

ON LINEAR INTERVALS IN THE ALT ν -TAMARI LATTICES

Cesar Ceballos^{*1} and Clément Chenevière^{†2}

¹*Institute of Geometry, TU Graz, Austria*
cesar.ceballos@tugraz.at

²*Institut de Recherche Mathématique Avancée, Université de Strasbourg, France*
and Fakultät für Mathematik, Ruhr-Universität Bochum, Germany
and Laboratoire Interdisciplinaire des Sciences du Numérique, Université Paris-Saclay, France
clement.cheneviere@lisn.fr

Submitted: Jul 31, 2023; Accepted: Jul 3, 2024; Published: Sep 30, 2024

© The authors. Released under the CC BY license (International 4.0).

Abstract. Given a lattice path ν , the ν -Tamari lattice and the ν -Dyck lattice are two natural examples of partial order structures on the set of lattice paths that lie weakly above ν . In this paper, we introduce a more general family of lattices, called alt ν -Tamari lattices, which contains these two examples as particular cases. Unexpectedly, we show that all these lattices have the same number of linear intervals.

Keywords. Lattices, intervals, Tamari

Mathematics Subject Classifications. 06A07, 06B05, 05A19

1. Introduction

The classical Tamari lattice is a partial order on Catalan objects which has inspired a vast amount of research in various mathematical fields [MHPS12]. One direction of research which has received a lot of attention in recent years regards its number of intervals [Cha07], which is conjectured to be equal to the dimension of the alternating component of an \mathfrak{S}_n -module in the study of trivariate diagonal harmonics [Hai94]. Motivated by this intriguing connection, Bergeron and Préville-Ratelle introduced a generalization of the Tamari lattice called the m -Tamari lattice, and conjectured that its number of intervals again coincides with the dimension of the alternating component of an \mathfrak{S}_n -module in higher trivariate diagonal harmonics [BPR12]. A formula for their enumeration and connections to representation theory can be found in [BMCP13, BMFPR11]. A further generalization of the Tamari lattice, which includes the m -Tamari lattice, is the ν -Tamari lattice introduced by Préville-Ratelle and Viennot [PRV17]. These lattices are indexed by a lattice

^{*}Supported by the Austrian Science Fund FWF, grants P 33278 and I 5788.

[†]Supported by the ANR Projects ANR-21-CE48-0020 and ANR-19-CE48-0011.

path ν , and their number of intervals is connected to the enumeration of non-separable planar maps as shown in [FPR17].

Inspired by the enumeration of intervals in the classical Tamari lattice and its generalizations, and guided by computer experimentation, Chapoton proposed to study the enumeration of the simpler class of linear intervals (intervals which are chains). This led to the work of the second author in [Che22], where he provides an explicit simple formula for the number of linear intervals in the classical Tamari lattice, and shows that their enumeration coincides with the enumeration of linear intervals in the Dyck lattice. The Dyck lattice, sometimes called the Stanley lattice, is perhaps the most natural poset on Dyck paths, defined by $P \leq Q$ if Q is weakly above P . In [Che22], the author also defines a new family of posets called alt Tamari posets, which contain the Tamari lattice and the Dyck lattice as particular cases. He shows that all alt Tamari posets have the same number of linear intervals of any given length.

In this paper, we generalize the results in [Che22] by introducing a new family of posets called alt ν -Tamari posets. We show that they are lattices, and that they all have the same number of linear intervals of any given length. Figure 1.1 and Figure 1.2 illustrate the three different alt ν -Tamari lattices for $\nu = ENEENN$. In each case, the number of linear intervals of length k is given by ℓ_k where $\ell = (\ell_0, \ell_1, \ell_2, \ell_3) = (16, 24, 16, 3)$. For instance, 16 represents the trivial intervals of length 0, which are just the elements of each poset; there are 24 linear intervals of length 1, which correspond to the cover relations (edges in the figures); there are 16 linear intervals of length 2, and 3 linear intervals of length 3. The fact that these numbers coincide is somewhat surprising, since the posets look quite different. As a warm up exercise, the reader is invited to find the 3 linear intervals of length 3 in each of the figures. Figure 2.1 illustrates the three different alt ν -Tamari lattices for $\nu = ENEEN$.

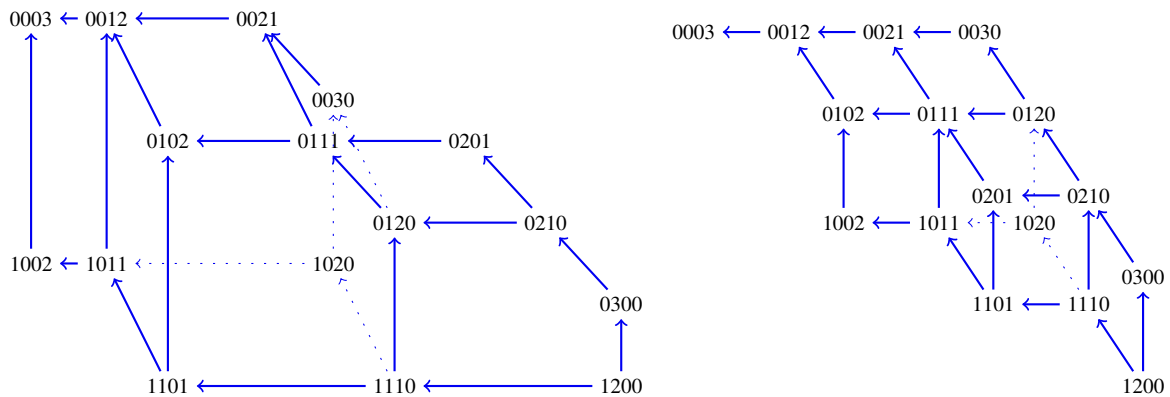


Figure 1.1: The ν -Tamari lattice and ν -Dyck lattice for $\nu = ENEENN$. They are the alt ν -Tamari lattices $\text{Tam}_\nu(\delta)$ for $\delta = (2, 0, 0)$ and $\delta = (0, 0, 0)$, respectively.

As the figures suggest, the alt ν -Tamari lattices possess a rich underlying geometric structure, which seems to be realizable as a polytopal complex in some Euclidean space. This was shown to be true for ν -Tamari lattices in [CPS19], where polytopal complex realizations induced by some arrangements of tropical hyperplanes are provided. The first author presented similar geometric realizations in this general context for alt ν -Tamari lattices [Ceb24].

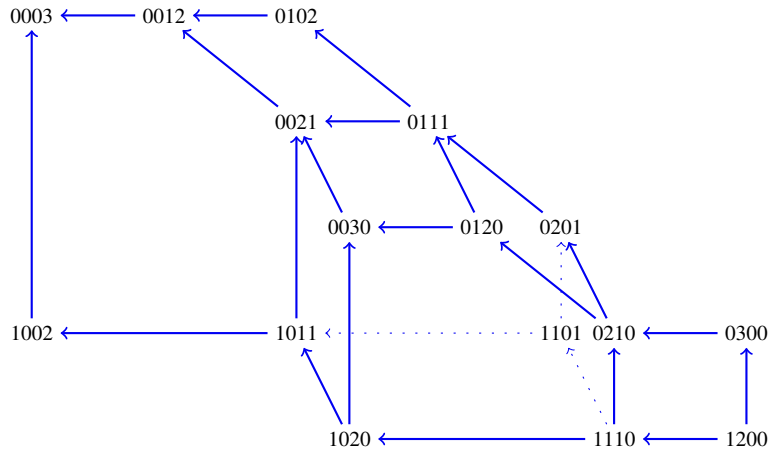


Figure 1.2: The alt ν -Tamari lattice $\text{Tam}_\nu(\delta)$ for $\nu = ENEENN$ and $\delta = (1, 0, 0)$.

2. The ν -Dyck lattice

Let ν be a lattice path on the plane from $(0, 0)$ to (m, n) , consisting of a finite number of north and east unit steps. We may represent a path ν as a word in the letters E and N for east and north steps respectively. We may as well represent ν as a sequence of non-negative integers $(\nu_0, \nu_1, \dots, \nu_n)$, where $n \in \mathbb{N}$ is the number of north steps of ν , ν_0 is the number of initial east steps, and $\nu_i \geq 0$ is the number of consecutive east steps immediately following the i -th north step of ν . In particular, $m = \nu_0 + \dots + \nu_n$ is the total number of east steps. For instance, the path $ENEENNENEENE$ would correspond to the sequence $(1, 2, 0, 1, 3)$, while $ENEENN$ corresponds to $(1, 2, 0, 0)$.

A ν -path μ is a lattice path using north and east steps, with the same endpoints as ν , that is weakly above ν . Alternatively, $\mu = (\mu_0, \dots, \mu_n)$ is a ν -path if and only if $\sum_{i=0}^j \mu_i \leq \sum_{i=0}^j \nu_i$ for all $0 \leq j \leq n$, with equality for $j = n$. The elements of the posets in Figure 1.1 and Figure 1.2 are labelled by ν -paths using this representation, where we omit the commas and parentheses for simplicity. For instance, the label 1200 is the minimal path $(1, 2, 0, 0)$, which corresponds to $\nu = ENEENN$.

Definition 2.1. The ν -Dyck lattice Dyck_ν is the poset on ν -paths where $P \leq Q$ if Q is weakly above P .

An example of the ν -Dyck lattice for $\nu = ENEEN$ is illustrated on the left of Figure 2.1.

Remark 2.2. The case where ν is $(NE)^n$ coincides with the classical Dyck lattice on Dyck paths of size n .

Remark 2.3. Covering relations $P \lessdot Q$ in the ν -Dyck lattice consist of transforming a valley EN in a peak NE in some ν -path P .

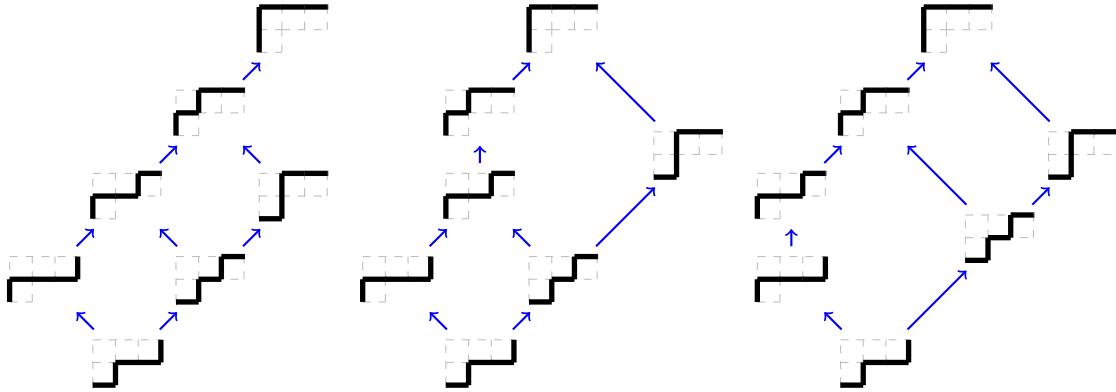


Figure 2.1: Examples of alt ν -Tamari lattices $\text{Tam}_\nu(\delta)$ for $\nu = ENEEN = (1, 2, 0)$. Left: the ν -Dyck lattice, for $\delta = (0, 0)$. Middle: the lattice for $\delta = (1, 0)$. Right: the ν -Tamari lattice, for $\delta = (2, 0)$. In each case, the number of linear intervals of length k is given by ℓ_k where $\ell = (\ell_0, \ell_1, \ell_2, \ell_3) = (7, 8, 4, 1)$. For instance, 7 represents the trivial intervals of length 0, which are just the elements of each poset; there are 8 linear intervals of length 1, which correspond to the edges; 4 linear interval of length 2, and 1 linear interval of length 3.

2.1. Left and right intervals in the ν -Dyck lattice

We focus on the special class of linear intervals of a poset. An interval $[P, Q]$ is *linear* if it is totally ordered, or equivalently if it is a chain of the form $P = P_0 < P_1 < \dots < P_\ell = Q$. The *length* of such a linear interval is defined to be ℓ . A linear interval of length zero (containing only one element) is said to be *trivial*. A linear interval of length one is by definition a covering relation since it contains two elements. The non-trivial linear intervals of the ν -Dyck lattice can be easily characterized into two different classes.

Definition 2.4. An interval $[P, Q]$ in Dyck_ν is a *left interval* if Q is obtained from P by transforming a subpath $E^\ell N$ into NE^ℓ for some $\ell \geq 1$. It is a *right interval* if Q is obtained from P by transforming a subpath EN^ℓ into $N^\ell E$ for some $\ell \geq 1$.

Lemma 2.5. *The linear intervals of length 2 are either left or right intervals.*

Proof. Let $P < Q < R$ be a linear interval of length 2. The covering relations $P < Q$ transforms a valley EN of P into a peak NE . If the next covering relation $Q < R$ happens at a valley of Q that is also a valley of P , then the $[P, R]$ is a square. Thus, this second covering relation must use either of the two steps of the peak NE that was created in Q . \square

Proposition 2.6. *The left and right intervals in the previous definition are linear intervals of length ℓ . Moreover, all non-trivial linear intervals in Dyck_ν are either left or right intervals.*

Proof. If $[P, Q]$ is an interval of this form with $\ell \geq 1$, then there exists only one maximal chain from P to Q . Indeed, there is only one valley of P that is not a valley of Q and thus, any maximal chain from P to Q starts at this valley. We then obtain an interval of the same form, but ℓ has decreased by 1 and we conclude by induction.

Lemma 2.5 proves that all linear intervals of height $k = 2$ are either left or right intervals. Suppose that $[P, Q]$ is a linear interval of height $k + 1 \geq 3$. It is linear so Q has only one lower cover Q' in $[P, Q]$. Then $[P, Q']$ is linear of height k and thus by induction, it is of the prescribed form.

Suppose that Q' is obtained from P by transforming a subpath $E^k N$ into NE^k , which creates this peak NE in Q' followed by $k - 1$ east steps. Then, Lemma 2.5 ensures that the covering relation $Q' \leq Q$ has to use the north step N of this peak and thus, Q is obtained from P by changing a subpath $E^{k+1}N$ into NE^{k+1} .

Suppose now that Q' is obtained from P by transforming a subpath EN^k into $N^k E$. Similarly, Lemma 2.5 ensures that $Q' \leq Q$ has to use the east step E of this peak and Q is obtained from P by transforming a subpath EN^{k+1} into $N^{k+1}E$. □

Corollary 2.7. *Left intervals of length ℓ in Dyck_ν are in bijection with ν -paths marked at a north step preceded by ℓ east steps. Right intervals of length ℓ in Dyck_ν are in bijection with ν -paths marked at an east step followed by ℓ north steps.*

3. The ν -Tamari lattice

The ν -Tamari lattices are a generalization of the Tamari lattice. They were defined in terms of ν -paths by Préville-Ratelle and Viennot in [PRV17]. An alternative description in terms of ν -trees was presented in [CPS20].

3.1. On ν -paths

For a lattice point p on a ν -path μ , define its ν -altitude $\text{alt}_\nu(p)$ to be the maximum number of horizontal steps that can be added to the right of p without crossing ν . Given a valley EN of μ , let p be the lattice point between the east and north steps. Let q be the next lattice point of μ such that $\text{alt}_\nu(q) = \text{alt}_\nu(p)$, and $\mu_{[p,q]}$ be the subpath of μ that starts at p and ends at q . Let μ' be the path obtained from μ by switching $\mu_{[p,q]}$ with the east step E that precedes it. The ν -rotation of μ at the valley p is defined to be $\mu \leq_\nu \mu'$. An example is illustrated in Figure 3.1.

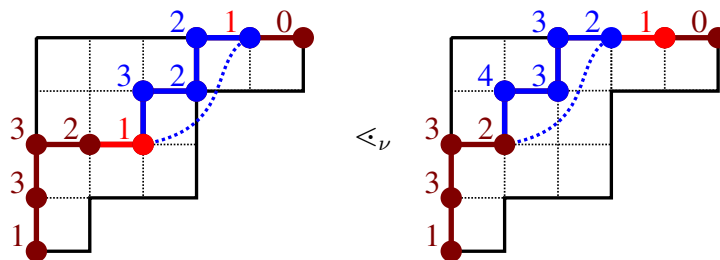


Figure 3.1: The rotation operation of a ν -path. Each node is labelled with its ν -altitude.

Definition 3.1. The ν -Tamari poset Tam_ν is the reflexive transitive closure of ν -rotations on ν -paths.

An example of the ν -Tamari poset for $\nu = ENEEN$ is illustrated on the right of Figure 2.1.

Theorem 3.2 ([PRV17, Theorem 1]). *The ν -Tamari poset is a lattice. The ν -rotations are exactly its covering relations.*

Another approach to define the ν -Tamari lattice is to introduce the notion of ν -elevation of a subpath as the difference of ν -altitude between its ending point and its starting point. We thus write $\text{elev}_\nu(E) = -1$ for an east step E and $\text{elev}_\nu(N_i) = \nu_i$ if N_i is the i -th north step of a ν -path μ . For any subpath A of μ , we then have $\text{elev}_\nu(A) = \sum_{a \in A} \text{elev}_\nu(a)$ as the sum of the ν -elevation of the steps of A .

The ν -excursion of a north step N of a ν -path μ is defined as the shortest subpath A of μ that starts with this N and such that $\text{elev}_\nu(A) = 0$. It follows from the definition of the ν -excursion that exchanging the east step E of a valley with the ν -excursion that follows it is exactly a covering relation in Tam_ν .

3.2. On ν -trees

One can also define a poset on ν -trees which is isomorphic to the ν -Tamari lattice.

We denote by F_ν the Ferrers diagram that lies weakly above ν in the smallest rectangle containing ν . Let L_ν denote the set of lattice points inside F_ν . We say that two points $p, q \in L_\nu$ are ν -incompatible if p is strictly southwest or strictly northeast of q , and the smallest rectangle containing p and q lies entirely in F_ν . Otherwise, p and q are said to be ν -compatible. A ν -tree is a maximal collection of pairwise ν -compatible elements in L_ν . In particular, the vertex at the top-left corner of F_ν is ν -compatible with everyone else, and belongs to every ν -tree. Connecting two consecutive elements (not necessarily at distance 1) in the same row or column allows us to visualize ν -trees as classical rooted binary trees [CPS20]. The vertex at top-left corner of F_ν is always the *root*. An example of a ν -tree and the rotation operation which we now describe is shown in Figure 3.2.

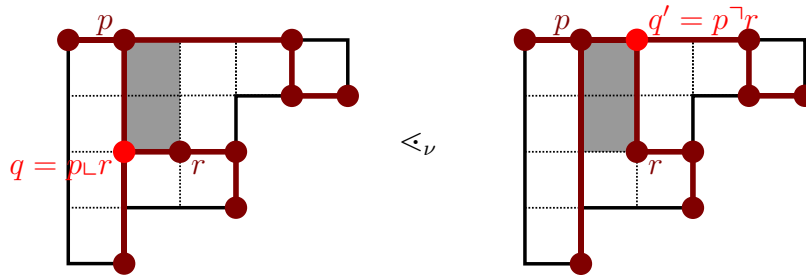


Figure 3.2: The rotation operation of a ν -tree.

Let T be a ν -tree and $p, r \in T$ be two elements which do not lie in the same row or same column. We denote by $p \square r$ the smallest rectangle containing p and r , and write $p \perp r$ (resp. $p \top r$) for the lower left corner (resp. upper right corner) of $p \square r$.

Let $p, q, r \in T$ be such that $q = p \perp r$ and no other elements besides p, q, r lie in $p \square r$. The ν -rotation of T at q is defined as the set $T' = (T \setminus \{q\}) \cup \{q'\}$, where $q' = p \top r$, and we write $T \leq_\nu T'$. As proven in [CPS20, Lemma 2.10], the rotation of a ν -tree is also a ν -tree.

Definition 3.3. The rotation poset of ν -trees Tam_ν^{tr} is the reflexive transitive closure of ν -rotations.

An example of the rotation poset of ν -trees for $\nu = ENEEN$ is illustrated on the right of Figure 3.3.

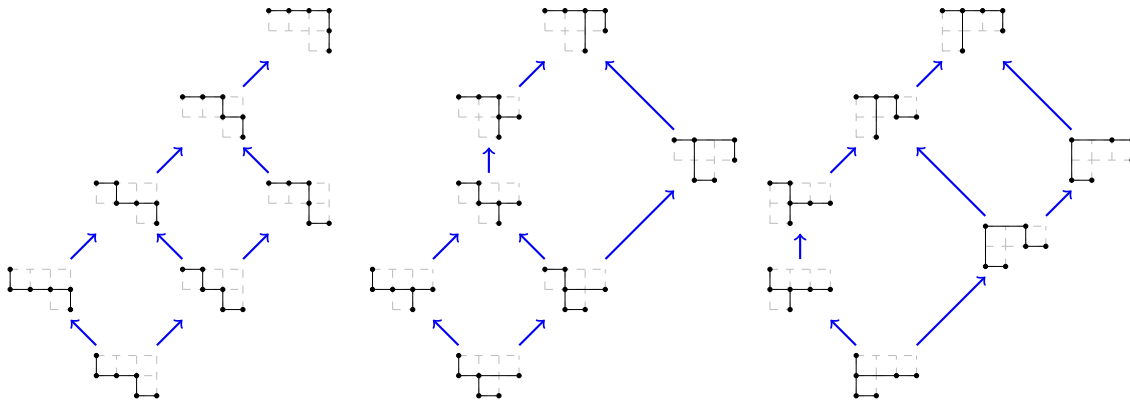


Figure 3.3: Examples of rotation lattices of (δ, ν) -trees for $\nu = ENEEN$. Left: the ν -Dyck lattice, for $\delta = (0, 0)$. Middle: the lattice for $\delta = (1, 0)$. Right: the ν -Tamari lattice, for $\delta = (2, 0)$.

Theorem 3.4 ([CPS20]). The ν -Tamari lattice is isomorphic to the rotation poset of ν -trees:

$$\text{Tam}_\nu \cong \text{Tam}_\nu^{tr}.$$

In particular, the rotation poset of ν -trees is a lattice.

A bijection between these two posets is given by the *right flushing bijection* flush_ν introduced in [CPS20]. This bijection maps a ν -path $\mu = (\mu_0, \dots, \mu_n)$ to the unique ν -tree with $\mu_i + 1$ nodes at height i . This tree can be recursively obtained by adding $\mu_i + 1$ nodes at height i from bottom to top, from right to left, avoiding forbidden positions. The forbidden positions are those above a node that is not the left most node in a row (these come from the initial points of the east steps in the path μ). In Figure 3.4, the forbidden positions are the ones that belong to the zigzag lines. Note that the order of the nodes per row is reversed.

The inverse flush_ν^{-1} of the right flushing bijection is called the *left flushing bijection*, and can be described similarly, adding points from left to right, from bottom to top, avoiding the forbidden position given by the zigzag lines. In other words, the left flushing bijection of a ν -tree is the ν -path that has as many nodes per row as the tree.

3.3. Left and right intervals in the ν -Tamari lattice

The description of the ν -Tamari lattice on ν -trees gives an easy description of its linear intervals.

Definition 3.5. An interval $[T, T']$ in Tam_ν^{tr} is a *left interval* if T' is obtained from T by applying $\ell > 0$ rotations at the first ℓ nodes of a consecutive sequence $q_0, \dots, q_{\ell-1}, q_\ell$ in the

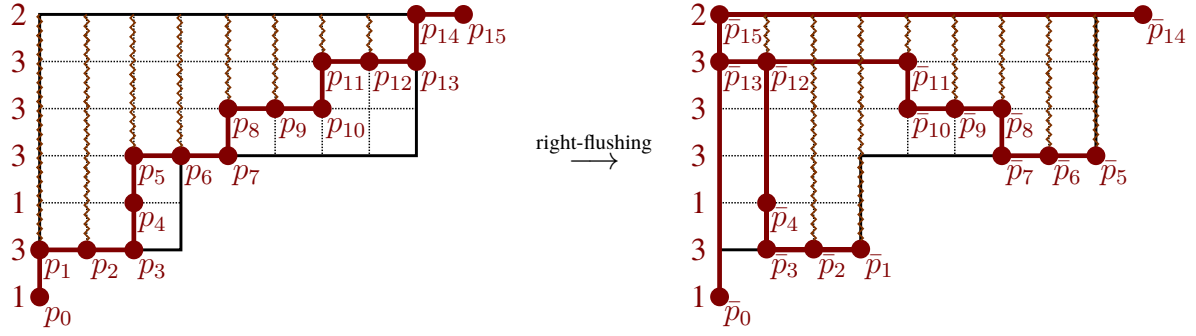


Figure 3.4: Right flushing bijection from ν -paths to ν -trees.

same row, from left to right. For example, applying two rotations at the first two nodes of the sequence $\bar{p}_{13}, \bar{p}_{12}, \bar{p}_{11}$ in Figure 3.4 (right). It is a *right interval* if T' is obtained from T by applying ℓ rotations at the first ℓ nodes of a consecutive sequence $q_0, \dots, q_{\ell-1}, q_\ell$ which are consecutive in the same column, from bottom to top. For example, applying two rotations at the first two nodes of the sequence $\bar{p}_3, \bar{p}_4, \bar{p}_{12}$ in Figure 3.4 (right).

Proposition 3.6. *The left and right intervals in the previous definition are linear intervals of length ℓ . Moreover, all non-trivial linear intervals in the rotation lattice on ν -trees are either left or right intervals.*

Proof. A left (resp. right) interval $[T, T']$ is indeed linear because there is a unique maximal chain from T to T' and ℓ is its length. Indeed, each element different from T' in the interval has only one upper cover that is below T' .

The converse is proven by induction, similarly as the simpler case of the Dyck lattice in Proposition 2.6.

Intervals of length 1 are all covering relations and thus both left and right intervals.

An interval which contains a maximal chain $T_0 \triangleleft T_1 \triangleleft T_2$ can be:

- a left interval when $T_1 \triangleleft T_2$ is the rotation in the same row immediately on the right of $T_0 \triangleleft T_1$,
- a right interval when $T_1 \triangleleft T_2$ is the rotation in the same column immediately above $T_0 \triangleleft T_1$,
- a pentagon when $T_1 \triangleleft T_2$ is the rotation in the same column immediately under $T_0 \triangleleft T_1$,
- a square otherwise.

In particular, this proves that all linear intervals of length 2 are exactly left or right intervals.

A linear interval $[T, T']$ of length $k \geq 3$ must contain a linear interval $[T, T'']$ of length $k - 1$. By induction, $[T, T'']$ must be either a left interval and in this case $[T, T']$ is a left interval as well or a right interval and in this case $[T, T']$ is also a right interval. \square

Remark 3.7. The left flushing of a left interval on the rotation lattice of ν -trees produces a left interval $[P, Q]$ of ν -paths in Tam_ν , where P is of the form AE^kBC with B some ν -excursion

Definition 4.1. Let δ be an increment vector with respect to ν . The *alt ν -Tamari poset* $\text{Tam}_\nu(\delta)$ is the reflexive transitive closure of δ -rotations on the set of ν -paths.

The three examples of alt ν -Tamari posets for $\nu = ENEEN = (1, 2, 0)$ are illustrated on Figure 2.1.

Remark 4.2. For a fixed path ν , there are two extreme choices for the increment vector δ . If $\delta_i = \nu_i$ for all $1 \leq i \leq n$, the alt ν -Tamari lattice coincides with the ν -Tamari lattice. If $\delta_i = 0$ for all $1 \leq i \leq n$, the alt ν -Tamari lattice coincides with the ν -Dyck lattice. We denote these two cases by δ^{\max} and δ^{\min} , respectively.

Another approach to define the alt ν -Tamari poset is to introduce the notion of δ -elevation of a subpath as the difference of the δ -altitude between its ending point and its starting point. We thus write $\text{elev}_\delta(E) = -1$ for an east step E and $\text{elev}_\delta(N_i) = \delta_i$ if N_i is the i -th north step of a ν -path μ . For any subpath A of μ , we then have $\text{elev}_\delta(A) = \sum_{a \in A} \text{elev}_\delta(a)$ as the sum of the δ -elevation of the steps of A .

The δ -excursion of a north step N of a ν -path μ is defined as the shortest subpath A of μ that starts with this N and such that $\text{elev}_\delta(A) = 0$. It follows from the definition of the δ -excursion that exchanging the east step E of a valley with the δ -excursion that follows it is exactly a covering relation in $\text{Tam}_\nu(\delta)$.

Remark 4.3. Note that for $\delta = \delta^{\max}$, the δ -altitude is the ν -altitude shifted by $-\nu_0$, but the δ -elevation and the ν -elevation are equal.

For a general increment vector δ with respect to ν , it is not a priori clear that $\text{Tam}_\nu(\delta)$ is a lattice. This is a consequence of the following proposition.

Proposition 4.4. Let $\check{\nu}_0 = \sum_{i=0}^n \nu_i - \sum_{i=1}^n \delta_i$ with $\delta_i \leq \nu_i$. Then $\check{\nu} = (\check{\nu}_0, \check{\nu}_1, \dots, \check{\nu}_n) = (\check{\nu}_0, \delta_1, \dots, \delta_n)$ is a path below ν whose endpoints are the same as ν . Moreover, the following properties hold:

1. δ -rotations of a ν -path μ coincide with $\check{\nu}$ -rotations of μ .
2. The alt ν -Tamari poset $\text{Tam}_\nu(\delta)$ is the restriction of $\text{Tam}_{\check{\nu}}$ to the subset of paths weakly above ν .
3. The covering relations of $\text{Tam}_\nu(\delta)$ are exactly the δ -rotations.
4. The alt ν -Tamari poset $\text{Tam}_\nu(\delta)$ is the interval $[\nu, 1^\nu]$ in $\text{Tam}_{\check{\nu}}$.

Here, $1^\nu = N^n E^m$ denotes the top path above ν and $\check{\nu}$, where $m = \nu_0 + \dots + \nu_n = \check{\nu}_0 + \dots + \check{\nu}_n$.

Proof. The first part of the statement follows from $\sum_{i=0}^j \check{\nu}_i - \sum_{i=0}^j \nu_i = \sum_{i=j+1}^n (\nu_i - \delta_i) \geq 0$, where equality holds for $j = n$.

Note that a ν -path μ is also a $\check{\nu}$ -path, and for a subpath A of μ we have $\text{elev}_\delta(A) = \text{elev}_{\check{\nu}}(A)$. Since the δ -rotations (resp. $\check{\nu}$ -rotations) are determined by the δ -elevation (resp. $\check{\nu}$ -elevation), then Item (1) follows. Items (2) and (3) follow from Item (1).

For Item (4) we need to show that the restriction of $\text{Tam}_{\check{\nu}}$ to the subset of paths weakly above ν is the interval $[\nu, 1^\nu]$ in $\text{Tam}_{\check{\nu}}$. In other words, we need to show that every ν -path μ

satisfies $\nu \leq_{\text{Tam}_{\check{\nu}}} \mu$. Note that this property does not hold for an arbitrary path $\check{\nu}$ below ν , but for our particular choice this is equivalent to show that $\nu \leq_{\text{Tam}_{\nu}(\delta)} \mu$ (by Item (1)). This holds because we can reach any ν -path μ by applying a sequence of δ -rotations: add the boxes between ν and μ one at a time from bottom to top, from right to left. Each of these steps corresponds to a δ -rotation because $\delta_i \leq \nu_i$. \square

Corollary 4.5. *The alt ν -Tamari poset is a lattice.*

Proof. By Proposition 4.4 (4), the alt ν -Tamari poset $\text{Tam}_{\nu}(\delta)$ is isomorphic to the interval $[\nu, 1^{\nu}]$ in $\text{Tam}_{\check{\nu}}$. Since an interval in a lattice is also a lattice, we deduce that $\text{Tam}_{\nu}(\delta)$ is a lattice. \square

Remark 4.6. If we chose any other path $\check{\nu}$ weakly below ν that does not satisfy $\check{\nu}_i \leq \nu_i$, for all $i > 0$, then the restriction of $\text{Tam}_{\check{\nu}}$ to the subset of ν -paths is not a lattice. It is an upper set but it has several minimal elements. The results on the number of linear intervals that are presented in the rest of the paper do not hold either with this weaker condition (see Remark 6.2).

We do not use the following proposition in this paper, but it is an interesting property that we would like to highlight.

Proposition 4.7. *Let δ and δ' be two increment vectors with respect to ν such that $\delta_i \leq \delta'_i$ for all i . If $P < Q$ in $\text{Tam}_{\nu}(\delta')$, then $P < Q$ in $\text{Tam}_{\nu}(\delta)$.*

In other words, whenever $\delta \leq \delta'$, the poset $\text{Tam}_{\nu}(\delta)$ is an *extension* of the poset $\text{Tam}_{\nu}(\delta')$, meaning that it can be obtained from $\text{Tam}_{\nu}(\delta')$ by adding some relations.

Proof. It is sufficient to prove that the result is true for covering relations in $\text{Tam}_{\nu}(\delta')$, namely that if $P \lessdot_{\delta'} Q$, then $P < Q$ in $\text{Tam}_{\nu}(\delta)$.

We can write $P = AEBC$ and $Q = ABEC$ for some δ' -excursion B . Note that for any north step N in B , the δ' excursion of N is a subword of B and that the δ -excursion of any north step is a prefix of its δ' -excursion since we have $\delta \leq \delta'$. Thus, we can build a chain of δ -rotations from P to Q by exchanging this east step E with either the next δ -excursions if it is followed by a north step or with the east step that follows it otherwise, which does not change the path. \square

4.2. On (δ, ν) -trees

The alt ν -Tamari lattice $\text{Tam}_{\nu}(\delta)$ is the interval $[\nu, 1^{\nu}]$ in $\text{Tam}_{\check{\nu}}$. So, it can be described as the rotation lattice of $\check{\nu}$ -trees that are above the $\check{\nu}$ -tree T_{ν} corresponding to ν in $\text{Tam}_{\check{\nu}}$. These trees can be described as maximal collections of pairwise compatible elements in a shape $F_{\delta, \nu}$ which we will now describe. This point of view is useful to show that all alt ν -Tamari lattices have the same number of linear intervals of any length.

Let δ, ν and $\check{\nu}$ as in Proposition 4.4. Let $F_{\check{\nu}}$ be the Ferrers diagram that lies weakly above $\check{\nu}$. We consider the lattice path $\hat{\nu}$ that starts at the lowest right corner of $F_{\check{\nu}}$ (the point with coordinates $(\check{\nu}_0, 0)$) which consists of the sequence of west and north steps

$$W^{\nu_0} N W^{\gamma_1} N W^{\gamma_2} \dots N W^{\gamma_n}, \quad \text{for } \gamma_i = \nu_i - \delta_i. \tag{4.1}$$

We define $F_{\delta, \nu}$ to be the subset of $F_{\check{\nu}}$ consisting of the boxes that are weakly above $\hat{\nu}$, and denote by $L_{\delta, \nu}$ its set of lattice points. A (δ, ν) -tree is a maximal collection of pairwise $\check{\nu}$ -compatible elements in $L_{\delta, \nu}$. An example is illustrated at the bottom of Figure 4.2.

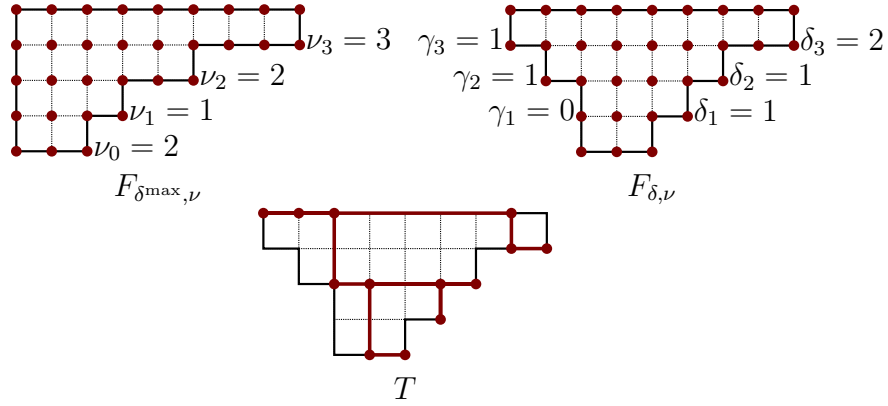


Figure 4.2: Top left: The Ferrers diagram $F_{\delta^{\max}, \nu}$ and its corresponding lattice points $L_{\delta^{\max}, \nu}$ for $\nu = EENENEENEEN = (2, 1, 2, 3, 0)$ and $\delta^{\max} = (1, 2, 3, 0)$. Top right: The Ferrers diagram $F_{\delta, \nu}$ and its corresponding lattice points $L_{\delta, \nu}$ for the same ν and $\delta = (1, 1, 2, 0)$; the path $\tilde{\nu} = EEEENENENEEN$ and $\hat{\nu} = WWNNWNWN$. Bottom: a (δ, ν) -tree.

Lemma 4.8. *The (δ, ν) -trees are exactly the $\tilde{\nu}$ -trees that are contained in $L_{\delta, \nu}$.*

Proof. A $\tilde{\nu}$ -tree that is contained in $L_{\delta, \nu}$ is automatically a (δ, ν) -tree by definition. So, we just need to check that (δ, ν) -trees are $\tilde{\nu}$ -trees. By definition, a (δ, ν) -tree is a maximal collection of pairwise $\tilde{\nu}$ -compatible elements in $L_{\delta, \nu}$. As $L_{\delta, \nu} \subseteq L_{\tilde{\nu}}$, we want to prove that the maximality in $L_{\delta, \nu}$ implies the maximality in $L_{\tilde{\nu}}$.

Recall that the paths ν and $\tilde{\nu}$ have the same starting point $(0, 0)$ and the same ending point (m, n) , and that every ν -tree and every $\tilde{\nu}$ -tree has exactly $m + n + 1$ nodes (equal to the number of lattice points in ν and $\tilde{\nu}$). Furthermore, the shape $F_{\delta, \nu}$ fits in the $m \times n$ box with the top corners being $(0, n)$ and (m, n) . In our example in Figure 6.1, $m = 11$ and $n = 7$. The ν -tree and the (δ, ν) -tree shown in this figure both have $m + n + 1 = 19$ nodes. We want to show that every (δ, ν) -tree has exactly $m + n + 1$ elements.

Let T be a (δ, ν) -tree. Label its elements p_0, p_1, \dots, p_r from bottom to top, from right to left. We will show that $r = m + n$, which implies that T has $m + n + 1$ elements as desired.

Let us reconstruct T recursively, by adding the elements p_0, p_1, \dots, p_r one at a time in order. Note that if p_i is not the left most element in its row, then all the lattice points above p_i are forbidden in the next steps, because they are incompatible with an element $p_j \in T$ that is to the left of p_i in the same row.

Now, when we add an element p_j in the process of reconstructing T , then p_j is necessarily located at the right most position of its row that is not forbidden by any element before. Otherwise, let $p_j \in T$ be the node with smallest label that does not satisfy that property, and let q be the right most lattice point in the same row that is not forbidden by any element p_i with $i < j$. In particular, q is on the right of p_j by assumption, and q is compatible with every p_i with $i < j$. Moreover, q is also compatible with $p_k \in T$ with $k > j$, otherwise p_k, p_j would be incompatible. So, we can add the element q to T , creating a new compatible set, contradicting the maximality of T .

Furthermore, the following relation holds,

$$j = \text{forb}(p_j) + \text{height}(p_j)$$

where $\text{height}(p_j)$ is the height of $p_j \in T$ and $\text{forb}(p_j)$ is the number of $p_i \in T$ with $i < j$ such that p_i is not the left most node of T in its row. That is, $\text{forb}(p_j)$ is the number of nodes before p_j that forbid the positions above them. This formula is clear because the j nodes p_0, \dots, p_{j-1} appearing before p_j either forbid positions above them (not the left most node of their row) or increase the height by one (the left most node of their row).

If we apply this formula to the last node p_r and assume that $r < m + n$, then

$$\text{forb}(p_r) + \text{height}(p_r) < m + n.$$

We can assume that $\text{height}(p_r) = n$ (maximum possible height), otherwise we could add the top left corner of $F_{\delta, \nu}$ to T , creating a bigger compatible set and contradicting the maximality of T . This implies that $\text{forb}(p_r) < m$, which means that on the top row there are still some lattice points that are not forbidden. Adding one of these points contradicts the maximality of T . As a consequence, we have proven that $r = m + n$ as desired. \square

We define the (δ, ν) -right flushing $\text{flush}_{\delta, \nu}$ as the restriction of the right flushing bijection $\text{flush}_{\check{\nu}}$ (with respect to $\check{\nu}$) to set of ν -paths (thought as the subset of $\check{\nu}$ -paths that are above ν).

Proposition 4.9. *The map $\text{flush}_{\delta, \nu}$ is a bijection between the set of ν -paths and the set of (δ, ν) -trees. Moreover, two ν -paths are related by a δ -rotation $\mu \leq_{\delta} \mu'$ if and only if the corresponding trees are related by a $\check{\nu}$ -rotation $T \leq_{\check{\nu}} T'$.*

Proof. By Proposition 4.4, δ -rotations of a ν -path μ coincide with $\check{\nu}$ -rotations of μ , and the right flushing bijection $\text{flush}_{\check{\nu}}$ transforms $\check{\nu}$ -rotations on paths to $\check{\nu}$ -rotations on the corresponding trees. Therefore, the second part of the proposition is clear. It remains to show that μ is a ν -path if and only if $\text{flush}_{\check{\nu}}(\mu)$ is a (δ, ν) -tree, or equivalently a $\check{\nu}$ -tree that is contained in $L_{\delta, \nu}$.

We start by proving the forward direction. First, note that the image of the bottom path $T_{\nu} = \text{flush}_{\check{\nu}}(\nu)$ is contained in $L_{\delta, \nu}$. More precisely, the shape $F_{\delta, \nu}$ has ν_k east steps on its boundary at height k . For $k > 0$, some of these east steps (exactly $\nu_k - \delta_k$) are on the left boundary, and some (exactly δ_k) are on the right boundary. The $\nu_k + 1$ nodes of T_{ν} at height k consist of the $\nu_k - \delta_k + 1$ points on the left boundary, and the δ_k end points of the east steps of the right boundary. At height $k = 0$, the $\nu_0 + 1$ nodes of T are all the lattice points at the bottom of $F_{\delta, \nu}$. This shows that T_{ν} is contained in $L_{\delta, \nu}$.

Now, every ν -path μ can be obtained by applying a sequence of $\check{\nu}$ -rotations to the bottom path ν . Its image $\text{flush}_{\check{\nu}}(\mu)$ is a $\check{\nu}$ -tree that can be obtained by applying the corresponding sequence of $\check{\nu}$ -rotations to the tree T_{ν} . Since T_{ν} is contained in $L_{\delta, \nu}$ and such rotations preserve this property, then $\text{flush}_{\check{\nu}}(\mu)$ is also contained in $L_{\delta, \nu}$. This finishes the proof of the forward direction.

The backward direction is equivalent to the following statement: if T is a $\check{\nu}$ -tree contained in $L_{\delta, \nu}$ then $\mu = \text{flush}_{\check{\nu}}^{-1}(T)$ is weakly above ν . This is equivalent to show that the number of nodes in T at heights less than or equal to k is at most $\nu_0 + \dots + \nu_k + k + 1$.

For $0 \leq k \leq n$, let L_k (resp. F_k) be the restriction of $L_{\delta, \nu}$ (resp. $F_{\delta, \nu}$) to the points with height less than or equal to k . The width of F_k is equal to $\nu_0 + \dots + \nu_k$. The maximal number of compatible lattice points inside L_k is equal to $\nu_0 + \dots + \nu_k + k + 1$. The restriction of T to L_k is a compatible set (not necessarily maximal). The result follows. \square

The (δ, ν) -right flushing bijection from ν -paths to (δ, ν) -trees is described in exactly the same way as in Section 3.2: we recursively add $\mu_i + 1$ nodes to the tree inside the shape $F_{\delta, \nu}$ from right to left, from bottom to top, while avoiding the forbidden positions above a node which is not the left most node in a row. Figure 6.1 shows the image of the path $\mu = (1, 0, 1, 1, 3, 2, 1, 2)$ for $\delta^{max} = (1, 0, 2, 2, 0, 3, 0)$ (left) and for $\delta = (0, 0, 1, 2, 0, 1, 0)$ (right), where the base path is $\nu = (3, 1, 0, 2, 2, 0, 3, 0)$.

Definition 4.10. The rotation poset of (δ, ν) -trees $\text{Tam}_{\nu}^{tr}(\delta)$ is the reflexive transitive closure of $\check{\nu}$ -rotations on (δ, ν) -trees.

The three examples of the rotation poset of (δ, ν) -trees for $\nu = ENEEN = (1, 2, 0)$ are illustrated on Figure 3.3.

Theorem 4.11. The alt ν -Tamari lattice is isomorphic to the rotation poset of (δ, ν) -trees:

$$\text{Tam}_{\nu}(\delta) \cong \text{Tam}_{\nu}^{tr}(\delta).$$

In particular, the rotation poset of (δ, ν) -trees is a lattice.

Proof. The alt ν -Tamari lattice is the poset on ν -paths whose covering relations are given by δ -rotations. The rotation poset of (δ, ν) -trees is poset on (δ, ν) -trees whose covering relations are $\check{\nu}$ -rotations. The result is then a consequence of Proposition 4.9. The lattice property was proven for the alt ν -Tamari lattice in Corollary 4.5. \square

5. Left and right intervals in the alt ν -Tamari lattice

Since $\text{Tam}_{\nu}(\delta)$ is an interval in Tam_{ν} , its linear intervals are linear intervals in Tam_{ν} . In terms of trees, this gives the following simple characterization.

Definition 5.1. An interval $[T, T']$ in $\text{Tam}_{\nu}^{tr}(\delta)$ is a *left interval* (resp. *right interval*) if $[T, T']$ is a left interval (resp. right interval) in $\text{Tam}_{\check{\nu}}^{tr}$.

Proposition 5.2. The non-trivial linear intervals in $\text{Tam}_{\nu}^{tr}(\delta)$ are either left or right intervals.

Proof. This is a direct consequence of Proposition 3.6. \square

In this section, we aim to characterize the left and right intervals in terms of certain row and (reduced) column vectors associated to the (δ, ν) -trees. This will be used in Section 6, to show that the number of linear intervals in the alt ν -Tamari lattice $\text{Tam}_{\nu}(\delta)$ is independent of the choice of δ .

5.1. Row vectors and left intervals

The row vector of a (δ, ν) -tree T is the vector

$$r(T) = (r_0, \dots, r_n),$$

where $r_i + 1$ is the number of nodes of T at height i .

Proposition 5.3. *A (δ, ν) -tree T is completely characterized by its row vector. Moreover, (r_0, \dots, r_n) is the row vector of some (δ, ν) -tree if and only if*

1. $r_i \geq 0$ for all i ,
2. $\sum_{i=0}^j r_i \leq \sum_{i=0}^j \nu_i$ for all j , and
3. $\sum_{i=0}^n r_i = \sum_{i=0}^n \nu_i$.

Proof. By Proposition 4.9, the map $\text{flush}_{\delta, \nu}$ is a bijection between the set of ν -paths and the set of (δ, ν) -trees. Moreover, this map preserves the number of points at each height, and therefore the row vector. Since ν -paths are characterized by their row vectors, then (δ, ν) -trees are characterized by their row vectors as well.

Furthermore, via the map $\text{flush}_{\delta, \nu}$, characterizing the row vectors of (δ, ν) -trees is equivalent to characterizing the row vectors of ν -paths. Condition (1) just says that every ν -path μ has at least one lattice point at each height. Condition (2) says that μ is weakly above ν , and Condition (3) says that μ and ν have the same ending points. \square

Given a (δ, ν) -tree T , we say that an ordered set $L = \{p, q_0, q_1, \dots, q_\ell\} \subseteq T$ is a horizontal \mathbb{L} of T if L is the restriction of T to a rectangle R of the grid, such that p is the top-left corner of R , and q_0, q_1, \dots, q_ℓ appear in this order on the bottom side of R with q_0 being its left-bottom corner and q_ℓ its right-bottom corner. Note that no other elements of T belong to R . We say that the length of L is equal to ℓ . We denote by $T + L$ the (δ, ν) -tree obtained from T by rotating the nodes $q_0, q_1, \dots, q_{\ell-1}$ in T in this order. An example of these concepts is illustrated in Figure 5.1.

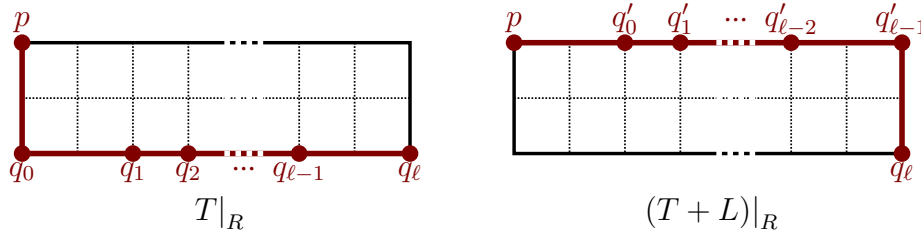


Figure 5.1: Schematic illustration of a horizontal \mathbb{L} and the tree $T + L$.

Lemma 5.4. *Let L be a horizontal \mathbb{L} of length ℓ of a (δ, ν) -tree T . Then, $[T, T + L]$ is a left interval of length ℓ in $\text{Tam}_\nu^{\text{tr}}(\delta)$. Moreover, every left interval of $\text{Tam}_\nu^{\text{tr}}(\delta)$ with bottom element T is of this form.*

Proof. This follows by the definition of left intervals. \square

Proposition 5.5. *Let T be a (δ, ν) -tree with row vector $r(T) = (r_0, \dots, r_n)$. The number of left intervals of length ℓ with bottom element T in $\text{Tam}_\nu^{tr}(\delta)$ is equal to*

$$|\{0 \leq i \leq n-1 : r_i \geq \ell\}|.$$

Proof. By Lemma 5.4, the left intervals of length ℓ with bottom element T are of the form $[T, T + L]$ where L is a horizontal \mathbb{L} of length ℓ of T . There is one such L for each $r_i \geq \ell$ with $0 \leq i \leq n-1$, where q_0, \dots, q_ℓ are the $\ell+1$ left most nodes of T at height i and p is the parent of q_0 in T . \square

The previous two results, Lemma 5.4 and Proposition 5.5, characterize the left intervals in $\text{Tam}_\nu^{tr}(\delta)$ with respect to the row vectors of (δ, ν) -trees. Our next goal is to have a similar characterization for the right intervals with respect to certain column vectors. As we will see, column vectors are not enough for such a characterization, and we will need to consider a notion of reduced column vectors. Before going into that, we first introduce column vectors and present some of their properties.

5.2. Column vectors

Given a path ν from $(0, 0)$ to (m, n) , the *reversed path* $\overleftarrow{\nu}$ is the path from $(0, 0)$ to (n, m) obtained by reading ν from right to left and replacing east steps by north steps and vice versa. Equivalently, $\overleftarrow{\nu} = (\overleftarrow{\nu}_0, \dots, \overleftarrow{\nu}_m)$ where $\overleftarrow{\nu}_i$ is the number of north steps of the path ν in column $m-i$. For instance, if $\nu = ENEEN = (1, 2, 0)$ then $\overleftarrow{\nu} = ENNEN = (1, 0, 1, 0)$. This notion is convenient to characterize column vectors.

In order to define the column vector of a (δ, ν) -tree, it is convenient to assign an order $j_0 \prec_\delta \dots \prec_\delta j_m$ to the columns of $L_{\delta, \nu}$, obtained by reading the columns from shortest to longest, from right to left, as illustrated in Figure 5.2 (left). See also the three examples in Figure 5.3.

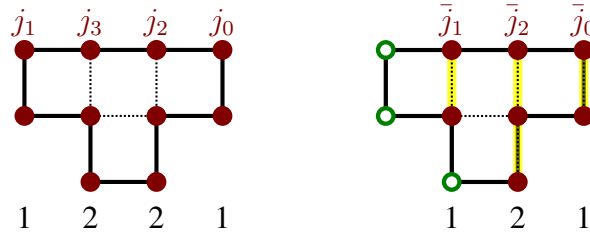


Figure 5.2: Left: the columns j_0, j_1, j_2, j_3 of $L_{\delta, \nu}$, their lengths are 1,1,2,2. Right: the reduced columns $\bar{j}_0, \bar{j}_1, \bar{j}_2$ of $L_{\delta, \nu}$ with lengths 1,1,2, where the relevant points are filled brown and the non-relevant points are unfilled green. In both cases, the columns are read from shortest to longest, from right to left.

The *column vector* of a (δ, ν) -tree T is the vector

$$c_\delta(T) = (c_0, \dots, c_m),$$

where $c_i + 1$ is the number of nodes of T in column j_i . For instance, the three (δ, ν) -trees (for the three choices of δ) in Figure 5.3, all have column vector $(0, 1, 0, 1)$. This means, in each of the cases, there are $0 + 1$ nodes of the tree in column j_0 , $1 + 1$ nodes in column j_1 , $0 + 1$ nodes in column j_2 , and $1 + 1$ nodes in column j_3 . Equivalently, c_i counts the number of edges of T in column j_i .

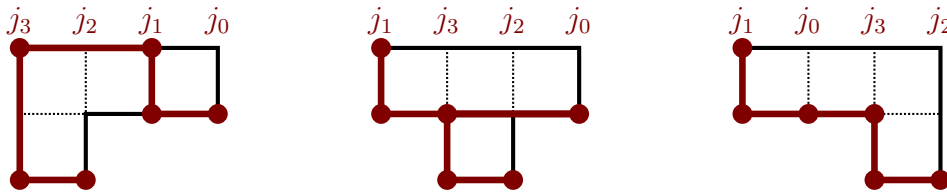


Figure 5.3: The columns j_0, j_1, j_2, j_3 of $L_{\delta, \nu}$ for $\nu = ENEEN$ and the three possible choices of $\delta = (2, 0), (1, 0)$ and $(0, 0)$. The columns are read from shortest to longest, from right to left. The column vector of the shown trees is $c_\delta(T) = (0, 1, 0, 1)$ in all three cases.

Proposition 5.6. *A (δ, ν) -tree T is completely characterized by its column vector. Moreover, (c_0, \dots, c_m) is the column vector of some (δ, ν) -tree if and only if*

1. $c_i \geq 0$ for all i ,
2. $\sum_{i=0}^j c_i \leq \sum_{i=0}^j \leftarrow v_i$ for all j , and
3. $\sum_{i=0}^m c_i = \sum_{i=0}^m \leftarrow v_i$.

We prove this proposition in several steps.

Lemma 5.7. *A (δ, ν) -tree T can be reconstructed from its column vector.*

Proof. Let T be a (δ, ν) -tree. Label the elements of the tree p_0, p_1, \dots, p_r from right to left, from bottom to top, as illustrated in Figure 5.4.

We reconstruct T recursively, by adding the elements p_0, p_1, \dots, p_r one at a time in order. Note that if p_i is not the top most element in its column, then all the lattice points on the left of p_i are forbidden in the next steps, because they are incompatible with an element $p_j \in T$ that is above p_i in the same column.

Now, when we add an element p_j in the process of reconstructing T , then p_j is necessarily located at the bottom most position of its column that is not forbidden by any element before. Otherwise, let $p_j \in T$ be the node with smallest label that does not satisfy that property, and let q be the bottom most lattice point in the same column that is not forbidden by any element p_i with $i < j$. In particular, q is below p_j by assumption, and q is compatible with every p_i with $i < j$. Moreover, q is also compatible with $p_k \in T$ with $k > j$, otherwise p_k, p_j would be incompatible. So, we can add the element q to T , creating a new compatible set, contradicting the maximality of T .

The tree T can therefore be constructed by adding nodes from right to left, from bottom to top, avoiding forbidden positions. The forbidden positions are those to the left of a node that is not the top most node in a column. The number of points in each column is determined by the column vector. \square

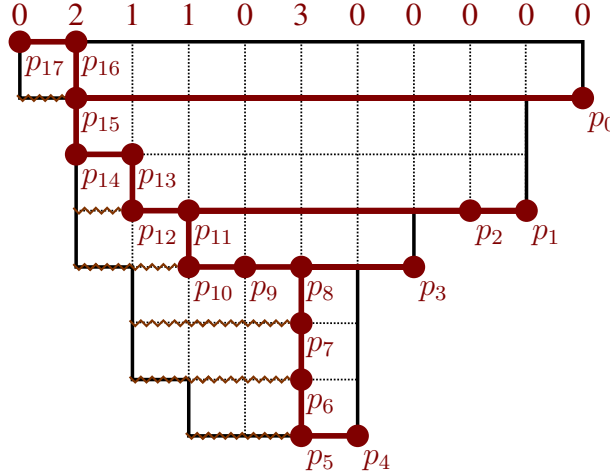


Figure 5.4: The down flushing algorithm.

We call the algorithm described in the previous proof the *down flushing algorithm*. Its input is a valid column vector (c_0, \dots, c_m) (or the number of nodes in each column), and its output is the unique (δ, ν) -tree such that $c_\delta(T) = (c_0, \dots, c_m)$. Figure 5.4 illustrate an example, where the labels on top represent the number of nodes, minus 1, in each column, and the forbidden positions are the ones that belong to the zigzag lines.

Lemma 5.8. *Let δ, δ' be two increment vectors with respect to ν , such that δ' is obtained by either adding or subtracting 1 to one of the entries of δ . For every (δ, ν) -tree T , there is a unique (δ', ν) -tree T' such that*

$$c_\delta(T) = c_{\delta'}(T').$$

Proof. Uniqueness follows by Lemma 5.7, so we just need to prove existence.

Let T be a (δ, ν) -tree with column vector $c_\delta(T) = (c_0, \dots, c_m)$, and assume that δ' is obtained by subtracting 1 to a non-zero entry δ_a of δ . This operation produces a small transformation to the columns of $L_{\delta, \nu}$. All the columns of length larger than $n - a$ are moved one step to the right, while the subsequent column (of length $n - a$) is moved one step to their left. All other columns stay the same. The result is the new set $L_{\delta', \nu}$. An example is illustrated in Figure 5.6.

Consider the labeling j_0, \dots, j_m of the columns of $L_{\delta, \nu}$ (and also of the columns of $L_{\delta', \nu}$) obtained by reading the columns from shortest to longest, from right to left, as before. Assume that j_{i_1} is the label of the column that was moved to the left under the small transformation that changes δ to δ' . We also consider the columns j_{i_2}, \dots, j_{i_k} , consisting of the columns of $L_{\delta, \nu}$, from right to left, of length bigger than $n - a$ that contain at least one node of T at height bigger than or equal to a . The restriction of the tree T to the nodes at height bigger than or equal to a in columns j_{i_1}, \dots, j_{i_k} is marked as a bold red path on the left of Figures 5.5 and 5.6. It is a subpath

of the unique path of the tree from column j_{i_1} to the root of the tree. We will describe a small transformation to T that produces a (δ', ν) -tree T' with the same column vector as T . The result of this is illustrated on the right of Figures 5.5 and 5.6, and affects the tree at the red marked nodes. The brown points in the columns between \bar{j}_{i_k} and \bar{j}_{i_1} are also moved one step to the right, together with their column.

Note that the columns j_{i_2}, \dots, j_{i_k} of $L_{\delta', \nu}$ are positioned one step to the right of columns j_{i_2}, \dots, j_{i_k} of $L_{\delta, \nu}$, while column j_1 was moved to some position to the left, see Figures 5.5 and 5.6.

Let A be the set of rows that contain at least one node of the marked bold red path of T . We apply the following transformation to T . For each node T in a column j_{i_b} , for $2 \leq b \leq k$, that belongs to A , we draw a node in T' in column j_{i_b} but shifted down $c_{j_{i_1}}$ positions within A . The $c_{j_{i_1}} + 1$ nodes in column j_{i_1} are moved to the top rows of A . All other nodes of T remain intact in their columns. A schematic illustration of this transformation is shown in Figure 5.5, and an explicit example in Figure 5.6.

The result is a (δ', ν) -tree T' with the same column vector as T : $c_\delta(T) = c_{\delta'}(T')$. The reason why this procedure works is guaranteed by a direct analysis of the down flushing algorithm. Moreover, we can also recover T from T' by a similar transformation in the reverse direction. \square




Figure 5.5: Schematic illustration of the transformation in the proof of Lemma 5.8.

Lemma 5.9. *Let δ, δ' be two increment vectors with respect to ν . For every (δ, ν) -tree T , there is a unique (δ', ν) -tree T' such that*

$$c_\delta(T) = c_{\delta'}(T').$$

Proof. Any two increment vectors with respect to ν can be connected by a sequence of increment vectors, such that each vector is obtained from the previous one by either adding or subtracting 1 to one of its entries. The result then follows by Lemma 5.8. \square

Proof of Proposition 5.6. By Lemma 5.7, a (δ, ν) -tree T is completely characterized by its column vector. Furthermore, the characterization of column vectors of (δ, ν) -trees is independent of the choice of increment vector δ , by Lemma 5.9. So, we just need to prove the three conditions of the proposition for one particular choice of δ . We choose the extreme case $\delta = \delta^{\max}$, where $\delta_i = \nu_i$. In this case (δ, ν) -trees are just the classical ν -trees.

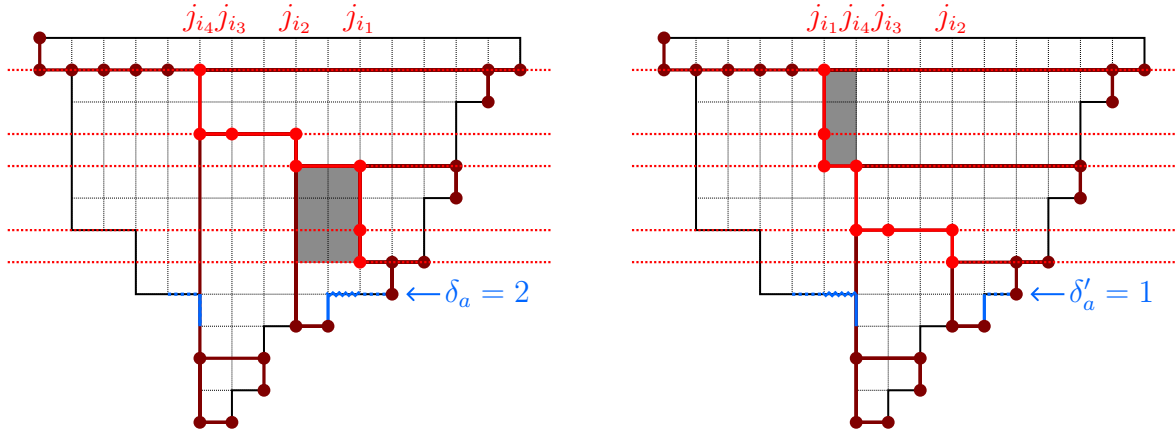


Figure 5.6: Example of the transformation in the proof of Lemma 5.8.

Classifying the column vectors of ν -trees is the same as classifying the row vectors of $\overleftarrow{\nu}$ -trees, because reversing the path transforms column vectors to row vectors, and vice versa. The three conditions of the proposition are then equivalent to the three conditions of Proposition 5.3 (for the extreme maximal case δ). \square

5.3. Reduced column vectors

We say that a lattice point $p \in L_{\delta, \nu}$ is *non-relevant* if it is the leftmost point of a row of $L_{\delta, \nu}$. All other points in $L_{\delta, \nu}$ are called *relevant*. Figure 5.2 (right) illustrates an example where the relevant points are filled brown, and the non-relevant points are unfilled green.

The *reduced columns* are the columns of relevant points in $L_{\delta, \nu}$. These are shown in yellow in Figure 5.2 (right). The three examples for $\nu = ENEEN$ and all possible choices of δ are shown in Figure 5.7. The reduced columns are colored yellow here as well for easier visualization.

In order to define the reduced column vector of a (δ, ν) -tree, it is convenient to assign an order $\bar{j}_0 \prec_{\delta} \cdots \prec_{\delta} \bar{j}_{m-1}$ to the reduced columns of $F_{\delta, \nu}$, obtained by reading the reduced columns from shortest to longest, from right to left, as illustrated in Figure 5.2 (right). See also the three examples in Figure 5.7.

The *reduced column vector* of a (δ, ν) -tree T is the vector

$$\bar{c}_{\delta}(T) = (\bar{c}_0, \dots, \bar{c}_{m-1}),$$

where $\bar{c}_i + 1$ is the number of nodes of T in reduced column \bar{j}_i . For instance, the three (δ, ν) -trees (for the three choices of δ) in Figure 5.7, all have reduced column vector $(0, 1, 0)$. This means, in each of the cases, there are $0 + 1$ nodes of the tree in reduced column \bar{j}_0 , $1 + 1$ nodes in reduced column \bar{j}_1 , and $0 + 1$ nodes in reduced column \bar{j}_2 . Note that the green nodes are non-relevant and do not belong to the reduced columns, and so are not counted here.

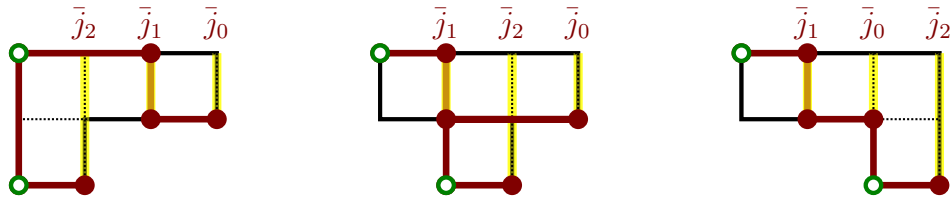


Figure 5.7: The ordering $\bar{j}_0 \prec_\delta \cdots \prec_\delta \bar{j}_2$ of the reduced columns of $L_{\delta, \nu}$ for $\nu = ENEEN$ and the three possible choices of $\delta = (2, 0), (1, 0)$ and $(0, 0)$. The reduced columns (colored yellow) are read from shortest to longest, from right to left. The reduced column vector of the shown trees is $\bar{c}_\delta(T) = (0, 1, 0)$ in all three cases.

Proposition 5.10. *A (δ, ν) -tree T is completely characterized by its reduced column vector. Moreover, $(\bar{c}_0, \dots, \bar{c}_{m-1})$ is the reduced column vector of some (δ, ν) -tree if and only if*

1. $\bar{c}_i \geq 0$ for all i ,
2. $\sum_{i=0}^j \bar{c}_i \leq \sum_{i=0}^j \bar{\nu}_i$ for all j .

The proof of this proposition follows the same steps as the proof of Proposition 5.6 for column vectors. We write all the (somewhat repeated) details for self containment.

Lemma 5.11. *A (δ, ν) -tree T can be reconstructed from its reduced column vector.*

Proof. We proceed in a similar way as in the proof of Lemma 5.7, with the small difference that we need to be careful what to do with the non-relevant positions, which are not counted by the reduced column vector.

Let $\bar{c}_\delta(T) = (\bar{c}_0, \dots, \bar{c}_{m-1})$ be the reduced column vector of T . Similarly as before, the tree T can be reconstructed by adding nodes from right to left, from bottom to top, avoiding the forbidden positions that are to the left of a node that is not the top most node of its column. Here comes the tricky part. When we want to add the nodes in column \bar{j}_i , there are two possible scenarios:

(1) If there are non-relevant positions in column \bar{j}_i that are not forbidden by any of the nodes added before in the process, then these non-relevant positions are automatically compatible with all the nodes of the tree T (the ones that were already added, and all the future ones). Therefore, all the non-relevant nodes in column \bar{j}_i that are not forbidden by any previously added node should be added to T . After this we proceed adding $\bar{c}_i + 1$ nodes from bottom to top in the positions that are not forbidden in column \bar{j}_i .

(2) If all the non-relevant positions in column \bar{j}_i are forbidden, then we simply proceed adding $\bar{c}_i + 1$ nodes from bottom to top in the positions that are not forbidden in that column.

This procedure reconstructs the tree T and only depends on the reduced column vector.

An example is illustrated in Figure 5.8. Note that the unfilled green point p_{14} is non-relevant, and was forced to be added to T because it is not forbidden by any of the previously added nodes p_1, \dots, p_{13} . At this step of the process, one proceeds adding the $1+1$ relevant points p_{15}, p_{16} in that column, which are counted by the corresponding entry plus one of the reduced column vector. □

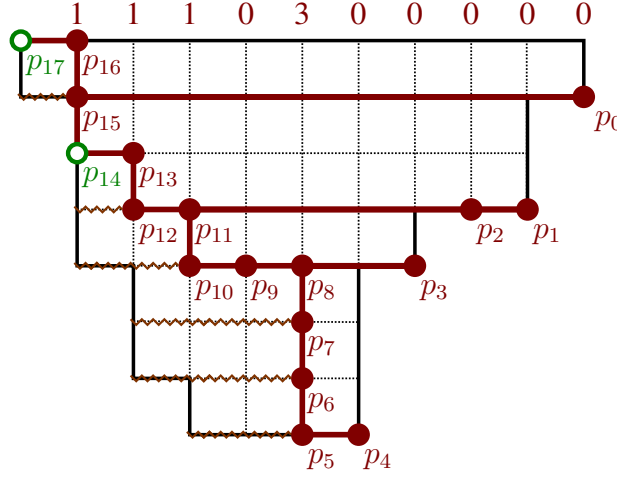


Figure 5.8: The reduced down flushing algorithm.

We call the algorithm described in the previous proof the *reduced down flushing algorithm*. Its input is a valid reduced column vector $(\bar{c}_0, \dots, \bar{c}_{m-1})$ (or the number of relevant nodes in each column), and its output is the unique (δ, ν) -tree such that $\bar{c}_\delta(T) = (\bar{c}_0, \dots, \bar{c}_{m-1})$.

Lemma 5.12. *Let δ, δ' be two increment vectors with respect to ν , such that δ' is obtained by either adding or subtracting 1 to one of the entries of δ . For every (δ, ν) -tree T , there is a unique (δ', ν) -tree T' such that*

$$\bar{c}_\delta(T) = \bar{c}_{\delta'}(T').$$

Moreover, the heights of the non-relevant nodes of T and T' coincide.

Proof. Uniqueness follows by Lemma 5.11, so we just need to prove existence.

Let T be a (δ, ν) -tree with reduced column vector $\bar{c}_\delta(T) = (\bar{c}_0, \dots, \bar{c}_{m-1})$, and assume that δ' is obtained by subtracting 1 to a non-zero entry δ_a of δ . This operation produces a small transformation to the reduced columns of $L_{\delta, \nu}$ (which is slightly different to the transformation in the proof on Lemma 5.8). All the reduced columns of length larger than $n - a$ are moved one step to the right, while the subsequent reduced column (of length $n - a$) is moved one step to their left. All other columns stay the same. An example is illustrated in Figure 5.10. Here, we have chosen the same example as in Figure 5.6, to highlight the differences with the transformation described in the proof of Lemma 5.8.

Consider the labeling $\bar{j}_0, \dots, \bar{j}_{m-1}$ of the reduced columns of $L_{\delta, \nu}$ (and also of the reduced columns of $L_{\delta', \nu}$) obtained by reading the reduced columns from shortest to longest, from right to left, as before. Assume that \bar{j}_{i_1} is the label of the reduced column that was moved to the left under the small transformation that changes δ to δ' . We also consider the columns $\bar{j}_{i_2}, \dots, \bar{j}_{i_k}$, consisting of the reduced columns of $L_{\delta, \nu}$, from right to left, of length bigger than $n - a$ that contain at least one node of T at height bigger than or equal to a . The restriction of the tree T to the nodes at height bigger than or equal to a in the reduced columns $\bar{j}_{i_1}, \dots, \bar{j}_{i_k}$ is marked as a bold red path on the left of Figures 5.9 and 5.10. It is a subpath of the unique path of the tree from column \bar{j}_{i_1} to the root of the tree.

Note that column \bar{j}_{i_4} , of length bigger than $n - a$ in Figure 5.6, is now a reduced column of length $n - a$. That is why there is no \bar{j}_{i_4} in our example in Figure 5.10.

We will describe a small transformation to T that produces a (δ', ν) -tree T' with the same reduced column vector as T . The result of this is illustrated on the right of Figures 5.9 and 5.10, and affects the red marked nodes of the tree. The brown and green points between the columns \bar{j}_{i_k} and \bar{j}_{i_1} are also moved one step to the right.

Note that the columns $\bar{j}_{i_2}, \dots, \bar{j}_{i_k}$ of $L_{\delta', \nu}$ are positioned one step to the right of columns $\bar{j}_{i_2}, \dots, \bar{j}_{i_k}$ of $L_{\delta, \nu}$, while column \bar{j}_{i_1} was moved to some position to the left, see Figures 5.9 and 5.10.

Let A be the set of rows that contain at least one node of the marked bold red path of T . We apply the following transformation to T . For each node T in a reduced column \bar{j}_{i_b} , for $2 \leq b \leq k$, that belongs to A , we draw a node in T' in the reduced column \bar{j}_{i_b} but shifted down $\bar{c}_{j_{i_1}}$ positions within A . The $\bar{c}_{j_{i_1}} + 1$ nodes in reduced column \bar{j}_{i_1} are moved to the top rows of A . All other relevant nodes of T remain intact in their reduced columns, and all non-relevant nodes remain intact in their “not reduced” columns. A schematic illustration of this transformation is shown in Figure 5.9, and an explicit example in Figure 5.10.

The result is a (δ', ν) -tree T' with the same reduced column vector as T : $\bar{c}_\delta(T) = \bar{c}_{\delta'}(T')$, and such that the heights of the non-relevant nodes are preserved. The reason why this procedure works is guaranteed by a direct analysis of the reduced down flushing algorithm. Moreover, we can also recover T from T' by a similar transformation in the reverse direction. \square

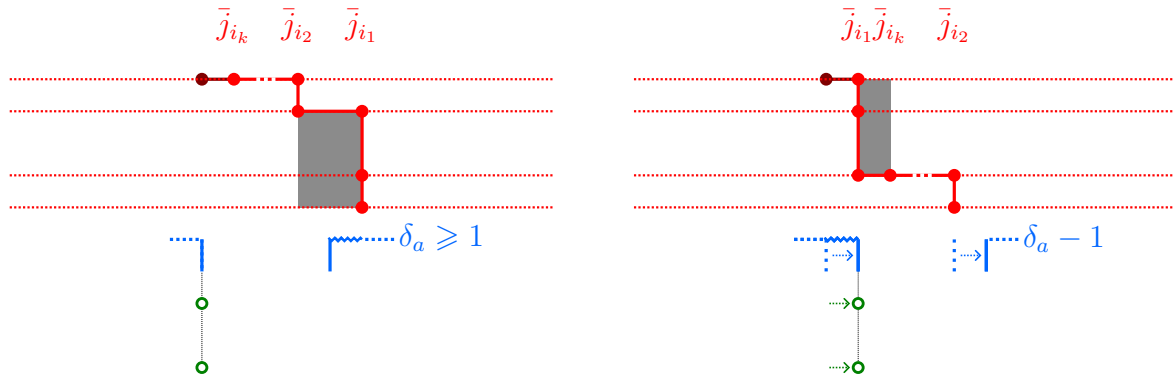


Figure 5.9: Schematic illustration of the transformation in the proof of Lemma 5.12.

Lemma 5.13. *Let δ, δ' be two increment vectors with respect to ν . For every (δ, ν) -tree T , there is a unique (δ', ν) -tree T' such that*

$$\bar{c}_\delta(T) = \bar{c}_{\delta'}(T').$$

Moreover, the heights of the non-relevant nodes of T and T' coincide.

Proof. Any two increment vectors with respect to ν can be connected by a sequence of increment vectors, such that each vector is obtained from the previous one by either adding or subtracting 1 to one of its entries. The result then follows by Lemma 5.12. \square

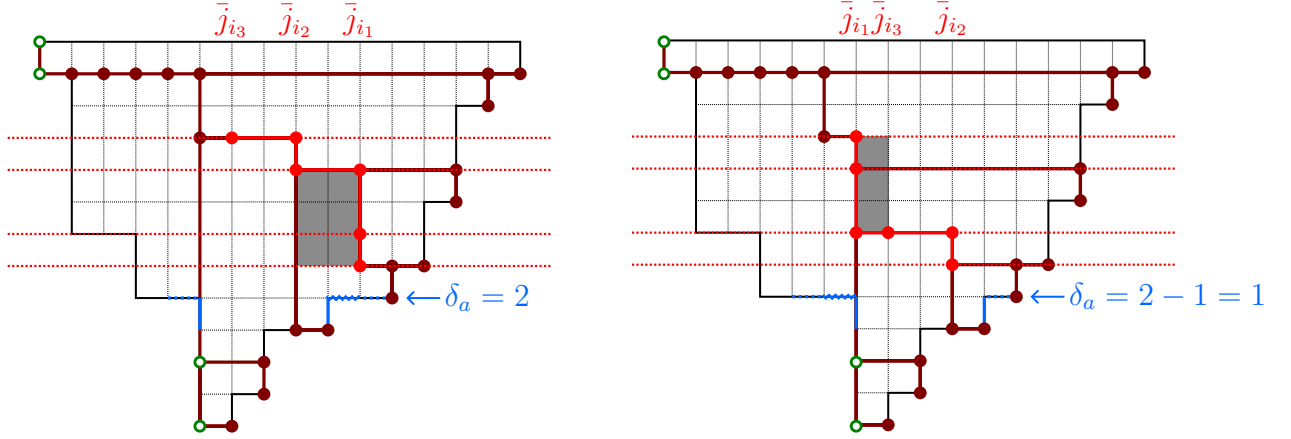


Figure 5.10: Example of the transformation in the proof of Lemma 5.12.

Proof of Proposition 5.10. A (δ, ν) -tree T is completely characterized by its reduced column vector by Lemma 5.11. Furthermore, the characterization of reduced column vectors of (δ, ν) -trees is independent of the choice of increment vector δ , by Lemma 5.13. So, we just need to prove the two conditions of the proposition for one particular choice of δ . We choose the extreme case $\delta = \delta^{\max}$, where $\delta_i = \nu_i$. In this case (δ, ν) -trees are just the classical ν -trees.

The reduced column vector $(\bar{c}_0, \dots, \bar{c}_{m-1})$ of a ν -tree T is obtained from the row vector (r_0, \dots, r_m) of the corresponding $\overleftarrow{\nu}$ -tree \overleftarrow{T} by removing its last entry r_m . The two conditions of the proposition are then equivalent to the first two conditions of Proposition 5.3 (for the extreme maximal case δ). The third condition of Proposition 5.3 is an equation involving the removed term r_m , and is irrelevant for Proposition 5.10. That is why we now have only two conditions. \square

5.4. Reduced column vectors and right intervals

We are finally ready to provide our characterization of right intervals in $\text{Tam}_\nu^{tr}(\delta)$ in terms of reduced column vectors.

Given a (δ, ν) -tree T , we say that an *ordered* set $L = \{p, q'_0, q'_1, \dots, q'_\ell\} \subseteq T$ is a *vertical \mathbb{L}* of T if L is the restriction of T to a rectangle $R \subseteq F_{\delta, \nu}$ of the grid, such that p is the top-left corner of R , and $q'_0, q'_1, \dots, q'_\ell$ appear in this order from top to bottom on the right side of R , with q'_0 being its top-right corner and q'_ℓ its bottom-right corner. Note that no other elements of T belong to R . We say that the length of L is equal to ℓ . We denote by $T - L$ the (δ, ν) -tree obtained from T by rotating down the nodes $q'_0, q'_1, \dots, q'_{\ell-1}$ in T in this order.

Note that the condition $R \subseteq F_{\delta, \nu}$ is crucial here, to guaranty that the result after applying these rotations is still contained in the Ferrers diagram $F_{\delta, \nu}$, otherwise $T - L$ would not be a (δ, ν) -tree. In particular, if $R \subseteq F_{\delta, \nu}$ then $q'_0, q'_1, \dots, q'_\ell$ are all relevant nodes in T , and contribute to the reduced column vector. Vice versa, if q'_ℓ is relevant then $p \perp q'_\ell \in F_{\delta, \nu}$ because of the reduced down flushing algorithm, and thus $R \subseteq F_{\delta, \nu}$.

An example of these concepts is illustrated in Figure 5.11.

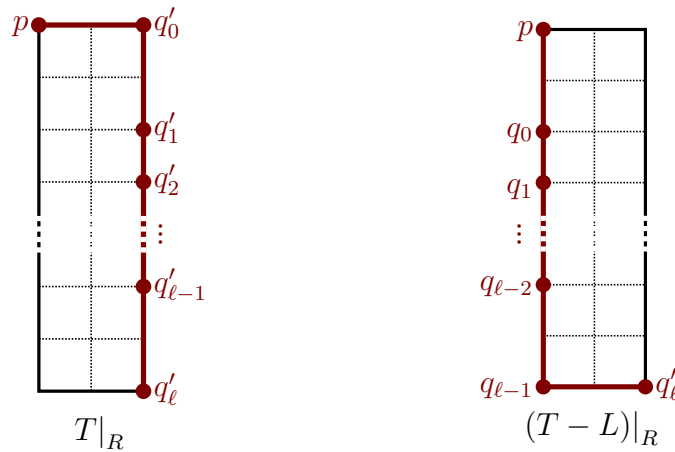


Figure 5.11: Schematic illustration of a vertical \mathbb{L} and the tree $T - L$.

Lemma 5.14. *Let L be a vertical \mathbb{L} of length ℓ of a (δ, ν) -tree T . Then, $[T - L, T]$ is a right interval of length ℓ in $\text{Tam}_\nu^{tr}(\delta)$. Moreover, every right interval of $\text{Tam}_\nu^{tr}(\delta)$ with top element T is of this form.*

Proof. This follows by the definition of right intervals. □

Proposition 5.15. *Let T be a (δ, ν) -tree with reduced column vector $\bar{c}_\delta(T) = (\bar{c}_0, \dots, \bar{c}_{m-1})$. The number of right intervals of length ℓ with top element T in $\text{Tam}_\nu^{tr}(\delta)$ is equal to*

$$|\{0 \leq i \leq m - 1 : \bar{c}_i \geq \ell\}|.$$

Proof. By Lemma 5.14, the right intervals of length ℓ with top element T are of the form $[T - L, T]$ where L is a vertical \mathbb{L} of length ℓ of T . There is one such L for each $\bar{c}_i \geq \ell$ with $0 \leq i \leq m - 1$, where q'_0, \dots, q'_ℓ are the $\ell + 1$ top most nodes of T at column \bar{j}_i and p is the parent of q_0 in T . □

6. Bijections between linear intervals

Using the tools developed in the previous section, we are now ready to prove one of our main results.

Theorem 6.1. *For a fixed path ν , all alt ν -Tamari lattices $\text{Tam}_\nu(\delta)$ have the same number of linear intervals of length ℓ .*

This is a direct consequence of Proposition 5.2 and Corollaries 6.5 and 6.7, which show that the number of left intervals and the number of right intervals of length ℓ are preserved for any choice of δ . Indeed, we prove more refined versions of these results in Propositions 6.4 and 6.6.

Remark 6.2. If we chose any other path $\tilde{\nu}$ weakly below ν that does not satisfy $\tilde{\nu}_i \leq \nu_i$, for all $i > 0$, then the restriction of $\text{Tam}_{\tilde{\nu}}$ to the subset of ν -paths does not satisfy the enumerative result of Theorem 6.1.

More precisely, this poset still has the same number of left intervals (the left flushing argument presented afterwards still works) as all alt ν -Tamari lattices. But based on computational experiments, it seems to have fewer right intervals. For instance, for $\nu = (1, 2, 0)$ and $\check{\nu} = (1, 2, 0)$, the distribution of linear intervals in the resulting poset is $(5, 5, 1)$ but the distribution of linear intervals in Tam_ν is $(5, 5, 2)$.

Remark 6.3. Theorem 6.1 generalizes the results obtained in [Che22] for the staircase $\nu = (NE)^n$. However, in this more general case, we usually do not have a closed formula counting the linear intervals of length ℓ similar to the one presented in [Che22].

In the m -Tamari lattice, where $\nu = (NE^m)^n$, one can adapt the decomposition given in [Che22] in order to find a closed formula for the number of right intervals of length ℓ :

$$m \binom{mn + n - \ell}{n - \ell - 1}.$$

We were not able to find a simple formula for the number of left intervals in this case. For $n = 5$ and $m = 2$, the distribution of left intervals in this lattice is $(728, 442, 222, 112, 47, 18, 5, 1)$. Since 47 is a prime number, no such a nice product formula seems to exist.

However, one can easily express the number of left intervals for general ν as a sum (and similarly for right intervals). Indeed, this number is independent of the choice of δ , so we can consider the case $\delta = (0, \dots, 0)$ giving rise to the ν -Dyck lattice. In this case, the number of left intervals is

$$\sum_{\substack{p, q \in L_\nu \\ q = p + (k, 1)}} \mathcal{C}(\nu, p) \tilde{\mathcal{C}}(\nu, q),$$

where $\mathcal{C}(\nu, p)$ (resp. $\tilde{\mathcal{C}}(\nu, q)$) is defined as the number of lattice paths weakly above ν , from the initial point of ν to p (resp. from q to the final point of ν). The numbers $\mathcal{C}(\nu, p)$ and $\tilde{\mathcal{C}}(\nu, q)$ can be calculated using Kreweras determinantal formula in [Kre65], which states that the number of partitions fitting a partition $\lambda = (\lambda_1, \dots, \lambda_\ell)$ is

$$\det \begin{pmatrix} \lambda_j + 1 \\ j - i + 1 \end{pmatrix}.$$

It would be interesting to compute the number of left or right intervals over all ν -paths of a given length n . Computational experiments seem to indicate that these numbers might be nice. For related work (counting all intervals) in connection with non-separable planar maps, see [FPR17].

6.1. The horizontal flushing and left intervals

We define the *horizontal flushing* $f_{\delta, \delta'}^h$ as the map between the set of (δ, ν) -trees and the set of (δ', ν) -trees characterized by the property

$$f_{\delta, \delta'}^h(T) = T' \quad \longleftrightarrow \quad r(T) = r(T').$$

That is, the map that preserves the row vector of the tree. This map is uniquely determined by this property, and can be computed as the composition

$$f_{\delta, \delta'}^h(T) = \text{flush}_{\delta', \nu} \circ \text{flush}_{\delta, \nu}^{-1},$$

which sends a (δ, ν) -tree to the unique ν -path with the same row vector, and then to the corresponding (δ', ν) -tree. In particular, $f_{\delta, \delta'}^h$ is a bijection, and can be described using a *horizontal flushing algorithm*:

If $r(T) = (r_0, \dots, r_n)$, then T' can be reconstructed by adding $r_i + 1$ nodes, from bottom to top, from right to left, avoiding the forbidden positions that are above the nodes that are not the left most nodes in their row.

This gives a natural correspondence between the horizontal \mathbb{L} 's of T and the horizontal \mathbb{L} 's of T' : an \mathbb{L} of length ℓ in row i ¹ of T corresponds to the unique horizontal \mathbb{L} of the same length in row i of T' . By abuse of notation, we denote by $f_{\delta, \delta'}^h(L) = L'$ the horizontal \mathbb{L} of T' associated to L , a horizontal \mathbb{L} of T .

Proposition 6.4. *Let T be a (δ, ν) -tree and $T' = f_{\delta, \delta'}^h(T)$ be its corresponding (δ', ν) -tree. We also denote by $L' = f_{\delta, \delta'}^h(L)$ the horizontal \mathbb{L} of T' associated to L , a horizontal \mathbb{L} of T .*

1. *The number of left intervals of length ℓ in $\text{Tam}_{\nu}^{tr}(\delta)$ with bottom element T is equal to the number of left intervals of length ℓ in $\text{Tam}_{\nu}^{tr}(\delta')$ with bottom element T' .*
2. *The map*

$$[T, T + L] \rightarrow [T', T' + L']$$

is a bijection between the left intervals of $\text{Tam}_{\nu}^{tr}(\delta)$ and the left intervals of $\text{Tam}_{\nu}^{tr}(\delta')$.

Proof. By Proposition 5.5, the number of left intervals with bottom element T depends only of the row vector $r(T)$. Since the $r(T) = r(T')$, then Item (1) follows. Item (2) is straight forward from the characterization of left intervals in Lemma 5.4. □

An example of the bijection between left intervals is illustrated in Figure 6.1. The maximal horizontal \mathbb{L} 's are marked red for easier visualization.

Corollary 6.5. *The number of left intervals of length ℓ in $\text{Tam}_{\nu}^{tr}(\delta)$ and $\text{Tam}_{\nu}^{tr}(\delta')$ are the same.*

Proof. This is a direct consequence of Proposition 6.4. □

6.2. The reduced vertical flushing and right intervals

We define the *vertical flushing* $f_{\delta, \delta'}^v$ as the map between the set of (δ, ν) -trees and the set of (δ', ν) -trees characterized by the property

$$f_{\delta, \delta'}^v(T) = T' \iff \bar{c}_{\delta}(T) = \bar{c}_{\delta'}(T').$$

That is, the map that preserves the reduced column vector of the tree.

¹here we mean that the bottom part of the \mathbb{L} is in row i

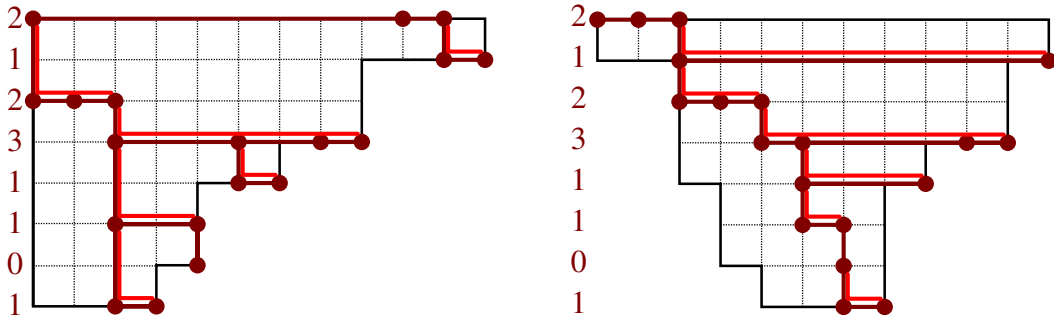


Figure 6.1: Bijection between left intervals for $\delta^{max} = (1, 0, 2, 2, 0, 3, 0)$ and $\delta = (0, 0, 1, 2, 0, 1, 0)$. Both trees have row vector $(1, 0, 1, 1, 3, 2, 1, 2)$, whose entries plus one count the number of nodes in each of the rows.

This map is uniquely determined by this property by Lemma 5.13. In particular, $f_{\delta, \delta'}^v$ is a bijection, and can be described using a *vertical flushing algorithm*:

If $\bar{c}_\delta(T) = (\bar{c}_0, \dots, \bar{c}_{m-1})$, then T' can be reconstructed by adding nodes, from right to left, from bottom to top, avoiding the forbidden positions that are to the left of the nodes that are not the top most nodes in their column. The difference here is that the number of nodes that we add to a column, whose reduced column is labeled \bar{j}_i , is not necessarily equal to $\bar{c}_{j_i} + 1$: we first add all the non-relevant nodes that are not forbidden by any of the previously added nodes; then we continue adding $\bar{c}_{j_i} + 1$ relevant nodes from bottom to top in the non-forbidden available positions.

This also gives a natural correspondence between the vertical \mathbb{L} 's of T and the vertical \mathbb{L} 's of T' : an \mathbb{L} of length ℓ in reduced column \bar{j}_i ² of T corresponds to the unique vertical \mathbb{L} of the same length in reduced column \bar{j}_i of T' . By abuse of notation, we denote by $f_{\delta, \delta'}^v(L) = L'$ the vertical \mathbb{L} of T' associated to L , a vertical \mathbb{L} of T .

Proposition 6.6. *Let T be a (δ, ν) -tree and $T' = f_{\delta, \delta'}^v(T)$ be its corresponding (δ', ν) -tree. We also denote by $L' = f_{\delta, \delta'}^v(L)$ the vertical \mathbb{L} of T' associated to L , a vertical \mathbb{L} of T .*

1. *The number of right intervals of length ℓ in $\text{Tam}_\nu^{tr}(\delta)$ with top element T is equal to the number of right intervals of length ℓ in $\text{Tam}_\nu^{tr}(\delta')$ with top element T' .*

2. *The map*

$$[T, T - L] \rightarrow [T', T' - L']$$

is a bijection between the right intervals of $\text{Tam}_\nu^{tr}(\delta)$ and the right intervals of $\text{Tam}_\nu^{tr}(\delta')$.

Proof. By Proposition 5.15, the number of right intervals with top element T depends only of the reduced column vector $\bar{c}_\delta(T)$. Since the $\bar{c}_\delta(T) = \bar{c}_{\delta'}(T')$, then Item (1) follows. Item (2) is straight forward from the characterization of right intervals in Lemma 5.14. \square

²here we mean that the right part of the \mathbb{L} is in reduced column \bar{j}_i

Examples of the bijection between right intervals are illustrated in Figures 6.2 and 6.3. The maximal vertical \mathbb{L} 's are marked red for easier visualization. The green nodes are the non-relevant nodes.

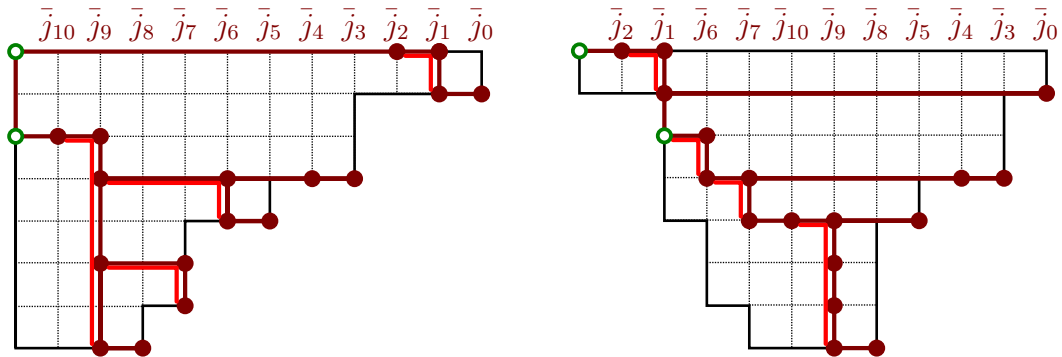


Figure 6.2: Bijection between right intervals for δ^{max} and $\delta = (0, 0, 1, 2, 0, 1, 0)$. Both trees have reduced column vector $(0, 1, 0, 0, 0, 0, 1, 1, 0, 3, 0)$, whose entries plus one count the number of relevant nodes in the reduced columns. The green non-relevant nodes are not counted.

Corollary 6.7. *The number of right intervals of length ℓ in $\text{Tam}_\nu^{tr}(\delta)$ and $\text{Tam}_\nu^{tr}(\delta')$ are the same.*

Proof. This is a direct consequence of Proposition 6.6. □

Acknowledgement

This project started during discussions at the SLC 88 conference in Strobl in 2022, and we are grateful for the pleasant and motivating atmosphere during the conference. We are also grateful to Wenjie Fang, Matias von Bell, and especially Frédéric Chapoton for helpful and interesting discussions.

References

[BMCPR13] Mireille Bousquet-Mélou, Guillaume Chapuy, and Louis-François Prévaille-Ratelle. The representation of the symmetric group on m -Tamari intervals. *Adv. Math.*, 247:309–342, 2013. doi:10.1016/j.aim.2013.07.014.

[BMFPR11] Mireille Bousquet-Mélou, Éric Fusy, and Louis-François Prévaille-Ratelle. The number of intervals in the m -Tamari lattices. *Electron. J. Combin.*, 18(2):Paper 31, 26, 2011. doi:10.37236/2027.

[BPR12] François Bergeron and Louis-François Prévaille-Ratelle. Higher trivariate diagonal harmonics via generalized Tamari posets. *J. Comb.*, 3(3):317–341, 2012. doi:10.4310/JOC.2012.v3.n3.a4.

- [Ceb24] Cesar Ceballos. A canonical realization of the alt ν -associahedron. 2024. arXiv:2401.17204.
- [Cha07] F. Chapoton. Sur le nombre d'intervalles dans les treillis de Tamari. *Sém. Lothar. Combin.*, 55:Art. B55f, 18, 2005/07. URL: <http://www.emis.de/journals/SLC/wpapers/s55chapoton.html>.
- [Che22] Clément Chenevière. Linear intervals in the Tamari and the Dyck lattices and in the alt-Tamari posets. 2022. arXiv:2209.00418.
- [CPS19] Cesar Ceballos, Arnau Padrol, and Camilo Sarmiento. Geometry of ν -Tamari lattices in types A and B . *Trans. Amer. Math. Soc.*, 371(4):2575–2622, 2019. doi:10.1090/tran/7405.
- [CPS20] Cesar Ceballos, Arnau Padrol, and Camilo Sarmiento. The ν -Tamari lattice via ν -trees, ν -bracket vectors, and subword complexes. *Electron. J. Combin.*, 27(1):Paper No. 1.14, 31, 2020. doi:10.37236/8000.
- [FPR17] Wenjie Fang and Louis-François Préville-Ratelle. The enumeration of generalized Tamari intervals. *European J. Combin.*, 61:69–84, 2017. doi:10.1016/j.ejc.2016.10.003.
- [Hai94] Mark D. Haiman. Conjectures on the quotient ring by diagonal invariants. *J. Algebraic Combin.*, 3(1):17–76, 1994. doi:10.1023/A:1022450120589.
- [Kre65] Germain Kreweras. Sur une classe de problèmes de dénombrement liés au treillis des partitions des entiers. *Cahiers du Bureau universitaire de recherche opérationnelle Série Recherche*, 6:9–107, 1965. URL: <http://eudml.org/doc/272631>.
- [MHPS12] Folkert Müller-Hoissen, Jean Marcel Pallo, and Jim Stasheff, editors. *Associahedra, Tamari lattices and related structures*, volume 299 of *Progress in Mathematics*. Birkhäuser/Springer, Basel, 2012. Tamari memorial Festschrift. doi:10.1007/978-3-0348-0405-9.
- [PRV17] Louis-François Préville-Ratelle and Xavier Viennot. The enumeration of generalized Tamari intervals. *Trans. Amer. Math. Soc.*, 369(7):5219–5239, 2017. doi:10.1090/tran/7004.

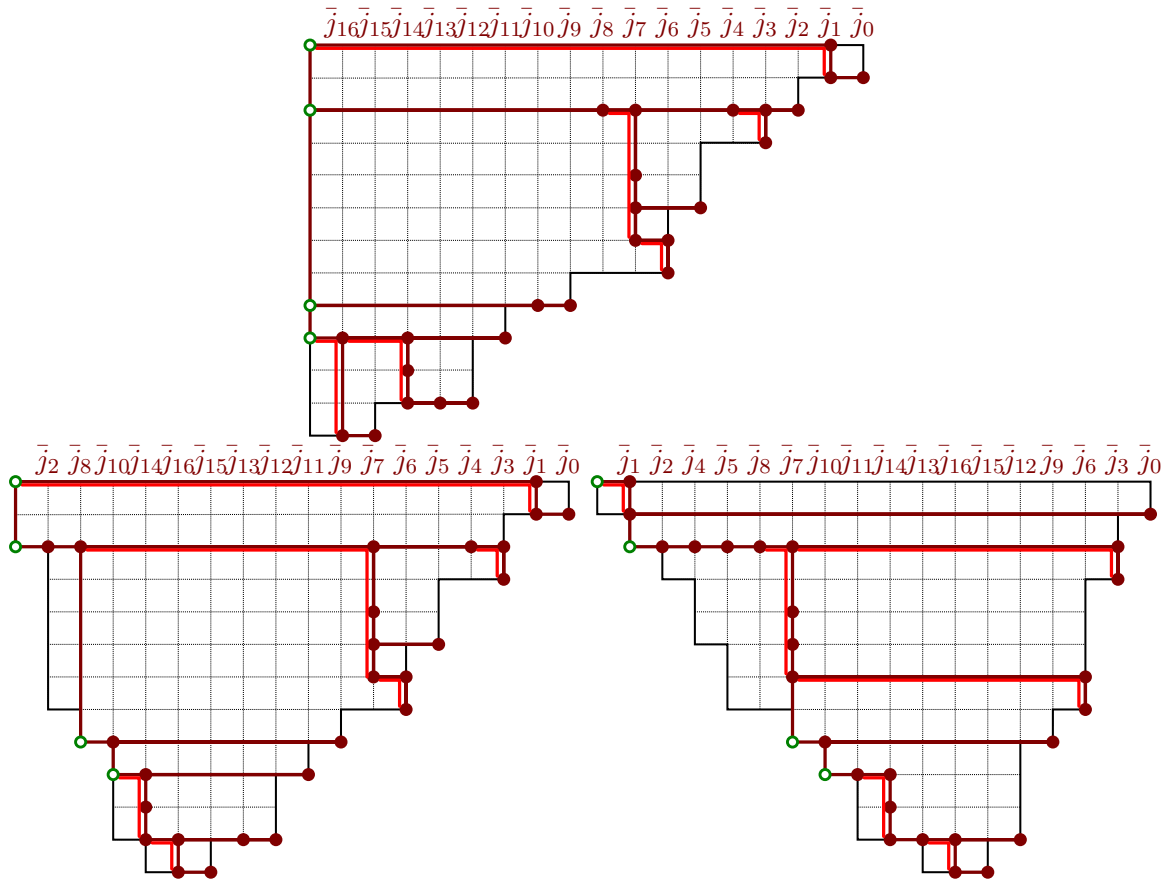


Figure 6.3: Bijection between right intervals for $\nu = (2, 3, 0, 1, 2, 3, 0, 1, 0, 2, 1, 2, 0)$, and for δ^{max} , $\delta = (2, 0, 1, 1, 2, 0, 1, 0, 2, 0, 2, 0)$, and $\delta' = (1, 0, 0, 1, 1, 0, 0, 0, 2, 0, 1, 0)$. The three trees have reduced column vector $(0, 1, 0, 1, 0, 0, 1, 3, 0, 0, 0, 0, 0, 0, 2, 0, 1)$.