}, \mu_{p_2, q_2} \in \mathcal{M}(\mathcal{T}_{\bar{v}_\alpha})$, where $t_{k_1} \leq q_1 < t_{k_1+1}$ and $t_{k_2} \leq q_2 < t_{k_2+1}$ for some $k_1, k_2 \in [r]$. We have $\mu_{p_2, q_2} \geq \mu_{p_2, q_2}^* \wedge \mu_{p_1, q_1}$ if and only if $t_{k_1} - p_1 \leq t_{k_2} - p_2 < t_{k_1}$ and $q_1 \leq q_2$.

Proof. From Lemma 2.14 it follows that

$$\begin{aligned} \underline{\mathbf{b}}(\mu_{p_1, q_1}) &= \left(n^{(t_{k_1} - p_1)}, q_1^{(p_1)}, n^{(n - t_{k_1})} \right), \\ \underline{\mathbf{b}}(\mu_{p_2, q_2}) &= \left(n^{(t_{k_2} - p_2)}, q_2^{(p_2)}, n^{(n - t_{k_2})} \right), \\ \underline{\mathbf{b}}(\mu_{p_2, q_2}^*) &= \left(n^{(t_{k_2} - p_2)}, q_2', q_2^{(p_2 - 1)}, n^{(n - t_{k_2})} \right), \end{aligned}$$

where $q_2' > q_2$. For $i \in [n]$ let a_i denote the i -th entry in $\underline{\mathbf{b}}(\mu_{p_1, q_1})$, let b_i denote the i -th entry in $\underline{\mathbf{b}}(\mu_{p_2, q_2})$, let c_i denote the i -th entry in $\underline{\mathbf{b}}(\mu_{p_2, q_2}^*)$.

In order to determine when $\mu_{p_2, q_2} \geq \mu_{p_2, q_2}^* \wedge \mu_{p_1, q_1}$ is satisfied, we have to figure out under which conditions

$$(9) \quad b_i \geq \min(c_i, a_i)$$

holds for all $i \in [n]$. By Lemma 2.14 it follows that $b_i = c_i$ for $i \neq t_{k_2} - p_2 + 1$ and $b_{t_{k_2} - p_2 + 1} < c_{t_{k_2} - p_2 + 1}$.

Hence, (9) holds trivially if $i \neq t_{k_2} - p_2 + 1$. It thus remains to consider the case $i = t_{k_2} - p_2 + 1$. We have $c_i = q_2' > q_2$, which implies that $\min(a_i, c_i) \leq b_i$ holds if and only if $a_i \leq b_i$. In particular, we need to have $a_i < n$, which is the case precisely when $t_{k_1} - p_1 < i \leq t_{k_1}$. \square

We may thus conclude the following description of the Galois graph of \mathcal{T}_{v_α} , which is illustrated in Figure 9 for $\alpha = (2, 2, 1)$.

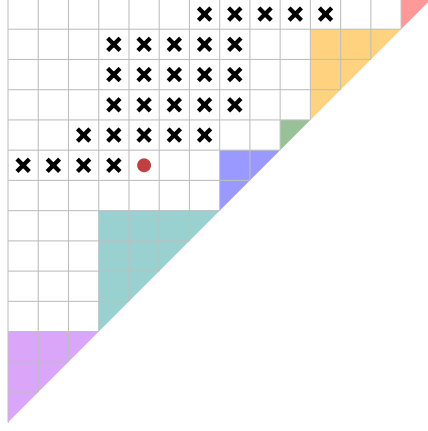


Figure 10. In the Galois graph of the ν -Tamari lattice given by the above bounce path, there exist a directed arrow from the box with a red dot to any of the boxes with a cross. (We identify boxes with the lattice points on their bottom right corner.)

Theorem 2.16. *Let α be an integer composition into r parts, and let ν_α be the corresponding α -bounce path. The Galois graph of \mathcal{T}_{ν_α} is isomorphic to the directed graph with vertex set*

$$\{(p, q) \mid 1 \leq p \leq t_k \leq q < t_{k+1} \text{ for some } k \in [r-1]\},$$

where there is a directed edge from $(p_1, q_1) \rightarrow (p_2, q_2)$ if and only if $(p_1, q_1) \neq (p_2, q_2)$ and $t_{k_1} - p_1 \leq t_{k_2} - p_2 < t_{k_1}$ and $q_1 \leq q_2$, where k_1 and k_2 are the unique indices in $[r]$ with $t_{k_1} \leq q_1 < t_{k_1+1}$ and $t_{k_2} \leq q_2 < t_{k_2+1}$.

Proof. We have seen in Theorem 2.9 that \mathcal{T}_{ν_α} is extremal and congruence uniform, and therefore, by Proposition A.4, its Galois graph is (isomorphic to) the directed graph on the vertex set $\mathcal{J}(\mathcal{T}_{\nu_\alpha})$, which has a directed edge $x \rightarrow y$ if and only if $x \neq y$ and $y \leq y_* \vee x$. Now, Theorem 2.10 implies that $\mathcal{T}_{\nu_\alpha} \cong \mathcal{T}_{\bar{\nu}_\alpha}$, and let us denote this isomorphism by Ψ . It follows that $\mu \in \mathcal{J}(\mathcal{T}_{\nu_\alpha})$ if and only if $\Psi(\mu) \in \mathcal{M}(\mathcal{T}_{\bar{\nu}_\alpha})$, and for $\mu_1, \mu_2 \in \mathcal{J}(\mathcal{T}_{\nu_\alpha})$ we have $\mu_2 \leq (\mu_2)_* \vee \mu_1$ if and only if $\Psi(\mu_2) \geq \Psi(\mu_2)_* \wedge \Psi(\mu_1)$. We may therefore use Lemma 2.15 to describe $\text{Gal}(\mathcal{T}_{\nu_\alpha})$.

Let V denote the set of pairs (p, q) with $1 \leq p \leq t_k \leq q < t_{k+1}$ for some $k \in [r-1]$. Let G be the directed graph with vertex set V , which has a directed edge precisely under the conditions given in the statement.

It is straightforward to verify that the elements of V are precisely the lattice points in the rectangle from $(0, 0)$ to (n, n) that do not have a coordinate equal to 0 or n , and that lie weakly above $\bar{\nu}_\alpha$. Therefore, for each such pair (p, q) we can find a unique $\bar{\nu}_\alpha$ -path that has a unique valley at (p, q) . Hence, the vertices of G are in bijection with the meet-irreducible elements of $\mathcal{T}_{\bar{\nu}_\alpha}$ via the map $(p, q) \mapsto \mu_{p,q}$.

By Lemma 2.15 we see that there exists a directed edge $(p_1, q_1) \rightarrow (p_2, q_2)$ in G if and only if $\mu_{p_1, q_1} \neq \mu_{p_2, q_2}$ and $\mu_{p_2, q_2} \geq \mu_{p_2, q_2}^* \wedge \mu_{p_1, q_1}$. In view of the first paragraph of this proof, this is exactly the case when there exists a directed edge $\Psi^{-1}(\mu_{p_1, q_1}) \rightarrow \Psi^{-1}(\mu_{p_2, q_2})$ in $\text{Gal}(\mathcal{T}_{\nu_\alpha})$. We conclude that G and $\text{Gal}(\mathcal{T}_{\nu_\alpha})$ are isomorphic. \square

We may identify the lattice points (p, q) above v_α with the box that has (p, q) as its lower right corner. From this perspective, Figure 10 illustrates Theorem 2.16, by indicating to which boxes an arrow from the highlighted box may go in the Galois graph.

Corollary 2.17. *Let α be an integer composition. In the Galois graph of \mathcal{T}_{v_α} the vertex (p, q) has precisely $p(|\alpha| - q) - 1$ outgoing arrows.*

Proof. This follows immediately from Theorem 2.16. \square

2.3. \mathcal{T}_α AND \mathcal{T}_{v_α} ARE ISOMORPHIC

In this section we prove that for every integer composition α the lattices \mathcal{T}_{v_α} and \mathcal{T}_α are isomorphic. In order to achieve this we make use of the bijective correspondence between the sets $\mathfrak{S}_\alpha(231)$ and \mathcal{D}_α that was first described in [26, Theorem 1.2], and recalled here in Sections 1.2.1 and 1.2.2.

Let Θ_1 and Θ_2 be the bijections from Theorems 1.8 and 1.13. Consequently, the map

$$\Theta: \mathcal{D}_{\bar{\alpha}} \rightarrow \mathfrak{S}_\alpha(231), \quad \mu \mapsto \Theta_2 \circ \Theta_1(\mu)$$

is a bijection.

Remark 2.18. Note that there is a slight subtlety to our approach. The bijection Θ maps elements of $\mathcal{D}_{\bar{\alpha}}$ to elements of $\mathfrak{S}_\alpha(231)$, and therefore *reverses* the composition. In fact, this comes in handy, because Theorems 2.7 and 2.10 state that reversing the composition corresponds to taking the lattice dual. Theorem 2.3 describes the Galois graph of $\mathcal{T}_{\bar{\alpha}}$, while Theorem 2.16 uses the dual perspective to describe the Galois graph of \mathcal{T}_{v_α} . Therefore, everything is in place for the map Θ to take effect. This is also the reason why we did not use the map $\Xi_{\text{perm}} \circ \Lambda_{\text{nn}}$ in this section.

We now show that the map Θ extends to an isomorphism from $\text{Gal}(\mathcal{T}_{v_\alpha})$ to $\text{Gal}(\mathcal{T}_\alpha)$. Let us write $w_{a,b}$ for the unique join-irreducible element of \mathcal{T}_α whose unique descent is (a, b) , and recall that $\mu_{p,q}$ is the unique meet-irreducible element of \mathcal{T}_{v_α} that has its unique valley at (p, q) .

Lemma 2.19. *Let $\mu_{p,q} \in \mathcal{M}(\mathcal{T}_{v_\alpha})$ such that $t_k \leq q < t_{k+1}$. The image of $\Theta(\mu_{p,q})$ is $w_{a,b} \in \mathcal{J}(\mathcal{T}_\alpha)$, where*

$$\begin{aligned} a &= s_{r-k-1} + q - n + s_{r-k} + 1, \\ b &= s_{r-k} + p. \end{aligned}$$

In particular, a belongs to the $(r-k)$ -th α -region, and b belongs to the i -th α -region for some $i > r-k$.

Proof. Since $\mu_{p,q}$ has a unique valley at (p, q) , it follows from Construction 1.7 that $\Theta_1(\mu_{p,q})$ has a unique bump at positions a, b with $a = s_{r-k-1} + \ell + 1$, where $q = t_k + \ell$ for a unique index $k \in [r]$, and $b = s_{r-k} + p$. (The value of b follows immediately, since there is no other bump in $\Theta_1(\mu_{p,q})$ that could potentially block elements.) By Construction 1.11 we get that $w = \Theta(\mu_{p,q})$ has (a, b) as its unique inversion, which proves the first claim.

By construction we have $s_{r-k} = n - t_k$, and since $t_k \leq q < t_{k+1}$ we conclude that $s_{r-k-1} < a \leq s_{r-k}$, meaning that a belongs to the $(r-k)$ -th α -region. Moreover, since $p \geq 1$, we conclude that $b > s_{r-k}$, which implies the last statement. \square

Theorem 2.20. *Let α be an integer composition. The map Θ induces an isomorphism from $\text{Gal}(\mathcal{T}_{V_\alpha})$ to $\text{Gal}(\mathcal{T}_\alpha)$.*

Proof. Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_r)$ with $n = |\alpha|$. We denote by G_{V_α} the Galois graph of \mathcal{T}_{V_α} , and by G_α that of \mathcal{T}_α . It follows from Lemma 2.19 and the fact that Θ is a bijection which sends descents to valleys, that the vertex sets of G_{V_α} and G_α are in bijection via Θ .

Now let $(p_1, q_1), (p_2, q_2)$ be two vertices of G_{V_α} such that $t_{k_1} \leq q_1 < t_{k_1+1}$ and $t_{k_2} \leq q_2 < t_{k_2+1}$ for some $k_1, k_2 \in [r]$.

By Lemma 2.19 we have $\Theta(\mu_{p_1, q_1}) = w_{a_1, b_1}$ and $\Theta(\mu_{p_2, q_2}) = w_{a_2, b_2}$, where

$$\begin{aligned} a_1 &= s_{r-k_1-1} + s_{r-k_1} + q_1 + 1 - n, & b_1 &= s_{r-k_1} + p_1, \\ a_2 &= s_{r-k_2-1} + s_{r-k_2} + q_2 + 1 - n, & b_2 &= s_{r-k_2} + p_2. \end{aligned}$$

Moreover, if we assume that $(p_1, q_1) \neq (p_2, q_2)$, then we see that $(a_1, b_1) \neq (a_2, b_2)$.

Suppose that there is a directed edge $(p_1, q_1) \rightarrow (p_2, q_2)$ in G_{V_α} . Theorem 2.16 then implies that $t_{k_1} - p_1 \leq t_{k_2} - p_2 < t_{k_2}$ and $q_1 \leq q_2$. Since $q_1 \leq q_2$, we conclude that $t_{k_1} \leq t_{k_2}$. Therefore, from the definition of t_k , we have two cases, $k_1 = k_2$ or $k_1 < k_2$.

(i) If $k_1 = k_2$, then it follows that a_1 and a_2 belong to the same α -region, and we must have $a_1 \leq a_2$. From the fact that $t_{k_1} - p_1 \leq t_{k_2} - p_2$ follows $p_2 \leq p_1$, and therefore $b_2 \leq b_1$. Lemma 2.19 implies further that $a_2 < b_2$; it follows now from Theorem 2.3 that there exists a directed edge $(a_1, b_1) \rightarrow (a_2, b_2)$ in G_α .

(ii) If $k_1 < k_2$, then Lemma 2.19 implies that a_1 belongs to the $(r - k_1)$ -th α -region, and a_2 belongs to the $(r - k_2)$ -th α -region. It is clear that $r - k_2 < r - k_1$, which implies $a_2 < a_1$. Moreover, we have

$$\begin{aligned} t_{k_1} - p_1 &= n - s_{r-k_1} - p_1 = n - (s_{r-k_1} + p_1) = n - b_1, \\ t_{k_2} - p_2 &= n - s_{r-k_2} - p_2 = n - (s_{r-k_2} + p_2) = n - b_2. \end{aligned}$$

By assumption we have $t_{k_1} - p_1 \leq t_{k_2} - p_2$, which implies $b_2 \leq b_1$. Finally, $t_{k_2} - p_2 < t_{k_1}$ implies $b_2 > s_{r-k_1}$. Therefore, Theorem 2.3 tells us that there is a directed edge $(a_1, b_1) \rightarrow (a_2, b_2)$ in G_α .

Conversely, let $(a_1, b_1), (a_2, b_2)$ be two vertices of G_α such that a_1 belongs to the $(r - k_1)$ -th α -region and a_2 belongs to the $(r - k_2)$ -th α -region.

By Lemma 2.19 we have $\Theta^{-1}(w_{a_1, b_1}) = \mu_{p_1, q_1}$ and $\Theta^{-1}(w_{a_2, b_2}) = \mu_{p_2, q_2}$, where

$$\begin{aligned} p_1 &= b_1 - s_{r-k_1}, & q_1 &= n + a_1 - s_{r-k_1-1} - s_{r-k_1} - 1, \\ p_2 &= b_2 - s_{r-k_2}, & q_2 &= n + a_2 - s_{r-k_2-1} - s_{r-k_2} - 1, \end{aligned}$$

and $t_{k_1} \leq q_1 < t_{k_1+1}$ and $t_{k_2} \leq q_2 < t_{k_2+1}$. Moreover, if we assume that $(a_1, b_1) \neq (a_2, b_2)$, then we see that $(p_1, q_1) \neq (p_2, q_2)$.

Suppose that there is a directed edge $(a_1, b_1) \rightarrow (a_2, b_2)$ in G_α . Theorem 2.3 implies first of all that $r - k_2 \leq r - k_1$.

(i) If $k_1 = k_2$, then Theorem 2.3 implies also that $a_1 \leq a_2 < b_2 \leq b_1$. It follows immediately that $q_1 \leq q_2$, and $t_{k_1} - p_1 \leq t_{k_1} - p_2$. Moreover, since $p_2 > 0$, we have $t_{k_1} - p_2 < t_{k_1}$. Theorem 2.16 states that there exists a directed edge $(p_1, q_1) \rightarrow (p_2, q_2)$ in G_{V_α} .

(ii) If $r - k_2 < r - k_1$, then Theorem 2.3 implies further that $s_{r-k_1} < b_2 \leq b_1$. This means exactly that $s_{r-k_1} < p_2 + s_{r-k_2} \leq p_1 + s_{r-k_1}$. It follows that $n - s_{r-k_1} > n - s_{r-k_2} - p_2 \geq n - s_{r-k_1} - p_1$, which means precisely that $t_{k_1} > t_{k_2} - p_2 \geq t_{k_1} - p_1$.

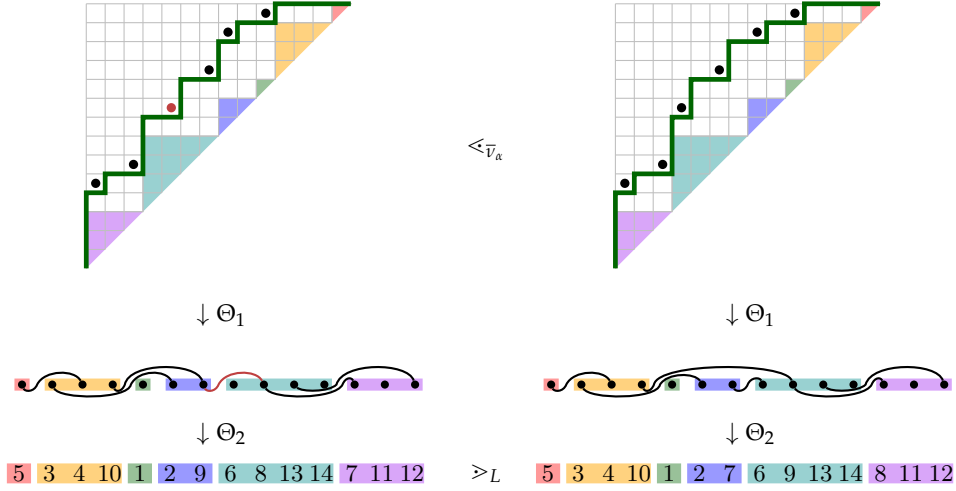


Figure 11. A cover relation in \mathcal{T}_α and the corresponding cover relation in $\mathcal{T}_{\bar{v}_\alpha}$ for $\alpha = (1, 3, 1, 2, 4, 3)$.

Moreover, $s_{r-k_1} > s_{r-k_2}$ implies $t_{k_1} < t_{k_2}$, and we conclude that $q_1 < t_{k_1+1} \leq t_{k_2} \leq q_2$. Theorem 2.16 now states that there exists a directed edge $(p_1, q_1) \rightarrow (p_2, q_2)$ in G_{v_α} .

We have thus shown that Θ induces an isomorphism of directed graphs from G_{v_α} to G_α , and the proof is complete. \square

We may now prove Theorem II, which states that \mathcal{T}_{v_α} and \mathcal{T}_α are isomorphic.

Theorem II. *For every integer composition α , the parabolic Tamari lattice \mathcal{T}_α is isomorphic to the v_α -Tamari lattice \mathcal{T}_{v_α} .*

Proof. Theorems 2.2 and 2.9 state that the lattices \mathcal{T}_α and \mathcal{T}_{v_α} are both extremal lattices, and Theorem 2.20 states that their Galois graphs are isomorphic. We may thus conclude the result as described in Remark A.2. \square

Corollary 2.21. *For every integer composition α , the lattice $\mathcal{T}_{\bar{v}_\alpha}$ is isomorphic to the dual of \mathcal{T}_α .*

Proof. This follows from Theorems 2.10 and II. \square

Remark 2.22. We may as well prove Theorem 2.7 by combining Theorems II and 2.10.

In fact, we suspect that the map Θ is an isomorphism from $\mathcal{T}_{\bar{v}_\alpha}$ to \mathcal{T}_α , but we have not succeeded proving this directly. See also Remark A.3. Figure 11 shows a cover relation in $\mathcal{T}_{v_{(3,4,2,1,3,1)}}$ and the corresponding cover relation in $\mathcal{T}_{(1,3,1,2,4,3)}$ under the map Θ .

Open Problem 2.23. *Show that Θ is a lattice isomorphism from $\mathcal{T}_{\bar{v}_\alpha}$ to \mathcal{T}_α .*

Part 3. The Zeta Map in Parabolic Cataland

In the last part of this article, we prove the Steep-Bounce Conjecture of Bergeron, Ceballos and Pilaud; see [5, Conjecture 2.2.8]. This conjecture arises as a combinatorial approach to understand an intriguing connection between the study of certain lattice walks in the positive quarter plane, and certain Hopf algebras with applications in the theory of multivariate diagonal harmonics.

We prove this conjecture by relating two families of nested Dyck paths via the left-aligned colorable trees (α -trees) from Section 1.1.4. Since a tree T can be an α -tree for potentially several choices of α , it is convenient to consider the following definition.

Definition 3.1. A *LAC tree* of size n is a pair (T, α) , where T is an α -tree for a composition α of n ($\alpha \models n$). For $n > 0$, we denote by

$$LAC_n \stackrel{\text{def}}{=} \{(T, \alpha) \mid \alpha \models n, T \in \mathbb{T}_\alpha\}$$

the set of all LAC trees of size n .

In contrast to the bijections defined in Section 1.2, where a composition α is fixed, the bijections we define in this section will consider all compositions of n at the same time.

3.1. LEVEL-MARKED DYCK PATHS

In view of Definition 1.3, a *Dyck path* is simply a $(1, 1, \dots, 1)$ -Dyck path. Let \mathcal{D}_n denote the set of all Dyck paths with $2n$ steps.

Definition 3.2. A *marked Dyck path* μ^\bullet is a Dyck path with two types of north-steps: *marked* ones denoted by N_\bullet , and *unmarked* ones denoted by N_\circ . It is *level-marked* if, for each lattice point \vec{p} on μ , the number of unmarked north-steps before \vec{p} does not exceed the number of east-steps before \vec{p} .

Remark 3.3. As an equivalent definition, μ^\bullet is level-marked if and only if for each lattice point (p, q) on μ^\bullet , there are at least $q - p$ marked north-steps before it.

The set of all level-marked Dyck paths is denoted by \mathcal{D}_n^\bullet . Level-marked Dyck paths have been introduced in [5, Section 2.2.3] under the name “colored Dyck paths”, and it was shown in that paper that \mathcal{D}_n^\bullet is in bijection with the set of lattice walks in $\mathbb{N} \times \mathbb{N}$ with $2n$ steps taken from the set $\{(0, 1), (-1, 1), (1, -1)\}$ that start at the origin and end on the x -axis. The bijection goes by sending N_\bullet to $(0, 1)$, N_\circ to $(-1, 1)$ and E to $(1, -1)$. An example of a level-marked Dyck path is shown in the middle of Figure 12.

We define the *right-to-left traversal* of a plane rooted tree T to be the depth-first search in T starting from the root, where children of the same node are visited from right to left.

The following construction describes how to obtain a marked Dyck path from an α -tree with $n = |\alpha|$.

Construction 3.4. Given an integer composition α , let $T \in \mathbb{T}_\alpha$ and C its corresponding left-aligned coloring. We construct a marked Dyck path μ_T^\bullet , starting from the origin, by performing a right-to-left traversal of T . Whenever we visit a




Figure 12. Example of left-aligned colorable trees, level-marked Dyck paths and steep pairs which are in bijection.

new node in T with a C -color we have not seen before, we add a marked north-step N_\bullet , otherwise we add an unmarked north-step N_\circ . When we visit a node that we have visited before, we add an east-step.

The left part of Figure 12 illustrates Construction 3.4. In the α -tree on the left of that figure, the first vertices per color in the right-to-left traversal are marked.

We now prove that the following map is well defined with the correct domain and image:

$$(10) \quad \Xi_{\text{dyck}}: \text{LAC}_n \rightarrow \mathcal{D}_n^\bullet, \quad (T, \alpha) \mapsto \mu_T^\bullet.$$

Proposition 3.5. *For a LAC tree (T, α) , the marked Dyck path μ_T^\bullet obtained by Construction 3.4 is indeed a level-marked Dyck path with $2|\alpha|$ steps.*

Proof. Since Construction 3.4 is simply a variant of the classical bijection between rooted plane trees and Dyck paths, it is clear that μ_T^\bullet is a Dyck path. To see that it is level-marked, we only need to consider the endpoint (p, q) of any north step in μ_T^\bullet which corresponds to the first visit of a node u in T . By Lemma 1.6(i), all nodes on the path from u to the root, including u and excluding the root, are of different colors. There are exactly $q - p$ such nodes, meaning that up to (p, q) we have seen at least $q - p$ colors in the traversal of Construction 3.4. Hence, at least $q - p$ marked north steps come before the point (p, q) , making μ_T^\bullet a level-marked Dyck path. \square

We now explain how to obtain a colored tree from a level-marked Dyck path.

Construction 3.6. For $n > 0$ and $\mu^\bullet \in \mathcal{D}_n^\bullet$, suppose that μ^\bullet has precisely r marked north-steps. We initialize our construction with the tree T_0 that consists of a single root node, the empty coloring f_0 , and a list of colors $L_0 = [0]$, where 0 stands for the color of the root. Throughout this construction we maintain a pointer on the colors. Initially, this pointer points to the 0 in L_0 . Moreover, since we traverse the tree during its construction, we say that the *current node* is the node that we are visiting in the current step, which is the root at the beginning.

We parse μ^\bullet as a word. Let x be the i -th letter of this word. We then update $T_{i-1}, f_{i-1}, L_{i-1}$ as follows.

- (i) If $x = N_\bullet$, then we extend T_{i-1} to T_i by adding a node v_i to the current node of T_{i-1} as its leftmost child. We set the current node of T_i to be v_i . We insert a new color c into L_{i-1} right after the pointer, and move the pointer to c . We extend f_{i-1} to f_i by setting $f_i(v_i) = c$.

- (ii) If $x = N_o$, then we extend T_{i-1} to T_i by adding a child v_i to the current node of T_{i-1} as its leftmost child. We set the current node of T_i to be v_i . We take L_i to simply be L_{i-1} with the pointer moved to the next color c' . We extend f_{i-1} to f_i by setting $f_i(v_i) = c'$.
- (iii) If $x = E$, then we take T_i to simply be T_{i-1} with the current node set to the parent of the current node of T_{i-1} . We take L_i to be L_{i-1} with the pointer moved to the previous color. Since μ^\bullet is a Dyck path, such a color (may be 0 for the root) must exist. We set $f_i = f_{i-1}$.

Let f'_i be the coloring of T_i that is obtained from f_i by renaming the colors according to first appearance in the left-to-right traversal of T_i . After $2n$ steps we have parsed μ^\bullet completely, and we output the colored tree $\mathcal{T}_{\mu^\bullet} = (T_{2n}, f'_{2n})$.

If $\mathcal{T}_{\mu^\bullet} = (T_{\mu^\bullet}, f')$ is the output of Construction 3.6, then we set α_i to be the number of nodes of T_{μ^\bullet} whose f' -color is i , and we set $\alpha_{\mu^\bullet} = (\alpha_1, \alpha_2, \dots, \alpha_r)$ for r the number of colors used by f' . We now prove that the following map is well defined:

$$(11) \quad \Lambda_{\text{dyck}}: \mathcal{D}_n^\bullet \rightarrow \text{LAC}_n, \quad \mu^\bullet \mapsto (T_{\mu^\bullet}, \alpha_{\mu^\bullet}).$$

Lemma 3.7. *At the end of the i -th step of Construction 3.6, the list L_i of colors (renamed by first appearances in left-to-right traversal) is increasing. In particular, if u is a non-root node of T_i and v is a descendant of u , then $f'_i(u) < f'_i(v)$.*

Proof. It follows from Construction 3.6 that the colors in the unique shortest path from the root of T_i to v exactly coincide with the colors in L_i from the beginning to the pointer, and they appear in the same order on that path and in L_i .

Now, suppose that there are two colors c_1, c_2 such that c_1 is after c_2 in L_i . Let u_1 be the first node in the LR-prefix order of T_i with $f'_i(u_1) = c_1$. There must be an ancestor u_2 of u_1 with $f'_i(u_2) = c_2$ by the argument in the first paragraph.

Since u_2 precedes u_1 in the LR-prefix order of T_i , it follows that $c_2 < c_1$ due to the renaming. This yields $f'_i(u_2) < f'_i(u_1)$ as desired. \square

Proposition 3.8. *For a level-marked Dyck path $\mu^\bullet \in \mathcal{D}_n^\bullet$ with r marked north-steps, let (T, f) be the colored tree obtained from Construction 3.6. The coloring f yields a composition α such that T is compatible with α .*

Proof. We suppose that μ^\bullet has r marked north-steps. Let $\mathcal{T}_{\mu^\bullet} = (T, f)$ be the colored tree obtained from μ^\bullet through Construction 3.6. It follows that f uses exactly r colors, and for $i \in [r]$ we let α_i be the number of nodes of T whose f -color is i . Since every non-root node of T receives a unique color, it is clear that $(\alpha_1, \alpha_2, \dots, \alpha_r)$ is a composition of n .

For $s \in [r]$, let T_s be the restriction of T to the nodes of f -color at most s . We prove by induction on s that the induced subtree T'_s of T whose nodes are colored at the end of step s in Construction 1.4 is precisely T_s .

The base case $s = 0$ is clear, since both T_0 and T'_0 consist of a single root node. Now suppose that $T_s = T'_s$. By Lemma 1.6(i), we see that T'_{s+1} is connected. By Lemma 3.7, the f -color of any node is bigger than those of all its ancestors, so T_{s+1} is also connected.

Now, let u be the last node with $f(u) = s + 1$ in the LR-prefix order of T . Consider a node $v \notin T'_s$ that is a child of some node in $T_s = T'_s$ while strictly preceding u in LR-prefix order. Suppose that v is created in step i of Construction 3.6, with

L_i the list of colors (after renaming) available just after. Note that every ancestor of v has f -color at most s , so that u is in particular not an ancestor of v . Therefore, v comes after u in the right-to-left traversal of T , which means that $s + 1$ is already present in L_i . The pointer of L_i cannot be after $s + 1$, because that would imply the existence of an ancestor of v of f -color $s + 1$, which is a contradiction. The pointer of L_i cannot be before $s + 1$, because otherwise by Lemma 3.7 we would have the contradiction $s + 1 > f(v) > s$ because $v \notin T_s$. It follows that $f(v) = s + 1$, which holds for all nodes of $T'_{s+1} \setminus T'_s$ that come before u in the LR-prefix order.

Since u is the last node in LR-prefix order that has f -color $s + 1$, no descendant of u belongs to T_{s+1} . Let X be the set of nodes in $T_{s+1} \setminus T_s$. We have just shown that X contains all the nodes of $T'_{s+1} \setminus T'_s$ that precede u in LR-prefix order. Moreover, all nodes in X are active in the $(s + 1)$ -st step of Construction 1.4. Since $|X| = \alpha_{s+1}$, all the vertices in X receive color $s + 1$ in Construction 1.4. We conclude that $T_{s+1} = T'_{s+1}$.

This completes the induction step, and we see that T is indeed compatible with α , and f is indeed the left-aligned coloring of T . \square

We conclude this subsection with the proof that Ξ_{dyck} is a bijection whose inverse is Λ_{dyck} .

Theorem 3.9. *For every positive integer n , the map Ξ_{dyck} is a bijection whose inverse is Λ_{dyck} .*

Proof. In view of Propositions 3.5 and 3.8 it remains to show that $\Xi_{\text{dyck}} \circ \Lambda_{\text{dyck}} = \text{id}$ and $\Lambda_{\text{dyck}} \circ \Xi_{\text{dyck}} = \text{id}$. We observe that the map Ξ_{dyck} without markings and colors is a small twist of the well-known bijection from the set of plane rooted trees with n non-root nodes to the set of Dyck paths with $2n$ steps, described for instance in [11, Appendix E.1].

To show that $\Xi_{\text{dyck}} \circ \Lambda_{\text{dyck}} = \text{id}$, for $\mu^\bullet \in \mathcal{D}_n^\bullet$, let $(T, \alpha) = \Lambda_{\text{dyck}}(\mu^\bullet)$ and $\hat{\mu}^\bullet = \Xi_{\text{dyck}}((T, \alpha))$. As mentioned before, the underlying Dyck paths μ and $\hat{\mu}$ are equal, so it remains to show that the marked north-steps are in the same positions. By construction, the marked north-steps of μ^\bullet determine the right-most nodes per color in (T, α) , which in turn determine the marked north-steps in $\hat{\mu}^\bullet$ in the same way. We conclude that $\mu^\bullet = \hat{\mu}^\bullet$.

To show that $\Lambda_{\text{dyck}} \circ \Xi_{\text{dyck}} = \text{id}$, for $(T, \alpha) \in \text{LAC}_n$, let $\mu^\bullet = \Xi_{\text{dyck}}((T, \alpha))$ and $(T', \alpha') = \Lambda_{\text{dyck}}(\mu^\bullet)$. As before, we conclude that $T = T'$, and it remains to show that $\alpha = \alpha'$. We observe already that α has the same number of parts as α' , as both are equal to the number of marked north steps in μ^\bullet by construction.

Let T_s, T'_s denote the subtrees of T and T' , respectively, that consist of all nodes of color at most s , including the root. We show that $T_s = T'_s$ by induction on s . The base case $s = 0$ is clear, since both T_0 and T'_0 consist of the root only. Now suppose that $T_s = T'_s$. At least one of the active nodes of T_s must belong to the set V of last nodes per color with respect to (T, α) . We denote by u the first active node in V in the LR-prefix order of T . It follows that all active nodes of T_s up to u in LR-prefix order receive color $s + 1$. The same reasoning works for T'_s as α and α' have the same number of parts, and we conclude that $T_{s+1} = T'_{s+1}$. This completes the induction step, and we have $\alpha = \alpha'$. \square

3.2. THE STEEP-BOUNCE THEOREM

Recall the definitions of steep and bounce Dyck paths from Section 1.1.3. A pair (μ_1, μ_2) of Dyck paths of the same length is *nested* if μ_1 always stays weakly below μ_2 . We denote by \mathcal{SP}_n the set of nested pairs (μ_1, μ_2) such that $\mu_1, \mu_2 \in \mathcal{D}_n$ and the top path μ_2 is steep. We refer to the elements of \mathcal{SP}_n as *steep pairs* for simplicity. Similarly, we denote by \mathcal{BP}_n the set of nested pairs (μ_1, μ_2) such that $\mu_1, \mu_2 \in \mathcal{D}_n$ and the bottom path μ_1 is bounce, and refer to its elements as *bounce pairs*.

The following construction encodes a level-marked Dyck path as a steep pair.

Construction 3.10 ([5, Section 2.2.3]). Given $n > 0$ and $\mu^\bullet \in \mathcal{D}_n^\bullet$, we first take $\mu_1 \in \mathcal{D}_n$ the path obtained from μ^\bullet by forgetting the marking. We then construct a Dyck path $\mu_2 \in \mathcal{D}_n$ by forgetting the east-steps, replacing each marked north-step N_\bullet by a north-step N , and each unmarked north-step N_\circ by a valley EN . We then append sufficiently many east-steps at the end of the so-created path to reach the coordinate (n, n) . We return the pair (μ_1, μ_2) .

Lemma 3.11. *For a level-marked Dyck path μ^\bullet , the pair (μ_1, μ_2) obtained by Construction 3.10 is a nested pair of Dyck paths, where μ_2 is steep.*

Proof. The fact that μ_2 is steep and is weakly above μ_1 is immediate from the construction. \square

It follows from [5, Section 2.2.3] that the map

$$(12) \quad \phi: \mathcal{D}_n^\bullet \rightarrow \mathcal{SP}_n, \quad \mu^\bullet \mapsto (\mu_1, \mu_2)$$

is a bijection. Let

$$(13) \quad \Xi_{\text{steep}}: \text{LAC}_n \rightarrow \mathcal{SP}_n, \quad (T, \alpha) \mapsto \phi \circ \Xi_{\text{dyck}}((T, \alpha))$$

be the composition of the maps Ξ_{dyck} and ϕ . Theorem 3.9 states that Ξ_{dyck} is a bijection, which implies that Ξ_{steep} is also a bijection. This map is illustrated in the right part of Figure 12, and also in the right part of Figure 13. We denote by Λ_{steep} the inverse of Ξ_{steep} .

Recall that for a fixed composition α , the map $\Xi_{\text{nn}}: \mathbb{T}_\alpha \rightarrow \mathcal{D}_\alpha$ from (7) is a bijection between the set of α -trees and the set of α -Dyck paths. This map can be naturally extended to a map on the set of LAC trees considering all compositions simultaneously as follows:

$$(14) \quad \Xi_{\text{bounce}}: \text{LAC}_n \rightarrow \mathcal{BP}_n, \quad (T, \alpha) \mapsto (\nu_\alpha, \Xi_{\text{nn}}(T)).$$

Remark 3.12. Although the maps Ξ_{nn} and Ξ_{bounce} are essentially the same, we assign them different names because they have different domains and codomains. On one hand, Ξ_{nn} is defined on the set of plane rooted trees that are compatible with a fixed α (α -trees), and its image is the set of Dyck paths that lie weakly above a fixed bounce path ν_α (α -Dyck paths). On the other hand, Ξ_{bounce} is defined on the set of pairs (T, α) where T is an α -tree and α is a composition of n (LAC trees), and its image is the set of bounce pairs of the form (ν_α, μ) .

Theorem 1.32 states that Ξ_{nn} is a bijection, which implies that Ξ_{bounce} is also a bijection. Let Λ_{bounce} be the inverse of Ξ_{bounce} . The left part of Figure 13 illustrates




Figure 13. An illustration of the Steep-Bounce zeta map Γ .

the map Ξ_{bounce} . We now define a bijection Γ by

$$(15) \quad \Gamma \stackrel{\text{def}}{=} \Xi_{\text{bounce}} \circ \Lambda_{\text{steep}}.$$

Note that Γ sends steep pairs to bounce pairs.

We are now in the position to prove the first main result of this section.

Theorem III (The Steep-Bounce Theorem). *For $n > 0$ and every $r \in [n]$, the map Γ is a bijection from*

- the set of nested pairs (μ_1, μ_2) of Dyck paths with $2n$ steps, where μ_2 is a steep path ending with exactly r east-steps with y -coordinate equal to n , to
- the set of nested pairs (μ'_1, μ'_2) of Dyck paths with $2n$ steps, where μ'_1 is a bounce path that touches the main diagonal $r + 1$ times.

Proof. Let $(\mu_1, \mu_2) \in \mathcal{SP}_n$ such that μ_2 ends with r consecutive east-steps, and $(\mu'_1, \mu'_2) = \Gamma(\mu_1, \mu_2)$. It is clear that μ'_1 touches the main diagonal precisely $r + 1$ times, since the parameter r describes precisely the number of colors in the LAC tree (T, α) that relates both pairs of Dyck paths. \square

We obtain a proof of [5, Conjecture 2.2.1] as a corollary. We refer the reader to [5] for any undefined terminology.

Corollary 3.13 ([5, Conjecture 2.2.1]). *The graded dimension of degree n of the Hopf algebra of the special family of pipe dreams considered in [5, Section 2.2] is equal to the number of walks in the quarter plane starting from the origin, ending on the x -axis, and consisting of $2n$ steps taken from the set $\{(-1, 1), (1, -1), (0, 1)\}$.*

An illustration of the bijection Γ can be found in Figure 13; the correspondence between all LAC trees, bounce pairs and steep pairs for $n = 3$ is shown in Figure 14.

3.3. A NEW INTERPRETATION OF THE ZETA MAP

Somewhat surprisingly, the map Γ turns out to generalize the classical zeta map in q, t -Catalan combinatorics. Before specifying this connection, we briefly recall the history of the zeta map. Readers familiar with this topic may skip the following section and proceed directly to Section 3.3.3.

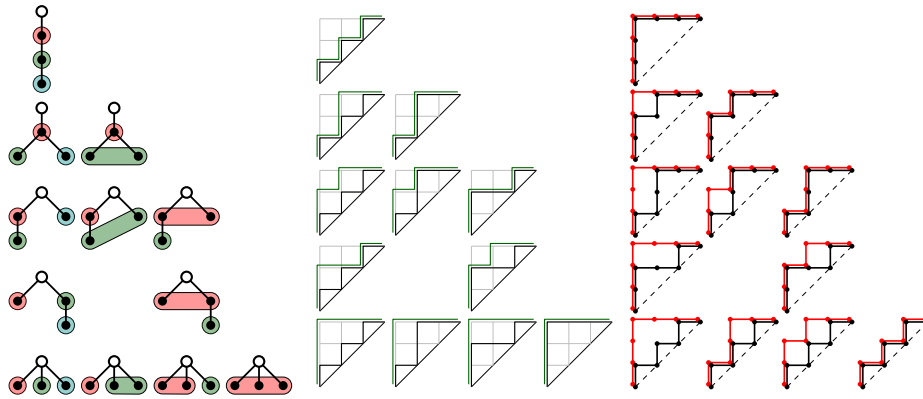


Figure 14. The LAC trees, bounce pairs, and steep pairs of size $n = 3$.

3.3.1. The Zeta Map. The classical zeta map is a certain bijection from \mathcal{D}_n onto itself. Its inverse was discovered in [1] in connection with nilpotent ideals in certain Borel subalgebras of $sl(n)$. It later became an important map in the theory of diagonal harmonics and q, t -Catalan combinatorics [17].

One remarkable ingredient in this theory is a bivariate polynomial $C_n(q, t)$ which can be defined as the bigraded Hilbert series of the alternating component of a certain module of diagonal harmonics. Surprisingly, the dimension of that space is equal to the *Catalan number* $c_n \stackrel{\text{def}}{=} \frac{1}{n+1} \binom{2n}{n}$, and therefore $C_n(q, t)$ is a polynomial with nonnegative integer coefficients whose evaluation at $q = t = 1$ recovers c_n .

Intensive attempts to understand these polynomials led to important developments [13, 14, 16, 20]. In particular, Haglund gave a combinatorial interpretation of the polynomial $C_n(q, t)$ in terms of a pair of statistics on Dyck paths known as *area* and *bounce* [16]. Shortly thereafter, Haiman announced another combinatorial interpretation in terms of two statistics *area* and *dinv*. Unexpectedly, these two pairs of statistics were different but gave rise to the same expression:

$$(16) \quad C_n(q, t) = \sum_{\mu \in \mathcal{D}_n} q^{\text{area}(\mu)} t^{\text{bounce}(\mu)} = \sum_{\mu \in \mathcal{D}_n} q^{\text{dinv}(\mu)} t^{\text{area}(\mu)}.$$

The *zeta map* ζ is a bijection on \mathcal{D}_n that explains this phenomenon by sending the pair $(\text{area}, \text{dinv})$ to the pair $(\text{bounce}, \text{area})$.

3.3.2. Some Statistics on Dyck Paths. Let us have a closer look at these statistics. First we need to describe how to obtain a steep and a bounce Dyck path from a given $\mu \in \mathcal{D}_n$.

Construction 3.14. Given $n > 0$ and $\mu \in \mathcal{D}_n$, we construct a Dyck path μ_{bounce} from μ as follows. We start at the origin and move north until we reach the last point that is still weakly below μ , and then move east until we hit the main diagonal. This process is repeated until we reach coordinate (n, n) . It is clear that $\mu_{\text{bounce}} \in \mathcal{D}_n$ is bounce, and we call it the *bounce path* of μ .

The *bounce parameters* of μ are the x -coordinates $0 = b_0 < b_1 < \dots < b_r = n$ of the contact points of μ_{bounce} and the main diagonal.

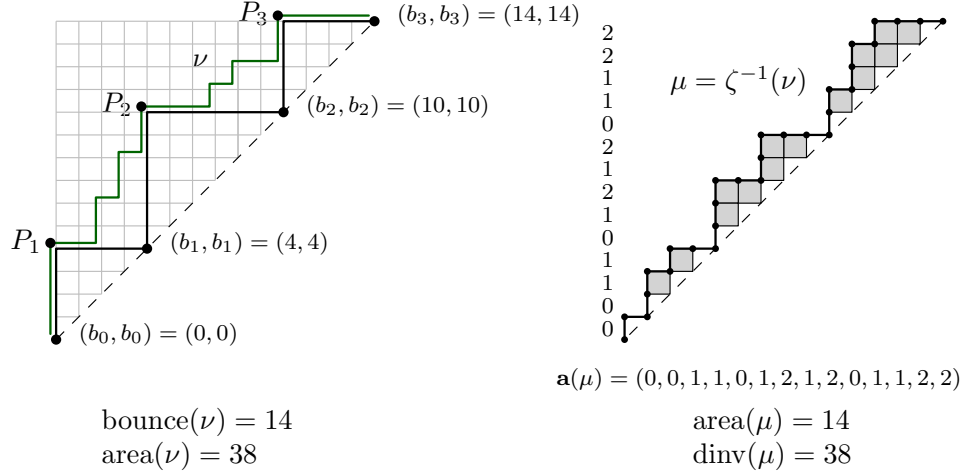


Figure 15. An example of the inverse of the zeta map.

Construction 3.15. Given $n > 0$ and $\mu \in \mathcal{D}_n$, we construct a Dyck path μ_{steep} from μ as follows. We start at the origin and move north until we reach the first point at which we can add an east-step while still remaining weakly above μ . We append this east-step, and repeat this process until we reach a lattice point with y -coordinate n . We add sufficiently many east-steps such that we end at (n, n) .

It is clear that $\mu_{\text{steep}} \in \mathcal{D}_n$ is steep, and we call it the *steep path* of μ .

Now for $\mu \in \mathcal{D}_n$, its *area vector* $\mathbf{a}(\mu) \stackrel{\text{def}}{=} (a_1, a_2, \dots, a_n)$ is the integer vector whose i -th entry a_i equals the number of unit boxes in row i that are below μ and above the main diagonal.

Definition 3.16. Given $\mu \in \mathcal{D}_n$ with area vector $\mathbf{a}(\mu) = (a_1, a_2, \dots, a_n)$ and bounce parameters b_0, b_1, \dots, b_r , we define *area*, *dinv*, and *bounce* of μ , respectively, by

$$\begin{aligned} \text{area}(\mu) &\stackrel{\text{def}}{=} a_1 + a_2 + \dots + a_n, \\ \text{dinv}(\mu) &\stackrel{\text{def}}{=} \left| \{(i, j) \mid 1 \leq i < j \leq n, a_i \in \{a_j, a_j + 1\}\} \right|, \\ \text{bounce}(\mu) &\stackrel{\text{def}}{=} \sum_{i=1}^{r-1} n - b_i. \end{aligned}$$

Figure 15 illustrates these statistics. We end this subsection by explicitly describing the inverse of the zeta map. We refer the reader to [1] and [17, Proof of Theorem 3.15] for the original sources, and to [2, 8, 34] for further information on generalizations.

Construction 3.17. Given $\nu \in \mathcal{D}_n$, let ν_{bounce} be its bounce path with bounce parameters b_0, b_1, \dots, b_r . For $k \in [r]$, we denote by $\vec{p}_k = (b_{k-1}, b_k)$ the peak of ν_{bounce} at height b_k . By construction, \vec{p}_k is also on ν . Clearly, the number of steps between \vec{p}_k and \vec{p}_{k+1} is $(b_k - b_{k-1}) + (b_{k+1} - b_k) = b_{k+1} - b_{k-1}$. For $k \in [r-1]$, the sequence $\mathbf{a}^{(k)}$ arises from the subpath of ν between \vec{p}_k and \vec{p}_{k+1} by replacing each east-step by $k-1$ and each north-step by k .




Figure 16. An illustration of Construction 3.17.

We fashion these $r - 1$ sequences together by interlacing the sequences $\mathbf{a}^{(k)}$ and $\mathbf{a}^{(k+1)}$ for all $k \in [r - 1]$, where we insert the values $k - 1$ and $k + 1$ relative to the values k , and where we never put $k - 1$ directly before $k + 1$. The resulting sequence \mathbf{a} has precisely n entries, $b_{k+1} - b_k$ of which are equal to k for $k \in \{0, 1, \dots, r - 1\}$.

This construction is illustrated in Figure 16.

Proposition 3.18 ([17, Proof of Theorem 3.15]). *For $v \in \mathcal{D}_n$, the sequence \mathbf{a} obtained by Construction 3.17 is the area-vector of some Dyck path $\mu \in \mathcal{D}_n$, and we have $\mu = \zeta^{-1}(v)$.*

3.3.3. The Steep-Bounce Zeta Map. In this section we prove Theorem IV, relating the zeta map ζ to our map Γ .

Let T be a plane rooted tree with n non-root nodes. A node of T is of *level* k if its graph distance from the root is $k + 1$. In other words, the level of a node is its depth minus one. Let $\mathbf{a}(T) = (a_1, a_2, \dots, a_n)$ be the vector obtained by reading levels of each node according to the right-to-left traversal of T . We define an integer partition $\alpha_T = (\alpha_1, \alpha_2, \dots)$ of n by setting α_i to the number of nodes of level $i + 1$. It then follows immediately that T is α_T -compatible. The pair (T, α_T) is a *horizontal LAC tree*. We have already seen a horizontal LAC tree in the middle of Figure 13.

Lemma 3.19. *For $n > 0$, the map Ξ_{steep} restricts to a bijection from*

- the set of horizontal LAC trees (T, α_T) where T has n non-root nodes, to
- the set of pairs $(\mu, \mu_{\text{steep}})$ where $\mu \in \mathcal{D}_n$.

Moreover, if $(\mu, \mu_{\text{steep}}) = \Xi_{\text{steep}}((T, \alpha_T))$, then $\mathbf{a}(\mu) = \mathbf{a}(T)$.

Proof. Recall that $\Xi_{\text{steep}} = \phi \circ \Xi_{\text{dyck}}$. Given a horizontal LAC tree (T, α) , let $\mu_{\bullet} = \Xi_{\text{dyck}}((T, \alpha))$ be the corresponding level-marked Dyck path, and $(\mu_1, \mu_2) = \phi(\mu_{\bullet})$ be the steep pair associated with μ_{\bullet} . We denote by $\mathbf{a}(\mu_1) = (a_1, a_2, \dots, a_n)$ the area vector of μ_1 . We recall that μ_1 is the Dyck path obtained from μ_{\bullet} by forgetting the markings. The proof proceeds in three steps.

We first prove that $\mathbf{a}(\mu_1) = \mathbf{a}(T)$. Recall that μ_1 is obtained from the right-to-left traversal of T : upon each visit to a new node, we add a north-step; every time we go back up one level, we add an east-step. The contribution of a north-step to the area vector is equal to the number of north-steps preceding it minus the number of east-steps preceding it; this is clearly equal to the level of the corresponding node of the tree.

We now prove that (T, α) is a horizontal LAC tree if and only if μ_\bullet is such that the j -th north-step is marked if and only if $a_i < a_j$ for all $i < j$.

Observe that for every $\mu \in \mathcal{D}_n$ there is a unique way to mark north-steps to satisfy the desired property. Moreover, every plane rooted tree T corresponds to a unique horizontal LAC-tree. Since the set of Dyck paths with $2n$ steps is in bijection with the set of plane rooted trees with n non-root nodes [11, Appendix E.1], and the map Ξ_{dyck} is a bijection, it remains to show that Ξ_{dyck} assigns to every horizontal LAC tree a marked Dyck path with the desired property.

Let (T, α) be the horizontal LAC tree derived from T . The color of a node in T is equal to its level. In the right-to-left traversal of the tree, each time we reach the first node of a certain level, we get a marked north-step N_\bullet in μ_\bullet . All other nodes correspond to regular north-steps N_\circ . Since a north-step contributes an entry to the area vector which is equal to the level of its corresponding node, the marked north-steps in μ_\bullet are exactly those that contribute an entry a_j such that $a_i < a_j$ for all $i < j$.

We finally prove that μ_2 is the steep path of μ_1 if and only if μ_\bullet is such that the j -th north-step is marked if and only if $a_i < a_j$ for all $i < j$.

This is straightforward in view of the following alternative description of the steep path of μ_1 . We parse μ_1 as follows. In the beginning, for each north-step we write down NE , until we reach one that contributes a 1 to the area vector of μ_1 , for which we write down NN . We then continue to append copies of NE until we reach a north-step that contributes a 2 to the area vector of μ_1 , for which we append NN again. We continue this process until a total of n letters N are written down, and we finish by adding the necessary number of letters E at the end.

The word so written corresponds to the steep path of μ_1 , and the valleys of $(\mu_1)_{\text{steep}}$ correspond precisely to the entries a_j of the area vector of μ_1 with $a_i < a_j$ for all $i < j$. By Construction 3.10, the valleys of μ_2 correspond precisely to the marked north-steps of μ_\bullet .

These three claims together show that (T, α) is a horizontal LAC tree if and only if $\Xi_{\text{steep}}((T, \alpha))$ is $(\mu_1, (\mu_1)_{\text{steep}})$, and that $\mathbf{a}(T) = \mathbf{a}(\mu_1)$. \square

Lemma 3.20. *For $n > 0$, the map Λ_{bounce} restricts to a bijection from*

- *the set of pairs (v_{bounce}, v) where $v \in \mathcal{D}_n$, to*
- *the set of horizontal LAC trees (T, α_T) where T has n non-root nodes.*

Moreover, if $(T, \alpha_T) = \Lambda_{\text{bounce}}(v_{\text{bounce}}, v)$, then $\mathbf{a}(\zeta^{-1}(v)) = \mathbf{a}(T)$.

Proof. Given a nested pair of Dyck paths (v_α, v) with $\alpha = (\alpha_1, \dots, \alpha_r)$ a composition of n , let $(T, \alpha) = \Lambda_{\text{bounce}}(v_\alpha, v)$. For $k \in [r]$, recall that $s_k = \alpha_1 + \alpha_2 + \dots + \alpha_k$ and that a_k is the number of active nodes at the beginning of the $(k+1)$ -st step of Construction 1.4 (not to be confused with components α_k in α). By definition, $s_0 = 0$, and a_0 is the number of children of the root. Since (T, α) is a horizontal LAC tree, we have $a_k = \alpha_{k+1}$. The proof proceeds in two steps.

We first show that (T, α) is a horizontal LAC tree if and only if $v_\alpha = v_{\text{bounce}}$.

We have already argued that the set of horizontal LAC trees with n non-root nodes is in bijection with the set of Dyck paths with $2n$ steps, and therefore also with the set of nested pairs of the form (v_{bounce}, v) with $v \in \mathcal{D}_n$. Thus, both sets have c_n elements. In order to prove the claim, it thus suffices to show that (T, α) being a horizontal LAC tree implies $v_\alpha = v_{\text{bounce}}$.

We already know that v_α lies weakly below v , and we need to show that each peak of v_α touches the path v at the initial point of an east-step of v . Equivalently, we need to show that, for each $k \in [r-1]$, the path v has a valley with coordinates (p, q) such that $s_{k-1} < p \leq s_k$ and $q = s_k$. This valley corresponds to the rightmost internal node of color k , which is the i -th node of color k for some $1 \leq i \leq \alpha_k$. Since T is a horizontal LAC tree, the rightmost child of this node is the rightmost node of color $k+1$, which is the α_{k+1} -st active node in the k -th step of Construction 1.4. By definition of Ξ_{bounce} , we see that $s_{k-1} < p = s_k - i + 1 \leq s_k$ and $q = s_k + a_k - \alpha_{k+1} = s_k$, as desired.

We now show that if (T, α) is a horizontal LAC tree, then $\mathbf{a}(\zeta^{-1}(v)) = \mathbf{a}(T)$.

Let $\mu = \zeta^{-1}(v)$ as described in Construction 3.17. The path v_{bounce} touches the diagonal at $(0, 0) = (s_0, s_0), (s_1, s_1), \dots, (s_r, s_r) = (n, n)$. Therefore, $\mathbf{a}(\mu)$ has exactly $s_{k+1} - s_k = \alpha_{k+1}$ entries equal to k for $0 \leq k \leq r-1$, which is exactly the number of entries equal to k in $\mathbf{a}(T)$.

For $1 \leq k \leq r-1$, we denote by $\mathbf{a}_k(\mu)$ and $\mathbf{a}_k(T)$, respectively, the restrictions of $\mathbf{a}(\mu)$ and $\mathbf{a}(T)$ to the values in $\{k-1, k\}$. It remains to show that $\mathbf{a}_k(\mu) = \mathbf{a}_k(T)$.

The vector $\mathbf{a}_k(T)$ can be obtained from the right-to-left traversal of T : each time we reach a new node at level $k-1$ we add a value $k-1$ and each time we reach a new node at level k we add a value k .

The vector $\mathbf{a}_k(\mu) = \mathbf{a}_k(\zeta^{-1}(v))$ can be obtained from the subpath of v between (s_{k-1}, s_k) and (s_k, s_{k+1}) by replacing each east-step with $k-1$ and each north-step by k .

Recall that $T = \Lambda_{\text{nn}}(v)$ as described in Construction 1.28. By Lemma 1.31, for $1 \leq i \leq s_k - s_{k-1} = \alpha_k$, the number of north-steps in v with x -coordinate $s_{k-1} + i$ is equal to the number of children of the i -th node of color k in the right-to-left traversal of T . From this we deduce that $\mathbf{a}_k(T) = \mathbf{a}_k(\mu)$. \square

Theorem IV. For $n > 0$, the map Γ restricts to a bijection from

- the set of pairs $(\mu, \mu_{\text{steep}})$, where μ is a Dyck path with $2n$ steps, to
- the set of pairs (v_{bounce}, v) , where v is a Dyck path with $2n$ steps.

Moreover, if $(v_{\text{bounce}}, v) = \Gamma(\mu, \mu_{\text{steep}})$, then $v = \zeta(\mu)$, where ζ is the zeta map from q, t -Catalan combinatorics.

Proof. By Lemmas 3.19 and 3.20, the map

$$\Gamma^{-1} = (\Xi_{\text{bounce}} \circ \Lambda_{\text{steep}})^{-1} = \Xi_{\text{steep}} \circ \Lambda_{\text{bounce}}$$

restricts to a bijection from

- the set of pairs (v_{bounce}, v) where $v \in \mathcal{D}_n$, to
- the set of pairs $(\mu, \mu_{\text{steep}})$ where $\mu \in \mathcal{D}_n$.

Moreover, if $(\mu, \mu_{\text{steep}}) = \Gamma^{-1}(v_{\text{bounce}}, v)$ then $\mathbf{a}(\mu) = \mathbf{a}(\zeta^{-1}(v))$, which implies that $\mu = \zeta^{-1}(v)$. \square

Remark 3.21. We can also see Equation (16) combinatorially through our bijections. Let (T, α) be a horizontal LAC tree, with $(\mu, \mu_{\text{steep}}) = \Xi_{\text{steep}}((T, \alpha))$, and $(v_\alpha, v) = \Xi_{\text{bounce}}((T, \alpha))$. Since μ is simply $\Xi_{\text{dyck}}((T, \alpha))$ without markings, we know that each north-step in μ corresponds to a non-root node in T , with the corresponding entry in the area vector $\mathbf{a}(\mu)$ to be the level of that node.

Therefore, we have (we sum over all non-root nodes)

$$\begin{aligned} \text{area}(\mu) &= \sum_{u \in T} \text{level}(u) = \sum_{u \in T} \sum_{v \text{ above } u} 1 \\ &= \sum_{v \in T} \sum_{u \text{ below } v} 1 = \sum_{i=1}^{r-1} (n - b_i) = \text{bounce}(v). \end{aligned}$$

The pairs contributing to $\text{dinv}(\mu)$ correspond to pairs of nodes (u, v) such that v comes before u in the right-to-left traversal of T and $\text{level}(v) - \text{level}(u) \in \{0, 1\}$. Now, if we order the nodes in T first by increasing level, and then by occurrence in the right-to-left traversal, then each node can be identified with an integer between 1 and n . It follows that the cells below v correspond precisely to the pairs contributing to $\text{dinv}(\mu)$, and we see that $\text{area}(v) = \text{dinv}(\mu)$.

3.3.4. The Zeta Map on Parking Functions. There is another generalization of the zeta map, due to Haglund and Loehr [19], which we can obtain as a special case of a labeled version of our bijection Γ . Before going into the details, let us introduce some definitions.

A *vertical labeling* of a Dyck path $\mu \in \mathcal{D}_n$ is an assignment of the numbers from 1 through n to the boxes directly to the right of the north steps of μ such that the numbers are increasing from below in each column. The set $\text{Vert}(\mathcal{D}_n)$ of pairs (μ, P) where P is a vertical labeling of $\mu \in \mathcal{D}_n$ is commonly referred to as the set of *parking functions* of length n .

A *diagonal labeling* of a Dyck path $v \in \mathcal{D}_n$ is an assignment of the numbers from 1 through n to the boxes on the main diagonal such that, for every consecutive pair of steps EN of v (valleys of v), the number in the same column as the E step is smaller than the number in the same row as the N step. We denote by $\text{Diag}(\mathcal{D}_n)$ the set of pairs (v, Q) where Q is a diagonal labeling of $v \in \mathcal{D}_n$.

In [19] [17, Conjecture 5.2], Haglund and Loehr conjectured two combinatorial interpretations of the bigraded Hilbert series of the module of diagonal harmonics:

$$\mathcal{DH}_n(q, t) = \sum_{(\mu, P) \in \text{Vert}(\mathcal{D}_n)} q^{\text{dinv}'(\mu, P)} t^{\text{area}(\mu)} = \sum_{(v, Q) \in \text{Diag}(\mathcal{D}_n)} q^{\text{area}'(v, Q)} t^{\text{bounce}(v)},$$

where dinv' and area' are generalizations of the dinv and area statistics that depend on the labelings P and Q . See [17, Chapter 5] for more details. Haglund and Loehr also provided a bijection

$$\tilde{\zeta}: \text{Vert}(\mathcal{D}_n) \rightarrow \text{Diag}(\mathcal{D}_n)$$

that explains the equivalence of their two interpretations, and their conjecture follows from the more general statement of the shuffle conjecture [18], which was proven by Carlsson and Mellit in [7].

The map $\tilde{\zeta}$ is a straightforward generalization of the zeta map ζ . It is defined by

$$\tilde{\zeta}(\mu, P) \stackrel{\text{def}}{=} (\zeta(\mu), Q),$$

where the order in which the labels in Q appear, in the north-east direction, is obtained by reading the labels in P along diagonals $y = x + i$, in the north-east




Figure 17. An example of the $\tilde{\zeta}$ map from parking functions to diagonally labeled Dyck paths, computed in terms of the map $\tilde{\Gamma}$.

direction, for $i \in \{0, 1, \dots, n-1\}$ [17, Chapter 5].² An example is illustrated in Figure 17.

We will show that $\tilde{\zeta}$ can be obtained in terms of a labeled version of our map Γ . Let T be a plane rooted tree with n non-root nodes. A *right increasing labeling* R of T is an assignment of the numbers from 1 through n to the non-root nodes of T that increases along right-most children, i.e. if v is the right-most child of u , then $R(v) > R(u)$. We define a *labeled horizontal LAC tree* as a triple (T, α_T, R) where (T, α_T) is a horizontal LAC tree and R is a right increasing labeling of T . We denote by \widetilde{LAC}_n the set of labeled horizontal LAC trees with n non-root nodes.

The restriction of the map Ξ_{steep} from horizontal LAC trees to pairs of the form $(\mu, \mu_{\text{steep}})$ (as in Lemma 3.19) can be extended naturally in the context of labeled horizontal LAC trees. We denote by $\text{Vert}(\widetilde{\mathcal{SP}}_n)$ the set of triples $(\mu, \mu_{\text{steep}}, P)$ where $\mu \in \mathcal{D}_n$ and P is a vertical labeling of μ . Let

$$(17) \quad \tilde{\Xi}_{\text{steep}}: \widetilde{LAC}_n \rightarrow \text{Vert}(\widetilde{\mathcal{SP}}_n)$$

be the map that sends a labeled horizontal LAC tree (T, α_T, R) to $(\mu, \mu_{\text{steep}}, P)$, where $(\mu, \mu_{\text{steep}}) = \Xi_{\text{steep}}((T, \alpha_T))$ and P is obtained from R as follows: the non-root nodes in T are in correspondence with the north steps of μ through the bijection Ξ_{steep} , and we assign the labels of the non-root nodes in R to their corresponding north steps. Since the labels in T increase along right-most children, the resulting labeling of the north steps in μ is a vertical labeling. And vice-versa, every vertical labeling of μ arises uniquely from a labeling of T that increases along right-most children. As a consequence, the map $\tilde{\Xi}_{\text{steep}}$ is a bijection. We denote by $\tilde{\Lambda}_{\text{steep}}$ its inverse.

Similarly, we also have a labeled version of the restriction of the map Ξ_{bounce} from horizontal LAC trees to pairs of the form $(\nu_{\text{bounce}}, \nu)$ (as in Lemma 3.20). We denote by $\text{Diag}(\widetilde{\mathcal{BP}}_n)$ the set of triples $(\nu_{\text{bounce}}, \nu, Q)$ where $\nu \in \mathcal{D}_n$ and Q is a diagonal labeling of ν . Let

$$(18) \quad \tilde{\Xi}_{\text{bounce}}: \widetilde{LAC}_n \rightarrow \text{Diag}(\widetilde{\mathcal{BP}}_n)$$

²We are using the version of the $\tilde{\zeta}$ map described by Haglund in [17, Chapter 5], which differs from Haglund and Loehr's original description by just reversing the path $\zeta(\mu)$ and the labels on the diagonal.

be the map that sends a labeled horizontal LAC tree (T, α_T, R) to $(\nu_{\text{bounce}}, \nu, Q)$, where $(\nu_{\text{bounce}}, \nu) = \Xi_{\text{bounce}}((T, \alpha_T))$ and Q is obtained from R as follows: the labels in Q , from bottom to top, are obtained by reading from right-to-left the labels of non-root nodes of color i in T , for $i \in \{1, 2, \dots\}$. Since the valleys of ν are in correspondence with pairs of nodes (u, v) in T such that v is the right-most child of u , it is not hard to see that Q is a diagonal labeling of ν and that the map Ξ_{bounce} is a bijection. We denote by $\tilde{\Lambda}_{\text{bounce}}$ its inverse.

Finally, we denote by

$$(19) \quad \tilde{\Gamma}: \text{Vert}(\widetilde{\mathcal{SP}}_n) \rightarrow \text{Diag}(\widetilde{\mathcal{BP}}_n)$$

the composition $\tilde{\Gamma} = \tilde{\Xi}_{\text{bounce}} \circ \tilde{\Lambda}_{\text{steep}}$. As a consequence of Theorem IV and the discussion above, we get the following corollary.

Corollary 3.22. *For $n > 0$, the map $\tilde{\Gamma}$ is a bijection from*

- *the set of triples $(\mu, \mu_{\text{steep}}, P)$ such that $\mu \in \mathcal{D}_n$ and P is a vertical labeling of μ , to*
- *the set of triples $(\nu_{\text{bounce}}, \nu, Q)$ such that $\nu \in \mathcal{D}_n$ and Q is a diagonal labeling of ν .*

Moreover, if $(\nu_{\text{bounce}}, \nu, Q) = \tilde{\Gamma}(\mu, \mu_{\text{steep}}, P)$, then $(\nu, R) = \tilde{\zeta}(\mu, P)$, where $\tilde{\zeta}$ is the Haglund-Loehr's zeta map from parking functions to diagonally labeled Dyck paths.

Remark 3.23. Corollary 3.22 is essentially a labeled version of Theorem IV. One might expect to have a generalization of Theorem III involving labeled versions of steep pairs and bounce pairs. Indeed, one can naturally define a vertically labeled steep pair as a triple (μ_1, μ_2, P) where (μ_1, μ_2) is a nested pair such that the top path μ_2 is steep and P is a vertical labeling of μ_1 . However, for the definition of diagonally labeled bounce pairs (ν_1, ν_2, Q) , we need (ν_1, ν_2) to be a nested pair such that the bottom path ν_1 is bounce as expected, but the requirement on the labeling Q depends on both ν_1 and ν_2 and is more involved.

Remark 3.24. Suppose that we remove the increasing condition in vertical labelings and the valley condition in diagonal labelings, but instead record as statistics the number of violations of these conditions, i.e. the number of descents from below to above along columns in vertical labelings, and the number of valleys with the entry below larger than the one above in diagonal labelings. Then our bijection $\tilde{\Gamma}$ generalizes naturally to these labelings while transferring the two statistics from one to the other. It is because they are relayed by labeled LAC trees without the right increasing condition, and both statistics correspond to the number of non-root nodes whose label is larger than that of its rightmost child. The zeta map for parking functions is simply the case where all these statistics are zero.

DISCUSSION

In this section we highlight some possible directions of future research.

Enumerative questions. From an enumerative point of view, Mishna and Rechinizer found in [24] that the generating function of the lattice walks mentioned after Definition 3.2, which are in bijection with the level-marked Dyck paths, is not differentially algebraic with respect to the variable marking the end of the walk. This is an indication that the enumeration of these objects is complicated. It

would be interesting to study the enumeration of certain sub-families with simple enumeration formulas.

Generalizations for other Coxeter groups. An intriguing follow-up question is to exhibit combinatorial models for the algebraic generalizations of parabolic 231-avoiding permutations, parabolic nonnesting partitions, and parabolic noncrossing partitions to other Coxeter groups in [26, Section 6]. Observe that these generalizations may depend on the choice of Coxeter element. These combinatorial models may then be used to generalize the other families described here accordingly. If such an extension is possible, it will hopefully lead to a Steep-Bounce zeta map in other types, which would possibly be related to the zeta map for irreducible crystallographic root systems described by Thiel in [32].

Fuß-Catalan and rational Catalan generalizations. Another natural extension would be to consider Fuß-Catalan and rational Catalan objects. Is it possible to define an analogue of LAC trees in this world? If such an adaptation exists, we may define a Steep-Bounce map on Fuß-Catalan or even rational Catalan objects, therefore potentially generalize the zeta map in rational Catalan combinatorics.

Generalizations of the zeta map and statistics for multivariate diagonal harmonics. One of the most remarkable properties of the zeta map is that it describes combinatorial statistics for the Hilbert series of certain modules in diagonal harmonics. As mentioned by Armstrong, Loehr and Warrington in [2], one point of view is that, rather than having two statistics area and dinv , we have only one statistic area and a nice map ζ . There are generalizations for multisets of variables that give rise to certain modules of multivariate diagonal harmonics [3]. One important and difficult problem in this area is to find combinatorial models and a multiset of combinatorial statistics to describe them. In the three variate case, these modules are closely related to intervals in the Tamari lattice. Two of the statistics have been conjectured in [4] but the third statistic is still missing. Furthermore, at the moment no generalization of the zeta map for Tamari intervals is known. Can the results in this paper be used to find the third missing statistic or a zeta map for intervals in the Tamari lattice? Can we push this forward for more sets of variables in connection to the results in [5]? A positive answer to these questions may lead to new understanding of multivariate diagonal harmonics.

ACKNOWLEDGEMENTS

The authors thank Nantel Bergeron, Myrto Kallipoliti, Marni Mishna, Vincent Pilaud, and Robin Sulzgruber for inspiring discussions.

APPENDIX A. LATTICE-THEORETIC BACKGROUND

In this appendix we provide some lattice-theoretic background that is important for the proof of Theorem II.

A.1. Basics. Let $\mathcal{P} = (P, \leq)$ be a finite partially ordered set (or *poset* for short). An element $x \in P$ is *minimal* if for every $y \in P$ with $y \leq x$ follows $y = x$. Dually, x is *maximal* if for every $y \in P$ with $x \leq y$ follows $x = y$. Then, \mathcal{P} is *bounded* if there exists a unique minimal and a unique maximal element, and we denote these bounds by $\hat{0}$ and $\hat{1}$, respectively.

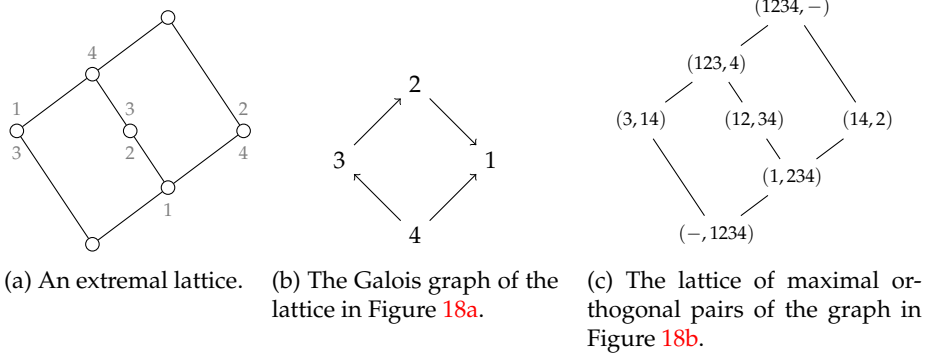


Figure 18. An extremal lattice with its associated Galois graph.

Two elements $x, y \in P$ form a *cover relation* if $x < y$ and there exists no $z \in P$ with $x < z < y$. We write $x < y$ in this case, and say that x is a *lower cover* of y and that y is a *upper cover* of x .

A *chain* of \mathcal{P} is a totally ordered subset of P . A chain is *saturated* if it can be written as a sequence of cover relations, and it is *maximal* if it is saturated and contains a minimal and a maximal element of \mathcal{P} . The *length* of \mathcal{P} is one less than the maximum cardinality of the maximal chains of \mathcal{P} , and it is denoted by $\ell(\mathcal{P})$.

The poset \mathcal{P} is a *lattice* if for every $x, y \in P$ there exists a least upper bound $x \vee y$ (the *join* of x and y), and a greatest lower bound $x \wedge y$ (the *meet* of x and y). It follows that every finite lattice is bounded.

If \mathcal{P} is a lattice, then an element $j \in P \setminus \{\hat{0}\}$ is *join irreducible* if for every $x, y \in P$ with $j = x \vee y$ we have $j \in \{x, y\}$. Dually, $m \in P \setminus \{\hat{1}\}$ is *meet irreducible* if for every $x, y \in P$ with $m = x \wedge y$ we have $m \in \{x, y\}$. We denote the set of all join-irreducible elements of \mathcal{P} by $\mathcal{J}(\mathcal{P})$, and the set of all meet-irreducible elements of \mathcal{P} by $\mathcal{M}(\mathcal{P})$. Moreover, we observe that if P is finite, then every $j \in \mathcal{J}(\mathcal{P})$ has a unique lower cover j_* , and every $m \in \mathcal{M}(\mathcal{P})$ has a unique upper cover m^* .

A.2. Extremal Lattices. Let $\mathcal{L} = (L, \leq)$ be a finite lattice. We follow [23] and say that \mathcal{L} is *extremal* if $|\mathcal{J}(\mathcal{L})| = \ell(\mathcal{L}) = |\mathcal{M}(\mathcal{L})|$. Extremal lattices admit a nice presentation as lattices that arise from certain directed graphs [23], see also [35].

Suppose that \mathcal{L} is extremal with $\ell(\mathcal{L}) = n$, and we choose a maximal chain $C : x_0 < x_1 < \dots < x_n$. We may totally order the join- and meet-irreducible elements of \mathcal{L} as j_1, j_2, \dots, j_n and m_1, m_2, \dots, m_n such that

$$j_1 \vee j_2 \vee \dots \vee j_i = x_i = m_{i+1} \wedge m_{i+2} \wedge \dots \wedge m_n$$

holds for all $i \in [n]$.

From this order we may define a directed graph with vertex set $[n]$, where we have a directed edge $i \rightarrow k$ if and only if $i \neq k$ and $j_i \not\leq m_k$. This graph is the *Galois graph* of \mathcal{L} , denoted by $\text{Gal}(\mathcal{L})$. Figure 18a shows an extremal lattice, and Figure 18b shows its Galois graph. The join- and meet-irreducible elements of the lattice in Figure 18a are labeled as described in the previous paragraph. We observe that this lattice has a unique maximal chain of maximum length.

Now let G be a simple directed graph with vertex set $[n]$ and edges $i \rightarrow k$ only if $i > k$. Two sets $X, Y \subseteq [n]$ form an *orthogonal pair* if $X \cap Y = \emptyset$ and there exists

no edge $i \rightarrow k$ in G with $i \in X$ and $k \in Y$. An orthogonal pair (X, Y) is *maximal* if whenever there exist $Y' \supseteq Y$ and $X' \supseteq X$ such that (X', Y') is an orthogonal pair, then $X' = X$ and $Y' = Y$. Let $\text{MO}(G)$ denote the set of maximal orthogonal pairs of G .

For $(X, Y), (X', Y') \in \text{MO}(G)$ we declare $(X, Y) \sqsubseteq (X', Y')$ if $X \subseteq X'$ (or equivalently $Y \supseteq Y'$). It can be verified that $\mathcal{L}(G) \stackrel{\text{def}}{=} (\text{MO}(G), \sqsubseteq)$ is a lattice, where the meet $(X, Y) \sqcap (X', Y')$ is the unique maximal orthogonal pair $(X \cap X', Z)$ and the join $(X, Y) \sqcup (X', Y')$ is the unique maximal orthogonal pair $(Z, Y \cap Y')$. We have the following result.

Theorem A.1 ([23, Section 4]). *If G is a simple directed graph on the vertex set $[n]$ which has edges $i \rightarrow k$ only if $i > k$, then $\mathcal{L}(G)$ is an extremal lattice of length n . Conversely, every finite extremal lattice \mathcal{L} is isomorphic to $\mathcal{L}(\text{Gal}(\mathcal{L}))$.*

The lattice of maximal orthogonal pairs of the directed graph in Figure 18b is shown in Figure 18c. In this figure, we have omitted set parentheses and commas for brevity.

Remark A.2. Let $\mathcal{L}, \mathcal{L}'$ be two extremal lattices with corresponding Galois graphs $\text{Gal}(\mathcal{L}), \text{Gal}(\mathcal{L}')$. If $f: \text{Gal}(\mathcal{L}) \rightarrow \text{Gal}(\mathcal{L}')$ is an isomorphism of directed graphs, then it is clear that (X, Y) is a maximal orthogonal pair of $\text{Gal}(\mathcal{L})$ if and only if $(f(X), f(Y))$ is a maximal orthogonal pair of $\text{Gal}(\mathcal{L}')$. It follows that $\mathcal{L}(\text{Gal}(\mathcal{L})) \cong \mathcal{L}(\text{Gal}(\mathcal{L}'))$, and in view of Theorem A.1 we conclude that $\mathcal{L} \cong \mathcal{L}'$.

Remark A.3. Observe that not every isomorphism of Galois graphs extends directly to an isomorphism of the corresponding extremal lattices. Consider for instance the pentagon lattice together with the bijection that swaps top and bottom and fixes everything else. This map is the identity on the Galois graph, but is not a lattice isomorphism.

A.3. Congruence-Uniform Lattices. Let us recall the following doubling construction due to Day [10]. Let $\mathcal{P} = (P, \leq)$ be a finite poset. For $X \subseteq P$ we define

$$P_{\leq X} \stackrel{\text{def}}{=} \{y \in P \mid y \leq x \text{ for some } x \in X\}.$$

Let $\mathbf{2}$ denote the unique two-element lattice on the ground set $\{0, 1\}$. The *doubling* of \mathcal{P} by X is the subposet of the direct product $\mathcal{P} \times \mathbf{2}$ induced by the set

$$\left(P_{\leq X} \times \{0\}\right) \uplus \left(\left((P \setminus P_{\leq X}) \cup X\right) \times \{1\}\right).$$

A lattice is *congruence uniform* if it can be constructed from the singleton lattice by a sequence of doublings by intervals [10, Theorem 5.1]. Figure 19 shows an example of a congruence-uniform lattice and the sequence of doublings from which it is constructed.

If a congruence-uniform lattice is also extremal, then we can realize the corresponding Galois graph as a directed graph whose vertices *are* the join-irreducible elements.

Proposition A.4 ([25, Corollary 2.17]). *Let $\mathcal{L} = (L, \leq)$ be an extremal, congruence-uniform lattice. The Galois graph of \mathcal{L} is the directed graph whose vertex set is $\mathcal{J}(\mathcal{L})$, and which has a directed edge $i \rightarrow k$ if and only if $i \neq k$ and $k \leq k_* \vee i$.*




Figure 19. A sequence of six interval doublings that yields a congruence-uniform lattice with twelve elements. At each step we double by the highlighted interval. Moreover, we have labeled the join-irreducible elements in the rightmost lattice.




Figure 20. The graph of the lattice in Figure 18a constructed according to the rules in Proposition A.4.

Let us illustrate that the construction in Proposition A.4 does not in general yield the Galois graph of an extremal lattice. Let \mathcal{L} be the extremal lattice from Figure 18a. We can check that \mathcal{L} is not congruence uniform. If we construct a directed graph on the join-irreducible elements of \mathcal{L} according to the rules in Proposition A.4, we get the graph in Figure 20, which is not the same graph as $\text{Gal}(\mathcal{L})$ shown in Figure 18b.

Conversely, the congruence-uniform lattice \mathcal{L} on the right of Figure 19 is also extremal. The directed graph $G(\mathcal{L})$ constructed according to the rules in Proposition A.4 is shown in Figure 21a. Figure 21b shows the lattice of maximal orthogonal pairs of $G(\mathcal{L})$, which is isomorphic to \mathcal{L} . This proves that $G(\mathcal{L})$ is indeed (isomorphic to) the Galois graph of \mathcal{L} . As an example, the map $(a, b, c, d, e, f) \mapsto (1, 2, 6, 4, 3, 5)$ realizes $G(\mathcal{L})$ in Figure 21a as a graph that satisfies the conditions of Theorem A.1.

REFERENCES

[1] George E. Andrews, Christian Krattenthaler, Luigi Orsina, and Paolo Papi, *ad-nilpotent \mathfrak{b} -ideals in $\mathfrak{sl}(n)$ having a fixed class of nilpotence: Combinatorics and enumeration*, Transactions of the American Mathematical Society **354** (2002), no. 10, 3835–3853.
 [2] Drew Armstrong, Nicholas A. Loehr, and Gregory S. Warrington, *Sweep maps: A continuous family of sorting algorithms*, Advances in Mathematics **284** (2015), 159–185.
 [3] François Bergeron, *Multivariate diagonal coinvariant spaces for complex reflection groups*, Advances in Mathematics **239** (2013), 97–108.
 [4] François Bergeron and Louis-François Prévaille-Ratelle, *Higher trivariate diagonal harmonics via generalized Tamari posets*, Journal of Combinatorics **3** (2012), 317–341.

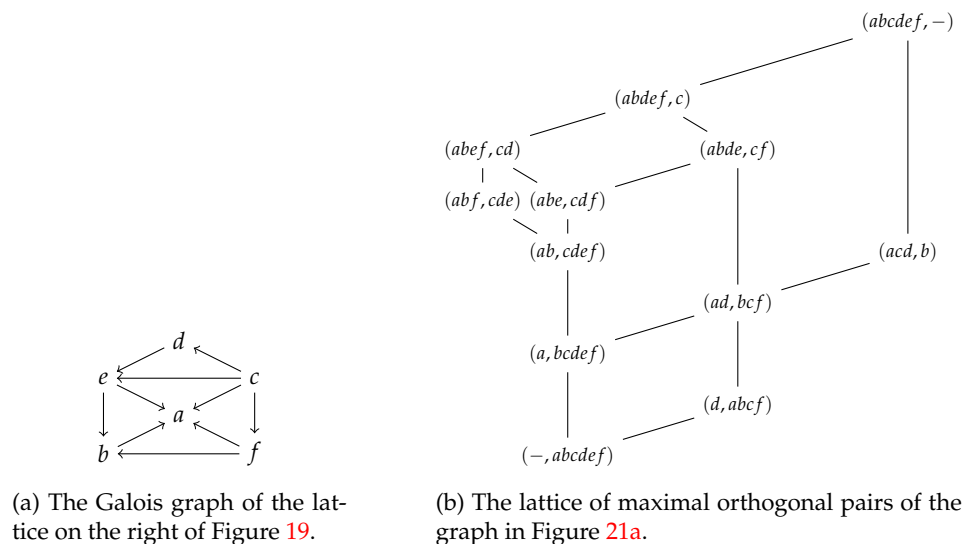


Figure 21. An illustration of Proposition A.4.

- [5] Nantel Bergeron, Cesar Ceballos, and Vincent Pilaud, *Hopf Dreams and Diagonal Harmonics* (2018), available at [arXiv:1807.03044](https://arxiv.org/abs/1807.03044).
- [6] Anders Björner and Michelle L. Wachs, *Shellable and nonpure complexes and posets II*, Transactions of the American Mathematical Society **349** (1997), 3945–3975.
- [7] Erik Carlsson and Anton Mellit, *A proof of the shuffle conjecture*, Journal of the American Mathematical Society **31** (2018), no. 3, 661–697.
- [8] Cesar Ceballos, Tom Denton, and Christopher R. H. Hanusa, *Combinatorics of the zeta map on rational Dyck paths*, Journal of Combinatorial Theory, Series A **141** (2016), 33–77.
- [9] Cesar Ceballos, Arnau Padrol, and Carmilo Sarmiento, *The v -Tamari lattice via v -trees, v -bracket vectors, and subword complexes* (2018), available at [arXiv:1805.03566](https://arxiv.org/abs/1805.03566).
- [10] Alan Day, *Characterizations of finite lattices that are bounded-homomorphic images or sublattices of free lattices*, Canadian Journal of Mathematics **31** (1979), 69–78.
- [11] Emeric Deutsch, *Dyck path enumeration*, Discrete Mathematics **204** (1999), 167–202.
- [12] Wenjie Fang and Louis-François Prévaille-Ratelle, *The enumeration of generalized Tamari intervals*, European Journal of Combinatorics **61** (2017), 69–84.
- [13] Adriano M. Garsia and James Haglund, *A proof of the q, t -Catalan positivity conjecture*, Discrete Mathematics **256** (2002), no. 3, 677–717.
- [14] Adriano M. Garsia and Mark Haiman, *A remarkable q, t -Catalan sequence and q -Lagrange inversion*, Journal of Algebraic Combinatorics **5** (1996), no. 3, 191–244.
- [15] Winfried Geyer, *On Tamari lattices*, Discrete Mathematics **133** (1994), 99–122.
- [16] James Haglund, *Conjectured statistics for the q, t -Catalan numbers*, Advances in Mathematics **175** (2003), no. 2, 319–334.
- [17] ———, *The q, t -Catalan Numbers and the Space of Diagonal Harmonics*, Vol. 41, American Mathematical Society, Providence, RI, 2008.
- [18] James Haglund, Mark Haiman, Nicholas A. Loehr, Jeffrey B. Remmel, and Alexander Ulyanov, *A combinatorial formula for the character of the diagonal coinvariants*, Duke Mathematical Journal **126** (2005), no. 2, 195–232.
- [19] James Haglund and Nicholas A. Loehr, *A conjectured combinatorial formula for the Hilbert series for diagonal harmonics*, Discrete Mathematics **298** (2005), no. 1–3, 189–204.
- [20] Mark Haiman, *Vanishing theorems and character formulas for the Hilbert scheme of points in the plane*, Inventiones Mathematicae **149** (2002), no. 2, 371–407.
- [21] Donald E. Knuth, *The Art of Computer Programming, Vol. 1, Fundamental Algorithms*, 3rd ed., Addison-Wesley, Reading, MA, 1973.
- [22] Germain Kreweras, *Sur les partitions non croisées d’un cycle*, Discrete Mathematics **1** (1972), 333–350.

- [23] George Markowsky, *Primes, irreducibles and extremal lattices*, *Order* **9** (1992), 265–290.
- [24] Marni Mishna and Andrew Rechnitzer, *Two non-holonomic lattice walks in the quarter plane*, *Theoretical Computer Science* **410** (2009), no. 38–40, 3616–3630.
- [25] Henri Mühle, *Noncrossing partitions, Tamari lattices, and parabolic quotients of the symmetric group* (2018), available at [arXiv:1809.01405](https://arxiv.org/abs/1809.01405).
- [26] Henri Mühle and Nathan Williams, *Tamari lattices for parabolic quotients of the symmetric group*, *The Electronic Journal of Combinatorics* **26** (2019), Article P4.34, 28 pages.
- [27] Folkert Müller-Hoissen, Jean Marcel Pallo, and Jim Stasheff, *Associahedra, Tamari Lattices and Related Structures*, Birkhäuser, Basel, 2012.
- [28] Louis-François Préville-Ratelle and Xavier Viennot, *The enumeration of generalized Tamari intervals*, *Transactions of the American Mathematical Society* **369** (2017), 5219–5239.
- [29] Christian Stump, Hugh Thomas, and Nathan Williams, *Cataland: Why the Fuß?* (2018), available at [arXiv:1503.00710](https://arxiv.org/abs/1503.00710).
- [30] Dov Tamari, *Monoïdes Préordonnés et Chaînes de Malcev*, Thèse de Mathématiques, Université de Paris, 1951.
- [31] ———, *The algebra of bracketings and their enumeration*, *Nieuw Archief voor Wiskunde* **10** (1962), 131–146.
- [32] Marko Thiel, *From Anderson to zeta*, *Advances in Applied Mathematics* **81** (2016), 156–201.
- [33] Hugh Thomas, *An analogue of distributivity for ungraded lattices*, *Order* **23** (2006), 249–269.
- [34] Hugh Thomas and Nathan Williams, *Sweeping up zeta*, *Selecta Mathematica. New Series* **24** (2018), no. 3, 2003–2034.
- [35] ———, *Rowmotion in slow motion*, *Proceedings of the London Mathematical Society* **119** (2019), 1149–178.
- [36] Nathan Williams, *Cataland*, Dissertation, University of Minnesota, 2013.

FACULTY OF MATHEMATICS, UNIVERSITY OF VIENNA, 1090 VIENNA, AUSTRIA.
Email address: cesar.ceballos@univie.ac.at

UNIVERSITÉ PARIS-EST MARNE-LA-VALLÉE, LIGM (UMR 8094), CNRS, ENPC, ESIEE PARIS, FRANCE.
Email address: wenjie.fang@u-pem.fr

TECHNISCHE UNIVERSITÄT DRESDEN, INSTITUT FÜR ALGEBRA, ZELLESCHER WEG 12–14, 01069 DRESDEN, GERMANY.
Email address: henri.muehle@tu-dresden.de