FPSAC 2012, Nagoya, Japan DMTCS proc. **AR**, 2012, 769–780

Dyck tilings, linear extensions, descents, and inversions (extended abstract)

Jang Soo Kim[1] Karola Mészáros[2†] Greta Panova[3‡] David B. Wilson[4]

[1] *University of Minnesota, Minneapolis, Minnesota 55455, USA*
[2] *University of Michigan, Ann Arbor, MI 48109, USA*
[3] *University of California at Los Angeles, Los Angeles, CA 90095, USA*
[4] *Microsoft Research, Redmond, WA 98052, USA*

Abstract. Dyck tilings were introduced by Kenyon and Wilson in their study of double-dimer pairings. They are certain kinds of tilings of skew Young diagrams with ribbon tiles shaped like Dyck paths. We give two bijections between "cover-inclusive" Dyck tilings and linear extensions of tree posets. The first bijection maps the statistic (area + tiles)/2 to inversions of the linear extension, and the second bijection maps the "discrepancy" between the upper and lower boundary of the tiling to descents of the linear extension.

Résumé. Les pavages de Dyck ont été introduits par Kenyon et Wilson dans leur étude du modèle des "double-dimères". Ce sont des pavages des diagrammes de Young gauches avec des tuiles en forme de rubans qui ressemblent à des chemins de Dyck. Nous donnons deux bijections entre les pavages de Dyck "couvre-inclusive" et les extensions linéaires de posets dont le diagramme de Hasse est un arbre. La première bijection transforme la statistique (aire + tuiles) / 2 en inversions de l'extension linéaire, et la deuxième bijection transforme la "discordance" entre la limite supériere et inférieure du pavage en descentes de l'extension linéaire.

Keywords: Dyck tiling, Dyck path, linear extension, tree poset, perfect matching

1 Introduction

Kenyon and Wilson introduced the notion of "cover-inclusive Dyck tilings" (defined below) while studying certain models from statistical physics. The probabilities of certain events in these models are given by formulas that involve counting these Dyck tilings [KW11a, KW11b]. We give two new bijections between Dyck tilings and perfect matchings, which when restricted to Dyck tilings with a certain "lower path," become bijections to linear extensions of a tree poset. The first bijection is compatible with the number of inversions of a linear extension, and gives a bijective proof of a formula that was conjectured by Kenyon and Wilson [KW11a, Conjecture 1] and proved non-bijectively by Kim [Kim11]. The second bijection is compatible with descents of the linear extension, and leads to a new enumeration formula.

†Supported by an NSF postdoctoral fellowship.
‡Supported by a Simons postdoctoral fellowship.

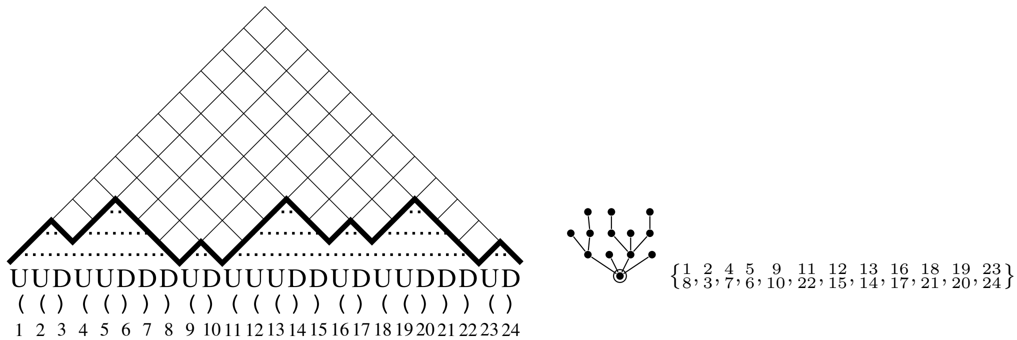

Fig. 1: A Dyck path λ of order $n = 12$ is shown in bold on the left. Above it is its associated (rotated) Young diagram. Below the Dyck path is its corresponding Dyck word, its balanced parentheses expression, and coordinates for the step positions. The horizontal dotted lines depict the chords between the matching up and down steps. In the middle is the rooted plane tree corresponding to the Dyck path, where each node other than the root corresponds to a chord. On the right is the set of chords of the Dyck path, where each chord is represented as the pair of coordinates of the matching up and down steps. (See [Sta99, exercise 6.19].) A linear extension of the chord poset of λ can be represented by placing the numbers $1, \dots, n$ on the chords of the diagram on the left (so that if one chord is nested within another, it gets a higher number), or equivalently by placing the numbers $0, \dots, n$ on the vertices of the rooted plane tree (so that the numbers increase when going up the tree), or equivalently by ordering the set of chords represented as pairs (so that if one pair is nested within another, it occurs later). See also Figure 4.

Dyck paths of order n are often defined as staircase lattice paths on an $n \times n$ square grid, from the lower-left corner to the upper-right corner, which do not go below the antidiagonal. Each such Dyck path has associated with it a Young diagram formed from the boxes above and to the left of it. If we rotate the lattice by $45°$ and dilate it by a factor of $\sqrt{2}$, then each step of the Dyck path is either $(+1, +1)$ (an up or "U" step) or $(+1, -1)$ (a down or "D" step). (See Figure 1.) This rotated form will be more convenient for us to work with. A Dyck path's sequence of U and D steps, when concatenated, forms a word which is called a Dyck word. If the U steps are written as "`(`" and the down steps are written as "`)`", then the Dyck word is a balanced parentheses expression.

We define a **chord** of a Dyck path λ to be the horizontal segment between an up step and the matching down step, as shown in Figure 1. The chords of a Dyck path λ naturally form a **chord poset** P_λ, where nesting is the order relation, i.e., one chord is above another chord in the partial order if its horizontal span lies within the span of the other chord, or equivalently, if the `(` and `)` corresponding to the first chord are nested within the `(` and `)` corresponding to the second chord. If we adjoin a bottom-most element to the chord poset, we call the result the tree poset, since its Hasse diagram is a rooted plane tree. See Figure 1 for an illustration of this partial order. There is a natural bijection between Dyck paths of order n and rooted plane trees with $n + 1$ vertices (see [Sta99, exercise 6.19]).

A **natural labeling** of a poset P with n elements is an order-preserving bijection $L : P \to [n]$, where $[n]$ denotes $\{1, 2, \dots, n\}$. For the tree poset associated with λ, it is more convenient to take a natural labeling of the chord poset P_λ, and then label root 0, as shown in Figure 2. A rooted tree embedded in the plane with a natural labeling is called an **increasing plane oriented tree**. As is well known, there are $(2n - 1)!!$ increasing plane oriented trees on $n + 1$ vertices (see [BFS92, Corollary 1(iv)]). Given a labeled tree L, if we delete all vertices with labels larger than $k \leq n$, the result is a labeled tree $L^{(k)}$ on

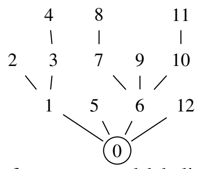

(a) Reference natural labeling L_0 of tree associated with λ, obtained from left-to-right depth-first search.

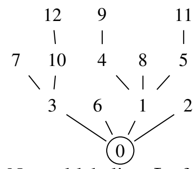

(b) Natural labeling L of tree.

$\frac{11}{22}, \frac{23}{24}, \frac{1}{8}, \frac{12}{15}, \frac{18}{21}, \frac{9}{10}, \frac{2}{3}, \frac{16}{17}, \frac{13}{14}, \frac{4}{7}, \frac{19}{20}, \frac{5}{6}$

(c) List of left and right endpoints of chords $\frac{\ell_1,\dots,\ell_n}{r_1,\dots,r_n}$ associated with λ in the order given by L.

$3, 7, 10, 12, 6, 1, 4, 9, 8, 5, 11, 2$

(d) Preorder word, list of labels of L encountered in left-to-right depth first search, $L \circ L_0^{-1}$.

$6, 12, 1, 7, 10, 5, 2, 9, 8, 3, 11, 4$

(e) Element of $\mathscr{L}(P_\lambda)$, inverse preorder word, standardization of $\ell_1, \dots, \ell_n$, $L_0 \circ L^{-1}$.

Fig. 2: Natural labelings of a rooted plane tree and their associated permutations.

$k + 1$ vertices (including the root labeled 0). Given $L^{(k-1)}$, there are $2k - 1$ possible positions where the vertex with label k may be attached to the labeled tree $L^{(k-1)}$. (Each time a new vertex gets added to the tree, the subsequent vertex has one less possible attachment location but three new ones: just before, just after, and on top of the new vertex.) Thus, to each labeled tree (P_λ, L) we can derive a sequence of attachment sites $p_1, \dots, p_n$, where $0 \le p_i < 2i - 1$. Any such sequence determines a poset P_λ together with a natural labeling L of P_λ, and the map from sequences to pairs (P_λ, L) is a bijection. In terms of the endpoints of the chords, this sequence of insertion locations is given by

$$p_i = \#\{j < i \ : \ \ell_j < \ell_i\} + \#\{j < i \ : \ r_j < \ell_i\},$$

where ℓ_i and r_i denote the left and right endpoints of the chord labeled i, as in Figure 2 c.

Let L_0 be the natural labeling of P_λ which orders the chords by their left endpoints (Figure 2 a). The **preorder word** of a natural labeling L of P_λ is $L \circ L_0^{-1}$, which is the permutation on $[n]$ obtained by reading the labels of L (excluding 0) in a left-to-right depth-first order (Figure 2 d). The inverse of the preorder word, $\sigma = L_0 \circ L^{-1}$, turns out to be a more natural object. It can also be obtained as the "standardization" of the sequence of left-endpoints $\ell_1, \dots, \ell_n$, i.e., σ is the permutation on $[n]$ for which $\sigma_i < \sigma_j$ iff $\ell_i < \ell_j$ (Figure 2 e). If ω is a natural labeling of the poset P, Stanley [Sta72] defines

$$\mathscr{L}(P, \omega) = \{\omega \circ L^{-1} : L \text{ is a natural labeling of } P\},$$

and these are sometimes called the linear extensions of the labeled poset (P, ω). We will abbreviate $\mathscr{L}(P_\lambda) = \mathscr{L}(P_\lambda, L_0)$.

It is also well-known that there are $(2n - 1)!!$ perfect matchings on the numbers $1, \dots, 2n$. Given the sequence $p_1, \dots, p_n$, one natural way to associate a perfect matching $\operatorname{match}(p_1, \dots, p_n)$ with it is to take $\operatorname{match}(p_1, \dots, p_{n-1})$, increment all the numbers that are bigger then p_n, and then adjoin pair $(p_n + 1, 2n)$. We define the **min-word** of the matching to be the list of the smaller item of each pair, sorted in order of the larger item in each pair, see Figure 4.

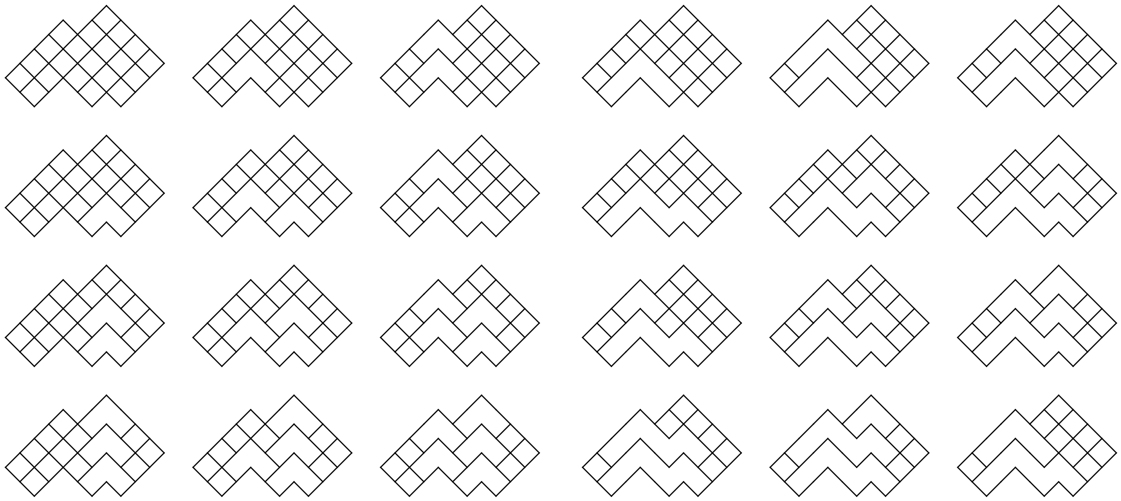

Fig. 3: All the cover-inclusive Dyck tilings of a particular skew shape. (This figure first appeared in [KW11a].)

Given two Dyck paths λ and μ of order n, if the path μ is at least as high as the path λ at each horizontal position, then we write $\mu \subset \lambda$ (since the Young diagram associated with μ is a subset of λ's Young diagram), and we write λ/μ for the skew Young diagram which consists of the boxes between λ and μ.

Dyck tiles are ribbon tiles (connected skew shapes that do not contain a 2×2 rectangle) in which the leftmost and rightmost boxes are at the same height, and no box within the tile is below these endpoints. (If each vertex in a Dyck path is replaced with a box, and the boxes are glued together, then the result is a Dyck tile, which explains the terminology.) A tiling of a skew Young diagram by Dyck tiles is a **Dyck tiling**. We say that one Dyck tile covers another Dyck tile if the first tile has at least one box whose center lies straight above the center of a box in the second tile. A Dyck tiling is called **cover-inclusive** if for each pair of its tiles, when the first tile covers the second tile, then the horizontal extent of the first tile is included as a subset within the horizontal extent of the second tile. We denote by $\mathcal{D}(\lambda, \mu)$ the set of all cover-inclusive Dyck tilings of shape λ/μ, and let

$$\mathcal{D}(\lambda, *) = \bigcup_{\mu} \mathcal{D}(\lambda/\mu).$$

Figure 3 shows all the cover-inclusive Dyck tilings of a certain skew shape.

It was observed empirically that there is a close connection between Dyck tilings and linear extensions. More specifically, for a Dyck path λ of order n, the total number of cover-inclusive Dyck tilings of skew shape λ/μ for some μ was conjectured [KW11a, Conjecture 1] and subsequently proven [Kim11] to be

$$|\mathcal{D}(\lambda, *)| = \frac{n!}{\prod_{\text{chords } c \text{ of } \lambda} |c|}, \tag{1}$$

where $|c|$ is the length of the chord c as measured in terms of number of up steps in λ between and including its ends. This formula has the form of the tree hook-length formula of Knuth [Knu73, pg. 70] for the number of linear extensions of the tree poset P_λ. These formulas call out for a bijection between cover-inclusive Dyck tilings and linear extensions. Here we give two such bijections. These bijections were in part inspired by a bijection due to Aval, Boussicault, and Dasse-Hartaut [ABDH12] (see also [ABN11]), to which one of our bijections specializes in the case where P_λ is the antichain. Our bijections actually provide refined enumeration formulas, which relate statistics (defined below) on the Dyck tilings to well-studied statistics on the permutations.

For two Dyck paths λ and μ, let $\mathbf{dis}(\lambda,\mu)$ be the "discrepancy" between λ and μ, i.e., the number of locations where λ has a down step while μ has an up step (which also equals half the number of locations where λ and μ step in opposite directions, i.e., half the Hamming distance between the Dyck words of λ and μ). For a skew shape λ/μ we define $\mathrm{dis}(\lambda/\mu) = \mathrm{dis}(\lambda,\mu)$.

For a Dyck tiling T of shape $\mathrm{sh}(T) = \lambda/\mu$, we define $\mathrm{dis}(T) = \mathrm{dis}(\lambda,\mu)$, and $\mathbf{tiles}(T)$ to be the number of tiles in T, and $\mathbf{area}(T)$ to be the number of boxes of the skew shape λ/μ. We define

$$\mathrm{art}(T) = (\underline{\mathrm{ar}}\mathrm{ea}(T) + \underline{\mathrm{t}}\mathrm{iles}(T))/2.$$

The art statistic is always integer-valued since each tile has odd area, and appears to be more natural than the area statistic.

Recall that the **inversion** statistic of a permutation σ on $[n]$ is defined by

$$\mathrm{inv}(\sigma) = \#\{(i,j) : 1 \leq i < j \leq n, \sigma(i) > \sigma(j)\}$$

and the **descent** statistic is defined by

$$\mathrm{des}(\sigma) = \#\{i < n : \sigma_i > \sigma_{i+1}\}.$$

In § 2 we give our two bijections, which we call DTS and DTR, which stand for "Dyck tiling strip" and "Dyck tiling ribbon" respectively, for reasons that will become apparent. The functions DTS and and DTR are bijections from the sequences $p_1,\dots,p_n$ to cover-inclusive Dyck tilings of order n (without restrictions on the lower path λ or upper path μ). As discussed above, these sequences $p_1,\dots,p_n$ are also in bijection with increasing plane oriented trees, so we can write $\mathrm{DTS}(\lambda,\sigma)$ and $\mathrm{DTR}(\lambda,\sigma)$ for the maps which take the labeled tree defined by λ and σ to the sequence $p_1,\dots,p_n$ and then to the Dyck tiling. These three bijections are compatible with one another in the sense that the lower paths of $\mathrm{DTS}(\lambda,\sigma)$ and $\mathrm{DTR}(\lambda,\sigma)$ are both λ. Thus, for a given Dyck path λ of order n, the maps $\mathrm{DTS}(\lambda,\cdot)$ and $\mathrm{DTR}(\lambda,\cdot)$ are bijections from $\mathscr{L}(P_\lambda)$ to $\mathcal{D}(\lambda,*)$, see Figure 4. Furthermore, if $\sigma \in \mathscr{L}(P_\lambda)$, then

$$\mathrm{art}(\mathrm{DTS}(\lambda,\sigma)) = \mathrm{inv}(\sigma)$$

and

$$\mathrm{dis}(\mathrm{DTR}(\lambda,\sigma)) = \mathrm{des}(\sigma).$$

It is natural to ask if there is a bijection from linear extensions in $\mathscr{L}(P_\lambda)$ to cover-inclusive Dyck tilings with lower path λ that simultaneously maps inv to art and des to dis, but as can be seen from Figure 4, no such bijection exists.

Björner and Wachs [BW89] gave the following q-analog of Knuth's tree hook-length formula:

$$\sum_{\sigma\in\mathscr{L}(P_\lambda)} q^{\mathrm{inv}(\sigma)} = \frac{[n]_q!}{\prod_{\text{vertices } v \,\in\, P_\lambda} [|\text{subtree rooted at } v|]_q},$$

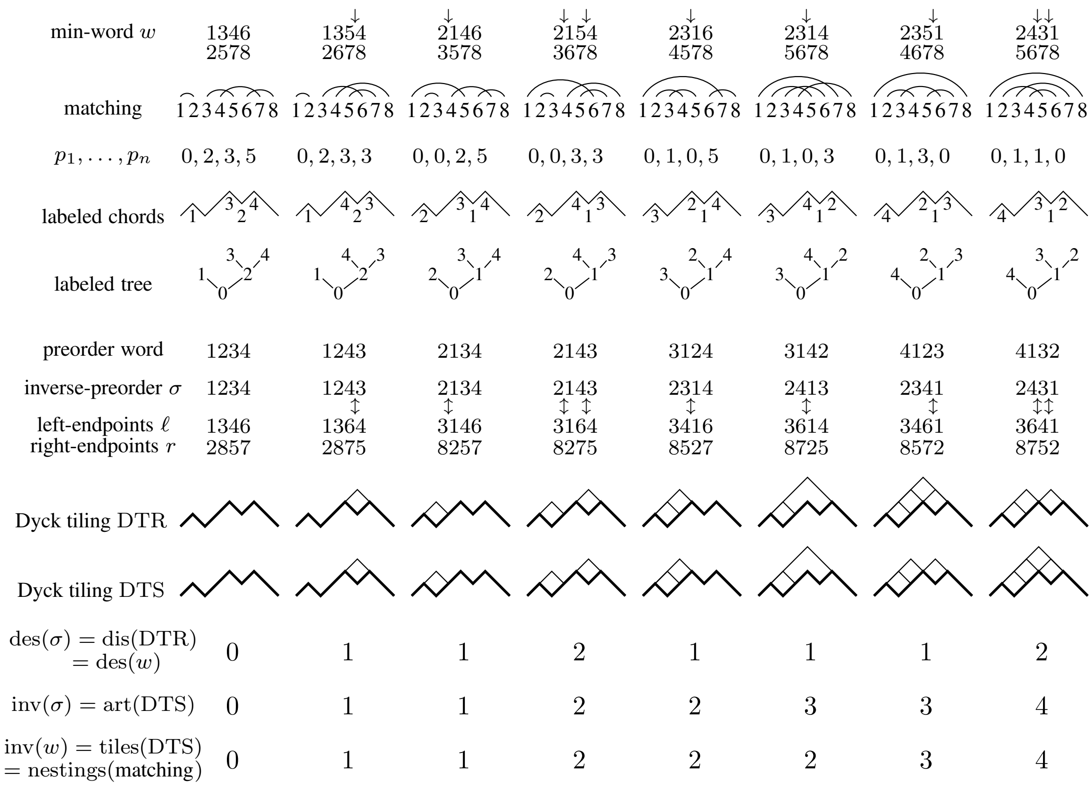

Fig. 4: Objects associated with the Dyck path λ =UDUUDUDD. The fourth row shows natural labelings of the chord poset P_λ. The fifth row shows essentially the same thing — natural labelings of the tree poset associated with λ. The third row shows the growth sequences $p_1, \dots, p_n$ which correspond to these labeled trees. The sixth row shows the labels from the tree poset listed in depth-first search order, where children are searched in left-to-right order. This is equivalent to listing the chord labels in order of the left endpoint. The seventh row shows the inverse of the permutation from the sixth row, with marks at the descents. These are the permutations σ of $\mathscr{L}(P_\lambda)$. The eighth row shows the left and right endpoints of the chords of λ, listed in the order of the natural labeling. The ninth and tenth rows shows cover-inclusive Dyck tilings $\mathrm{DTS}(\lambda, \sigma)$ and $\mathrm{DTR}(\lambda, \sigma)$ whose lower path is λ. The second row shows perfect matchings on $1, \dots, 2n$ that correspond to the sequences $p_1, \dots, p_n$ in the third row. The first row shows these same perfect matchings, represented as a $2 \times n$ array of numbers, where the columns are sorted, and the bottom row is sorted, together with markings at the descents in the top row. The top row of each $2 \times n$ array is the min-word w. The eleventh row gives $\mathrm{des}(w) = \mathrm{des}(\sigma) = \mathrm{des}(\ell) = \mathrm{dis}(\mathrm{DTR})$. The twelfth row gives $\mathrm{inv}(\sigma) = \mathrm{inv}(\ell) = \mathrm{art}(\mathrm{DTS})$. The thirteenth row gives $\mathrm{inv}(w) = \mathrm{tiles}(\mathrm{DTS}) = \mathrm{nestings}(\text{matching})$.

where $[n]_q = 1 + q + \cdots + q^{n-1}$ and $[n]_q! = [1]_q \cdots [n]_q$. Using this formula together with the DTS bijection, we get a bijective proof of the following theorem, originally proven in [Kim11] using inductive computation:

Theorem 1 *[KW11a, Conjecture 1][Kim11] Given a Dyck path λ of order n, we have*

$$\sum_{\textit{Dyck tilings } T \in \mathcal{D}(\lambda,*)} q^{\mathrm{art}(T)} = \frac{[n]_q!}{\prod_{\textit{chords } c \textit{ of } \lambda} [|c|]_q}. \tag{2}$$

In turn, the DTR bijection implies

Theorem 2

$$\sum_{\textit{Dyck tilings } T \in \mathcal{D}(\lambda,*)} z^{\mathrm{dis}(T)} = \sum_{\sigma \in \mathscr{L}(P_\lambda)} z^{\mathrm{des}(\sigma)}. \tag{3}$$

It is evident that the generating function (2) depends only on the poset structure of P_λ, and not upon how P_λ is embedded in the plane. The descent statistic for labeled trees is also invariant under different embeddings of the tree, as shown in [Sta72, Thm. 9.1 and Prop. 14.1]. Thus the generating function (3) also only depends on the poset structure of P_λ, and not its embedding in the plane. We conjecture that the tiles statistic is also invariant under different embeddings of P_λ. By contrast, neither the area statistic nor any of the joint statistics between art, dis and tiles have this embedding-invariance property.

The Dyck tilings $\mathrm{DTS}(p_1,\dots,p_n)$ and $\mathrm{DTR}(p_1,\dots,p_n)$ can also be understood in terms of the perfect matching $\mathrm{match}(p_1,\dots,p_n)$, as we discuss in § 2. We shall see, for example, that

$$\mathrm{dis}(\mathrm{DTR}(p_1,\dots,p_n)) = \mathrm{des}(\text{min-word}(\mathrm{match}(p_1,\dots,p_n)))$$

and

$$\mathrm{tiles}(\mathrm{DTS}(p_1,\dots,p_n)) = \mathrm{inv}(\text{min-word}(\mathrm{match}(p_1,\dots,p_n))).$$

An interesting special case for our bijections is when the lower path λ is the "zig-zag" path

$$\mathrm{zigzag}_n = \text{UDUDUD}\dots\text{UD} = (\mathrm{UD})^n,$$

so that P_λ is the antichain, and $\mathscr{L}(P_\lambda)$ is the set of all permutations. In this case, the bijection DTR restricts to the bijection between permutations and Dyck tableaux in [ABDH12].

We show that when $\lambda = \mathrm{zigzag}_n$,

$$\mathrm{art}(\mathrm{DTR}(\mathrm{zigzag}_n, \sigma)) = \mathrm{mad}(\sigma),$$

where mad is a Mahonian statistic on permutations defined in [CSZ97] . The composition of the two Dyck tiling bijections $\mathrm{DTS}^{-1}(\mathrm{zigzag}_n,\cdot) \circ \mathrm{DTR}(\mathrm{zigzag}_n,\cdot)$ gives a bijection mapping mad to inv which is different than the one given in [CSZ97].

If we further restrict to the case where both $\lambda = \mathrm{zigzag}_n$ and all the Dyck tiles have size 1, then the Dyck tiling is determined by its upper path μ. We show that in this case, both $\mathrm{DTS}(\mathrm{zigzag}_n,\cdot)$ and $\mathrm{DTR}(\mathrm{zigzag}_n,\cdot)$ restrict to (classical) bijections from permutations avoiding the pattern 231 to Dyck paths. All proofs and detailed discussions of these results can be found in [KMPW12].

2 Dyck tiling bijections

Recall that λ/μ denotes the skew shape between the lower Dyck path λ and upper Dyck path μ. We choose coordinates so that when λ and μ are of order n, they start at $(-n, 0)$, each step is either $(+1, +1)$

or $(+1,-1)$, and the ending location is $(+n,0)$. We coordinatize a Dyck tile by the Dyck path formed from the lower corners of the boxes contained within the Dyck tile. Column $s \in \mathbb{Z}$ refers to the set of points (s,y).

Given a Dyck path ρ and a column s, we define the **spread** of ρ at s to be the Dyck path ρ' whose points are

$$\begin{aligned} &\{(x,y)+(-1,0) \ : \ (x,y)\in\rho \text{ and } x\le s\}\cup\\ &\{(x,y)+(0,1) \quad\ : \ (x,y)\in\rho \text{ and } x=s\}\cup\\ &\{(x,y)+(+1,0) \ : \ (x,y)\in\rho \text{ and } x\ge s\}. \end{aligned}$$

Notice that the spread of ρ at s makes sense whether or not Dyck path ρ overlaps column s.

We define the spread of a Dyck tile at column s by taking the spread of the Dyck path that coordinatizes the Dyck tile. We define the spread of a Dyck tiling at column s by taking the spread of the upper and lower bounding Dyck paths as well as of all the Dyck tiles within the tiling.

Given a Dyck tiling T of λ/μ, a column s is **eligible** if:

1. The upper boundary μ contains an up step that ends in column s.
2. The intersection of μ with column s is not the top corner of a Dyck tile of T containing just one box.

There is always at least one eligible column of a Dyck tiling T, since the leftmost step of μ ends at an eligible column. We define the **special column** of a Dyck tiling to be its rightmost eligible column.

We now define two growth processes on Dyck tilings, one of which is used in the DTS bijection, the other in the DTR bijection. These processes are very similar — they both involve spreading the Dyck tiling at a certain column s and adding boxes to the right of the new column. In the **strip-grow** process, we add a "broken strip" of one-box tiles from the growth site to the right boundary of the tile, so that a portion of μ is pushed up and left at a $45°$ angle; see Figures 5 and 6. In the **ribbon-grow** process, we add a "ribbon" of one-box tiles from the growth site right-wards to the special column, so that a portion of μ is pushed up; see Figures 5 and 6.

Formally, given a Dyck tiling T of order n, and a column s such that $-n \le s \le n$, we define **strip-grow**(T,s) to be the Dyck tiling formed by first spreading T at s to get T'', and then adding a NE "broken strip" of one-box tiles to each up (NE) step of the upper boundary of T from the growth site until the right boundary of T'' to obtain $T' = \text{strip-grow}(T,s)$.

Similarly, we define **ribbon-grow**(T,s) to be the Dyck tiling formed by first spreading T at s to get T'', and then if T'''s special column Q is to the right of s, adding a ribbon of one-box tiles on top of T'' at columns that are strictly between columns s and Q to obtain $T' = \text{ribbon-grow}(T,s)$. Notice that the upper boundary of tiling T'' has a down step starting at column s, and an up step ending at column Q, so adding this ribbon of one-box tiles to T'' results in a valid Dyck tiling T'.

If T is cover-inclusive, then both $\text{strip-grow}(T,s)$ and $\text{ribbon-grow}(T,s)$ are also cover-inclusive. The maps strip-grow and ribbon-grow are the **grow maps**.

Column s is eligible for T'. Because we added a strip or a ribbon to the spread of T to get T', there are no eligible columns of T' to the right of s, so s is T''s special column in both growth processes.

The bijections DTS and DTR are given by repeated application of the strip-grow and ribbon-grow maps respectively. More precisely, we first do a minor change of coordinates,

$$p_i = (i-1) + s_i,$$

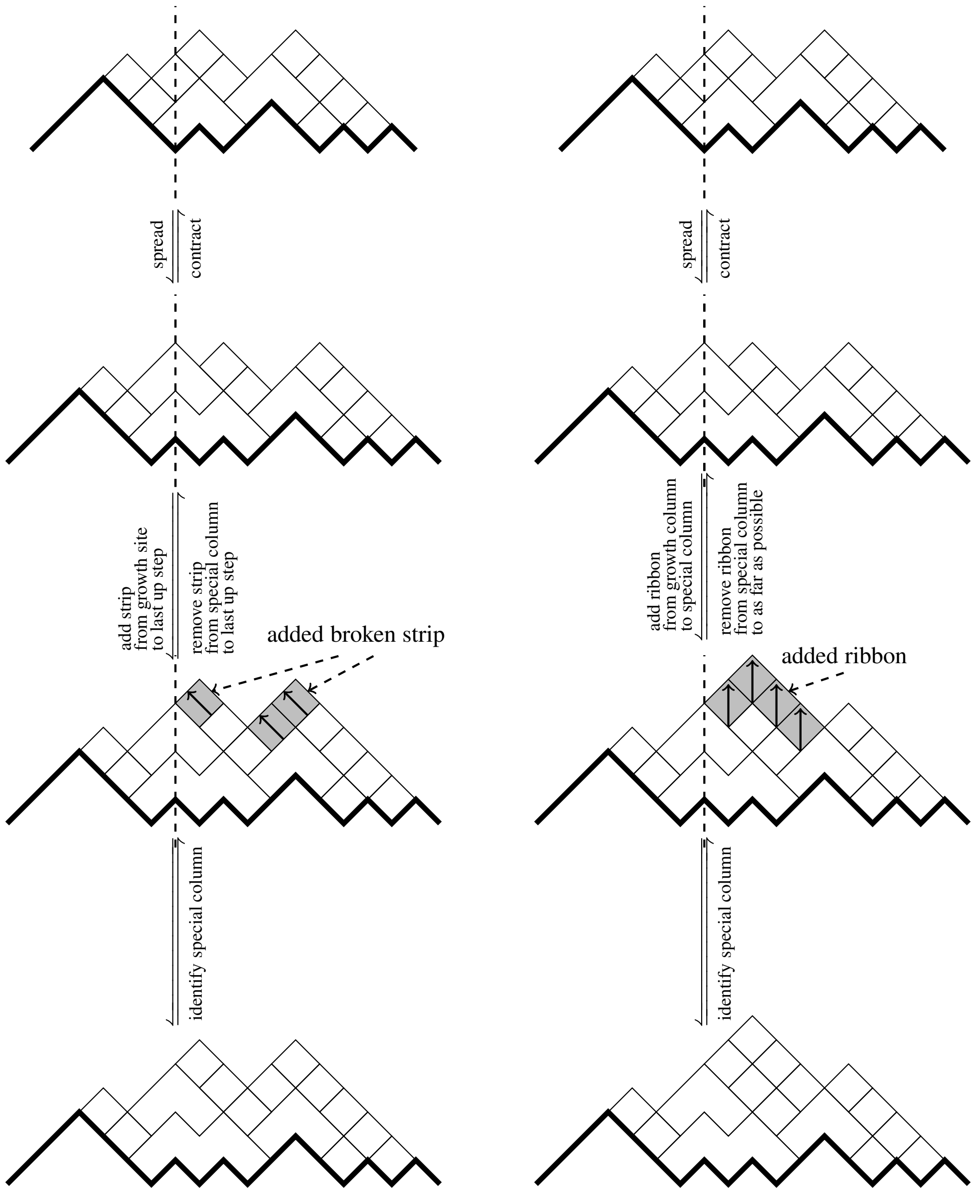

Fig. 5: The growth process for the DTS bijection (left column) and the DTR bijection (right column). The spread and contract steps are the same for both bijections, as is the definition of the special column. In the DTS bijection, a "broken strip" of one-box tiles is always added to each up step to the right of the growth site, which has the effect of pushing up-and-left the upper boundary (as indicated by the arrows). In the DTR bijection, when the growth is at a column to the left of the special column, a "ribbon" of one-box tiles is added from the growth site to the special column, which has the effect of pushing up the upper boundary of the tiling (as indicated by the arrows).

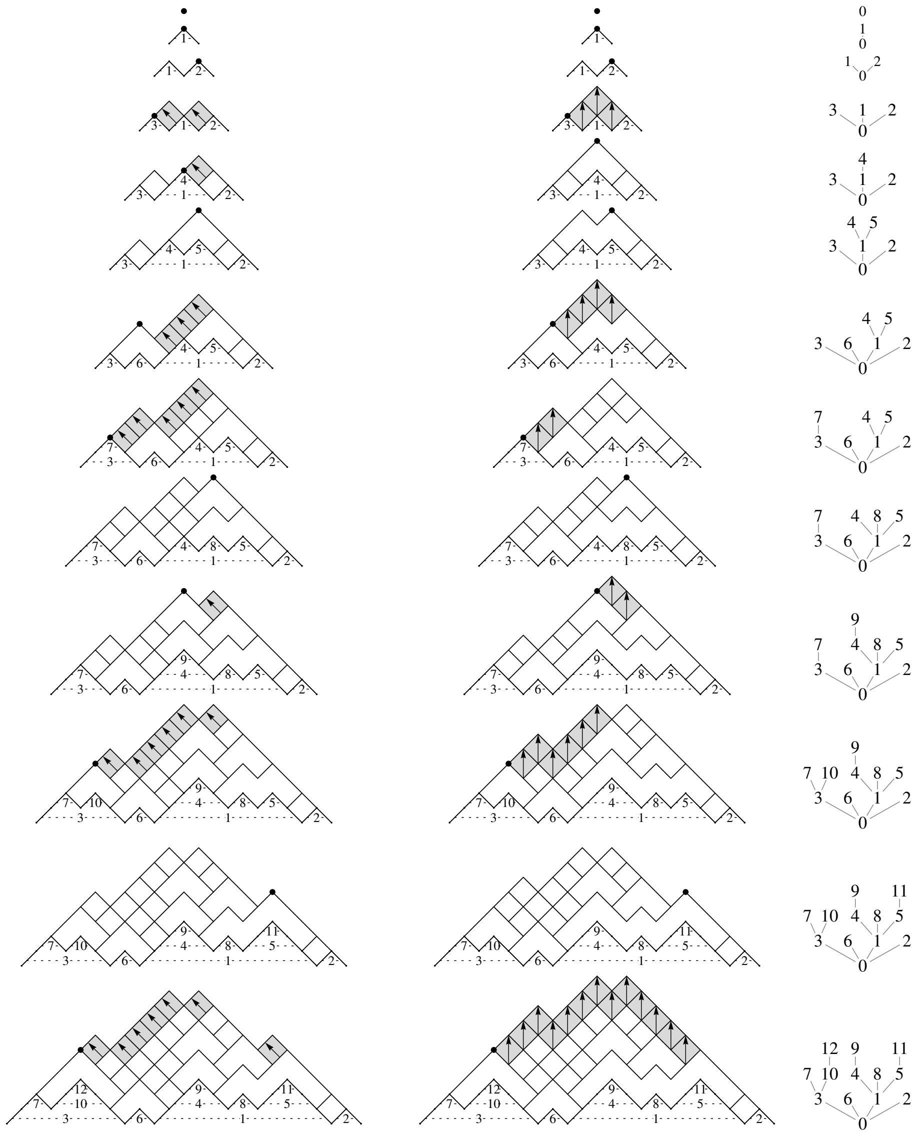

Fig. 6: Example showing the bijections from increasing plane trees (linear extensions of a chord poset) to Dyck tilings, with the DTS bijection on the left, and the DTR bijection on the right. The increasing plane tree is built up one vertex at a time in numerical order of the vertex labels, while the Dyck tilings are built up, starting from the empty Dyck tiling, by a sequence of the growth steps illustrated in Figure 5, with newly added tiles shown in gray. At each stage k, the linear extension of the chord poset $P_{\lambda^{(k)}}$ is shown together with the k^{th} Dyck tiling, whose lower boundary is $\lambda^{(k)}$. The preorder word of the final linear extension is $3, 7, 10, 12, 6, 1, 4, 9, 8, 5, 11, 2$, whose inverse is $6, 12, 1, 7, 10, 5, 2, 9, 8, 3, 11, 4$, which has 6 descents and 34 inversions. The discrepancy between the upper and lower paths of the final tiling on the right is 6, and $(\text{area} + \text{tiles})/2$ of the final tiling on the left is 34.

so that $0 \leq p_i \leq 2(i-1)$, and then define

$$\mathrm{DTS}(p_1,\dots,p_n) = \begin{cases} \varnothing & n=0 \\ \text{strip-grow}(\mathrm{DTS}(p_1,\dots,p_{n-1}), p_n-(n-1)) & n>0, \end{cases}$$

and

$$\mathrm{DTR}(p_1,\dots,p_n) = \begin{cases} \varnothing & n=0 \\ \text{ribbon-grow}(\mathrm{DTR}(p_1,\dots,p_{n-1}), p_n-(n-1)) & n>0. \end{cases}$$

Theorem 3 *The maps* DTS *and* DTR *are bijections from integer sequences $p_1,\dots,p_n$ with $0 \leq p_i \leq 2(i-1)$ to cover-inclusive Dyck tilings of order n.*

Next we compare the bijections DTS and DTR to the bijection from integer sequences $p_1,\dots,p_n$ to increasing plane oriented trees. The strip-grow and ribbon-grow maps introduce a new chord in the lower boundary of the Dyck tiling; the existing chords may be stretched, but their relative order is unchanged. (By induction, the lower boundaries of $\mathrm{DTS}(p_1,\dots,p_n)$ and $\mathrm{DTR}(p_1,\dots,p_n)$ are the same.) If we keep track of the order in which the chords are introduced, the result is a natural labeling L of the chord poset P_λ of the lower boundary λ, which together comprise an increasing plane oriented tree, as shown in Figure 6. In fact, this increasing plane oriented tree is the one given by the standard bijection from sequences $p_1,\dots,p_n$ to increasing plane oriented trees. We can represent an increasing plane oriented tree by a Dyck path λ and a permutation $\sigma \in \mathscr{L}(P_\lambda)$, as shown in Figure 2. Given such a pair (λ,σ), we define

$$\mathrm{DTS}(\lambda,\sigma) = \mathrm{DTS}(\text{sequence } p_1,\dots,p_n \text{ which yields labeled tree defined by } (\lambda,\sigma))$$

and

$$\mathrm{DTR}(\lambda,\sigma) = \mathrm{DTR}(\text{sequence } p_1,\dots,p_n \text{ which yields labeled tree defined by } (\lambda,\sigma)).$$

We therefore obtain the following theorem:

Theorem 4 *For each Dyck path λ, the maps $\mathrm{DTS}(\lambda,\cdot)$ and $\mathrm{DTR}(\lambda,\cdot)$ are bijections from linear extensions in $\mathscr{L}(P_\lambda)$ to cover-inclusive Dyck tilings with lower path λ.*

Next we compare statistics of the Dyck tiling T to statistics of the permutation σ. To do this, it is convenient to view an increasing plane oriented tree as having its edges labeled rather than its vertices, and when the tree is represented by a Dyck path λ, for the labels to reside on the up steps of λ. The preorder word (σ^{-1}) is the listing of these labels in order. Both grow maps will insert a new chord into λ at a position s. Let n be the order of Dyck path λ, and let λ' and σ' be the Dyck path and linear extension associated with $\text{strip-grow}(T,s)$ or $\text{ribbon-grow}(T,s)$. The word $(\sigma')^{-1}$ is just σ^{-1} with $n+1$ inserted at the location which is the number of up steps of λ to the left of s.

Notice that in $\text{strip-grow}(T,s)$, the art statistic increases by the number of up-steps in the lower path λ to the right of the insertion point s. This number also equals the number of elements appearing to the right of $n+1$ in $(\sigma')^{-1}$, which is $\mathrm{inv}((\sigma')^{-1}) - \mathrm{inv}(\sigma^{-1})$. Hence we arrive at the following Theorem, which also proves Theorem 1.

Theorem 5 *For each Dyck path λ and linear extension $\sigma \in \mathscr{L}(P_\lambda)$,*

$$\mathrm{art}(\mathrm{DTS}(\lambda,\sigma)) = \mathrm{inv}(\sigma).$$

Notice that spreading does not change the discrepancies between lower and upper path, and that the ribbon addition introduces exactly one extra discrepancy, right after the new column s. The ribbon is added exactly when n is a descent in $(\sigma')^{-1}$, which gives the following statement and Theorem 2.

Theorem 6 *For each Dyck path λ and linear extension $\sigma \in \mathscr{L}(P_\lambda)$,*

$$\mathrm{dis}(\mathrm{DTR}(\lambda, \sigma)) = \mathrm{des}(\sigma).$$

References

[ABDH12] Jean-Christophe Aval, Adrien Boussicault, and Sandrine Dasse-Hartaut. Dyck tableaux. *Theoretical Computer Science*, 2012. To appear. arXiv:1109.0370.

[ABN11] Jean-Christophe Aval, Adrien Boussicault, and Philippe Nadeau. Tree-like tableaux. In *23rd International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2011)*, Discrete Math. Theor. Comput. Sci. Proc., AO, pages 63–74. Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2011. arXiv:1109.0371. MR2820698

[BFS92] François Bergeron, Philippe Flajolet, and Bruno Salvy. Varieties of increasing trees. In *CAAP '92 (Rennes, 1992)*, volume 581 of *Lecture Notes in Comput. Sci.*, pages 24–48. Springer, Berlin, 1992. MR1251994 (94j:68233)

[BW89] Anders Björner and Michelle L. Wachs. q-hook length formulas for forests. *J. Combin. Theory Ser. A*, 52(2):165–187, 1989. MR1022316 (91e:05013)

[CSZ97] Robert J. Clarke, Einar Steingrímsson, and Jiang Zeng. New Euler-Mahonian statistics on permutations and words. *Adv. in Appl. Math.*, 18(3):237–270, 1997. MR1436481 MR(97m:05008)

[Kim11] Jang Soo Kim. Proofs of two conjectures of Kenyon and Wilson on Dyck tilings. 2011. arXiv:1108.5558.

[KMPW12] Jang Soo Kim, Karola Mészáros, Greta Panova, and David B. Wilson. Dyck tilings, linear extensions, and descents. 2012. arXiv:1205.6578.

[Knu73] Donald E. Knuth. *The art of computer programming. Volume 3*. Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1973. Sorting and searching, Addison-Wesley Series in Computer Science and Information Processing. MR0445948 (56 #4281)

[KW11a] Richard W. Kenyon and David B. Wilson. Double-dimer pairings and skew Young diagrams. *Electron. J. Combin.*, 18(1):Paper 130, 2011. arXiv:1007.2006. MR2811099

[KW11b] Richard W. Kenyon and David B. Wilson. Spanning trees of graphs on surfaces and the intensity of loop-erased random walk on $\mathbb{Z}^2$. 2011. arXiv:1107.3377.

[Sta72] Richard P. Stanley. *Ordered structures and partitions*. American Mathematical Society, Providence, R.I., 1972. Memoirs of the American Mathematical Society, No. 119.

[Sta99] Richard P. Stanley. *Enumerative Combinatorics, Volume 2*. Cambridge Studies in Advanced Mathematics #62. Cambridge University Press, Cambridge, 1999.