BIJECTIVE COUNTING OF PLANE BIPOLAR ORIENTATIONS AND SCHNYDER WOODS

ÉRIC FUSY, DOMINIQUE POULALHON, AND GILLES SCHAEFFER

Abstract. A bijection Φ is presented between plane bipolar orientations with prescribed numbers of vertices and faces, and non-intersecting triples of upright lattice paths with prescribed extremities. This yields a combinatorial proof of the following formula due to R. Baxter for the number Θ_{ij} of plane bipolar orientations with i non-polar vertices and j inner faces:

$$\Theta_{ij} \;=\; 2\,\frac{(i+j)!\,(i+j+1)!\,(i+j+2)!}{i!\,(i+1)!\,(i+2)!\,j!\,(j+1)!\,(j+2)!}.$$

In addition, it is shown that Φ specializes into the bijection of Bernardi and Bonichon between Schnyder woods and non-crossing pairs of Dyck words.

This is the extended and revised journal version of a conference paper with the title "Bijective counting of plane bipolar orientations", which appeared in Electr. Notes in Discr. Math. pp. 283-287 (proceedings of Eurocomb'07, 11-15 September 2007, Sevilla).

1. Introduction

A *bipolar orientation* of a graph is an acyclic orientation of its edges with a unique *source* s and a unique *sink* t, *i.e.,* such that s is the only vertex without incoming edge, and t the only one without outgoing edge; the vertices s and t are the *poles* of the orientation. Alternative definitions, characterizations, and several properties are given by De Fraysseix *et al* in [19]. Bipolar orientations are a powerful combinatorial structure and prove insightful to solve many algorithmic problems such as planar graph embedding [18, 8] and geometric representations of graphs in various flavours (*e.g.,* visibility [20], floor planning [17], straight-line drawing [21, 13]). Thus, it is an interesting issue to have a better understanding of their combinatorial properties.

This article focuses on the enumeration of bipolar orientations in the planar case: we consider bipolar orientations on planar maps, where a *planar map* is a connected graph embedded in the plane (i.e., drawn with no edge-intersection, the drawing being considered up to isotopy). A *plane bipolar orientation* is a pair (M, X), where M is a planar map and X is a bipolar orientation of M having its poles incident to the outer face of M, see Figure 1. Let Θ_{ij} be the number of plane bipolar orientations with i non-pole vertices and j inner faces. R. Baxter proved in [1, Eq 5.3] that Θ_{ij} satisfies the following simple formula:

$$\Theta_{ij} \;=\; 2\,\frac{(i+j)!\,(i+j+1)!\,(i+j+2)!}{i!\,(i+1)!\,(i+2)!\,j!\,(j+1)!\,(j+2)!}. \tag{1}$$

Nevertheless his methodology relies on quite technical algebraic manipulations on generating functions, with the following steps: the coefficients Θ_{ij} are shown to satisfy an explicit recurrence (expressed with the help of additional "catalytic" parameters), which is translated to a functional equation on the associated generating functions. Then, solving the recurrence requires to solve the functional equation: Baxter guessed and checked the solution, while more recently M. Bousquet-Mélou described a direct computation way based on the so-called "obstinate kernel method" [6].

The aim of this article is to give a direct bijective proof of Formula (1). Our main result, Theorem 1, is the description of a bijection between plane bipolar orientations and certain triples of lattice paths, illustrated in Figure 1.

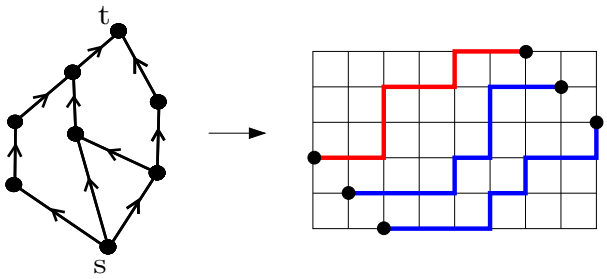

FIGURE 1. A plane bipolar orientation and the associated triple of non-intersecting upright lattice paths.

Theorem 1. *Plane bipolar orientations with i non-pole vertices and j inner faces are in bijection with non-intersecting triples of upright lattice paths on $\mathbb{Z}^2$ with respective origins $(-1,1)$, $(0,0)$, $(1,-1)$, and respective endpoints $(i-1,j+1)$, (i,j), $(i+1,j-1)$.*

This constitutes a proof of Formula (1), since the latter is easily derived from Theorem 1 using the Gessel–Viennot Lemma [14, 15]:

Lemma 2 (Gessel–Viennot)**.** *Let k be a positive integer, $\mathcal{A} = \{A_1, \ldots, A_k\}$ and $\mathcal{B} = \{B_1, \ldots, B_k\}$ be two sets of points on the $\mathbb{Z}^2$ lattice, such that any k-tuple of non-intersecting upright lattice paths with starting points in $\mathcal{A}$ and endpoints in $\mathcal{B}$ necessarily join together A_p and B_p for any index p. Then the number of such k-tuples is:*

$$\Theta = \mathrm{Det}(M),$$

where M is the $k \times k$ matrix such that M_{pq} is the number of upright lattice paths from A_p to B_q.

By Theorem 1, Θ_{ij} is equal to the number of triples of non-intersecting lattice paths from $A_1 = (-1,1), A_2 = (0,0), A_3 = (1,-1)$ to $B_1 = (i-1,j+1), B_2 = (i,j), B_3 = (i+1,j-1)$. Hence,

$$\Theta_{ij} \;=\; \begin{vmatrix} \binom{j+i}{i} & \binom{j+i}{i+1} & \binom{j+i}{i+2} \\ \binom{j+i}{i-1} & \binom{j+i}{i} & \binom{j+i}{i+1} \\ \binom{j+i}{i-2} & \binom{j+i}{i-1} & \binom{j+i}{i} \end{vmatrix} \;=\; \frac{2\ (i+j)!\ (i+j+1)!\ (i+j+2)!}{i!\ (i+1)!\ (i+2)!\ \ j!\ (j+1)!\ (j+2)!}.$$

The second main result of this paper is to show that our bijection extends in a natural way a bijection that has been recently described by Bernardi and Bonichon [2] (which itself reformulates an original construction due to Bonichon [4]) to count another well-known and powerful combinatorial structure related to planar maps, namely Schnyder woods on triangulations [11, Chapter 2]. Actually our construction draws much of its inspiration from the one in [2]. We recover the correspondence between these Schnyder woods and non-crossing pairs of Dyck paths, which easily yields the formula

$$S_n \;=\; C_n C_{n+2} \;-\; C_{n+1}^2 \;=\; \frac{6\ (2n)!\ (2n+2)!}{n!\ (n+1)!\ (n+2)!\ (n+3)!} \tag{2}$$

for the number S_n of Schnyder woods on triangulations with n inner vertices (where C_n denotes the nth Catalan number $(2n)!/(n!(n+1)!)$).

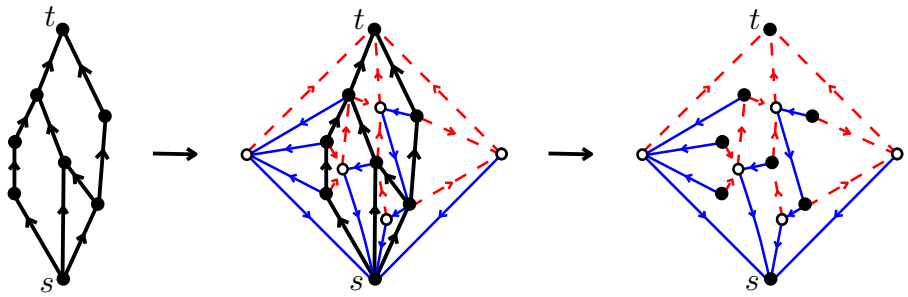

FIGURE 2. From a plane bipolar orientation to a separating decomposition.

Recent related work. Felsner et al [12] have very recently exhibited a whole collection of combinatorial structures that are bijectively related with one another, among which plane bipolar orientations, separating decompositions on quadrangulations, Baxter permutations, and triples of non-intersecting paths. Though very close in spirit, our bijection is not equivalent to the one exhibited in [12]. In particular, the restriction of this bijection to count Schnyder woods is a bit more involved than our one and is not equivalent to the bijection of Bernardi and Bonichon [2].

Even more recently, Bonichon et al [5] have described a simple and direct bijection between plane bipolar orientations and Baxter permutations. These Baxter permutations are known to be encoded by non-intersecting triples of lattice paths since work by Dulucq and Guibert [10]. Combining the bijections in [5] and [10] leads to yet another bijection (almost equivalent to the one in [12]) between plane bipolar orientation and non-intersecting triple of paths.

The main steps to encode a plane bipolar orientation by a non-intersecting triple of paths. At first (Section 2), we recall a well-known bijective correspondence between plane bipolar orientations and certain decompositions of quadrangulations into two spanning trees, which are called *separating decompositions*. The next step (Section 3.1) is to encode such a separating decomposition by a triple of words with some prefix conditions: the first two words encode one of the two trees T, in a slight variation on well known previous results for the 2-parameter enumeration of plane trees or binary trees (counted by the so-called Narayana numbers). The third word encodes the way the edges of the other tree shuffle in the tree T. The last step (Section 3.2) of the bijection is to represent the triple of words as a triple of upright lattice paths, on which the prefix conditions translate into a non-intersecting property.

2. Reduction to counting separating decompositions on quadrangulations

A *quadrangulation* is a planar map with no loop nor multiple edge and such that all faces have degree 4. Such maps correspond to *maximal* bipartite planar maps, *i.e.,* bipartite planar maps that would not stay bipartite or planar if an edge were added between two of their vertices.

Let $O = (M, X)$ be a plane bipolar orientation; the *quadrangulation* Q of M is the bipartite map obtained as follows: say vertices of M are black, and put a white vertex in each face of M; it proves convenient in this particular context to define a special treatment for the outer face, and put two white vertices in it, one on the left side and one on the right side of M when the source and sink are drawn at the bottom and at the top, respectively. These black and white vertices are the vertices of Q, and the edges of Q correspond to the incidences between vertices and faces of M. This construction, which can be traced back to Brown and Tutte [7], is illustrated in Figure 2. It is well known that Q is indeed a quadrangulation: to each edge e of M corresponds an inner (*i.e.,* bounded) face of Q (the unique one containing e in its interior), and our particular treatment of the outer face also produces a quadrangle.

If M is endowed with a bipolar orientation O, this classical construction can be enriched to transfer the orientation on Q, as shown in Figure 2. Notice that O (or, in general, any plane

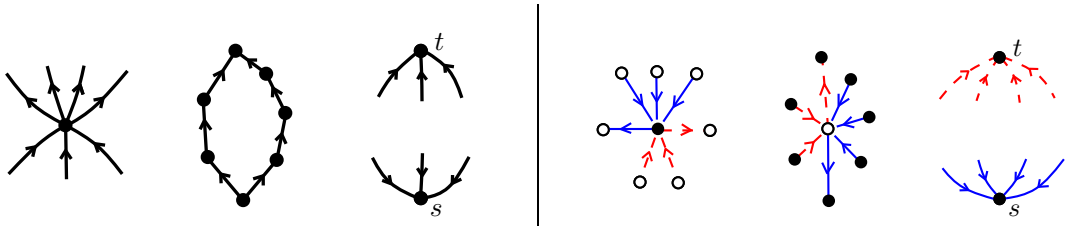

(a) in a plane bipolar orientation, (b) and in a separating decomposition.

FIGURE 3. The local rules.

bipolar orientation) satisfies the two following local conditions [9] illustrated in Figure 3(a), as easily proved using the acyclicity of the orientation and the Jordan curve theorem:

- edges incident to a non-pole vertex are partitioned into a non-empty block of incoming edges and a non-empty block of outgoing edges,
- dually, the contour of each inner face f consists of two oriented paths (one path has f on its left, the other one has f on its right); the common extremities of the paths are called the two *extremal vertices* of f.

A *separating decomposition* of Q is an orientation and bicoloration of its edges, say in red or blue, that satisfy the following local conditions illustrated in Figure 3(b) (in all figures, red edges are dashed):

- each inner vertex has exactly two outgoing edges, a red one and a blue one;
- around each inner black (white, *resp.*) vertex, the incoming edges in each color follow the outgoing one in clockwise (counterclockwise, *resp.*) order;
- all edges incident to s are incoming blue, and all edges incident to t are incoming red.

Given an inner face f of M, let us orient the two corresponding edges of Q from the white vertex w_f corresponding to f to the extremal vertices of f, and color respectively in red and blue the up- and the down-edges. The other edges incident to w_f are oriented and colored so as to satisfy the circular order condition around w_f. This defines actually a separating decomposition of Q, and this mapping from plane bipolar orientations to separating decompositions is one-to-one, as proved by an easy extension of [9, Theorem 5.3]:

Proposition 3. *Plane bipolar orientations with i non-pole vertices and j inner faces are in bijection with separating decompositions on quadrangulations with $i+2$ black vertices and $j+2$ white vertices.*

Accordingly, encoding plane bipolar orientations *w.r.t.* the numbers of vertices and faces is equivalent to encoding separating decompositions *w.r.t.* the numbers of black and white vertices.

3. Encoding a separating decomposition by a triple of non-intersecting paths

Separating decompositions have an interesting property: as shown in [3, 16], blue edges form a tree spanning all vertices but t, and red edges form a tree spanning all vertices but s. Moreover, the orientation of the edges corresponds to the natural orientation toward the root in both trees (the root is s for the blue tree and t for the red tree).

3.1. From a separating decomposition to a triple of words.

Let D be a separating decomposition with $i+2$ black vertices and $j+2$ white vertices, and let T_{blue} be its blue tree. A *clockwise* (or shortly *cw*) *traversal* of a tree is a walk around the tree with the outer face on the left. We define the *contour word* W_Q of Q as the word on the alphabet $\{a, \underline{a}, b, \underline{b}, c, \underline{c}\}$ that encodes the clockwise traversal of T_{blue} starting at s in the following manner (see Figure 4): letter a (b, *resp.*) codes the traversal of an edge e of T_{blue} from a black to a white vertex (from a white to a black

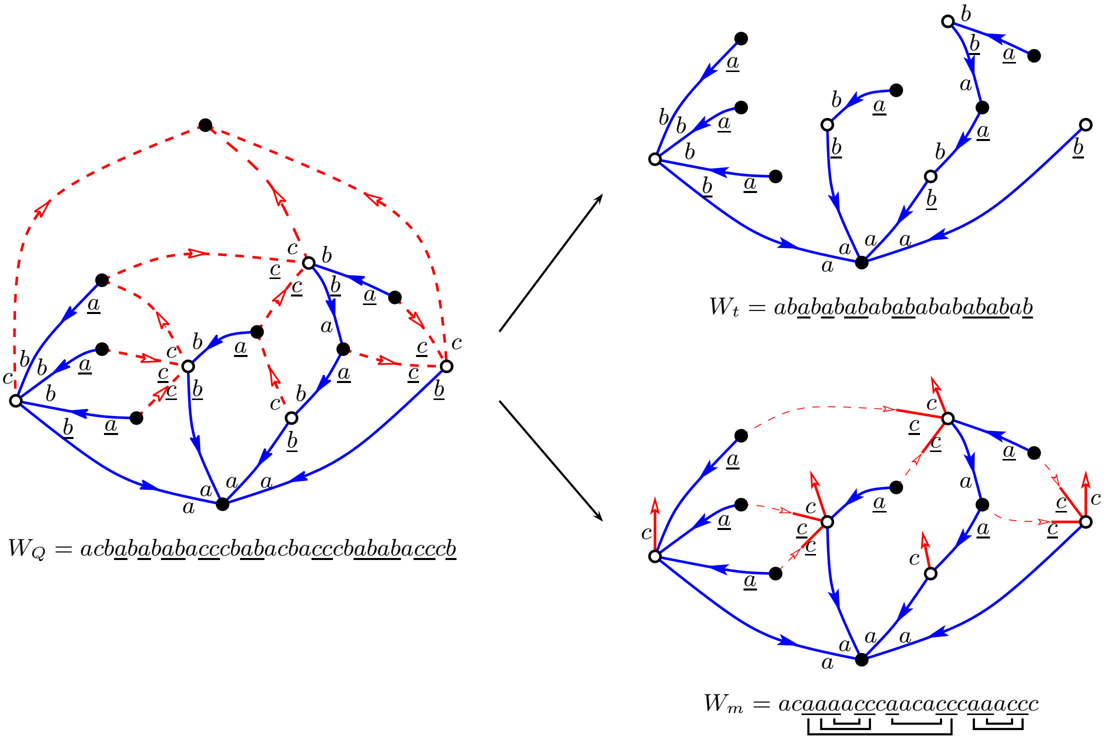

FIGURE 4. The words encoding a separating decomposition.

one, *resp.*), and the letter is underlined if it corresponds to the second traversal of e; letter c codes the crossing of red edge at a white vertex, and is underlined it if the edge is incoming.

We shall consider three subwords of W_Q: for any ℓ in $\{a, b, c\}$, let W_ℓ denote the subword obtained by keeping only the letters in the alphabet $\{\ell, \underline{\ell}\}$. In order to describe the properties of these words, we also introduce the *tree-word* W_{t} and the *matching word* W_{m}, that are respectively obtained from W_Q by keeping the letters in $\{a, \underline{a}, b, \underline{b}\}$, and in $\{a, \underline{a}, c, \underline{c}\}$.

3.1.1. *The tree-word encodes the blue tree.* Observe that W_{t} corresponds to a classical Dyck encoding of T_{blue}, in which the two alphabets $\{a, \underline{a}\}$ and $\{b, \underline{b}\}$ are used alternatively to encode the bicoloration of vertices. Hence W_{t} is just obtained by interlacing W_a and W_b starting with a, and each prefix of W_{t} has at least as many non-underlined letters as underlined letters.

Let us count precisely the number of occurrences of letters a, $\underline{a}$, b and $\underline{b}$ in W_{t}. For this purpose, let us associate each edge of a tree with its extremity that is farther from the root. From the defining rules it follows that the two traversals of edges corresponding to black vertices are encoded by b and $\underline{a}$, while those of edges corresponding to white vertices are encoded by a and $\underline{b}$. In other words, each occurrence of a letter a, $\underline{a}$, b, $\underline{b}$ corresponds to the first visit to a white vertex, last visit to a black vertex, first visit to a black vertex, and last visit to a white vertex, respectively. As T_{blue} has i non-root black vertices and $j + 2$ white vertices, the word W_a has $j + 2$ occurrences of a and i occurrences of $\underline{a}$, shortly written $W_a \in \mathfrak{S}(a^{j+2}\underline{a}^i)$. Similarly, $W_b \in \mathfrak{S}(b^i\underline{b}^{j+2})$. Furthermore, the fact that each prefix of W_{t} has at least as many non-underlined letters as underlined letters translates into the following property for the pair (W_a, W_b):

Property 1. *For $1 \leq k \leq i$, the number of a's on the left of the kth occurrence of $\underline{a}$ in W_a is strictly larger than the number of $\underline{b}$'s on the left of the kth occurrence of b in W_b.*

Proof. For each k, let $N_a(k)$ ($N_{\underline{b}}(k)$) be the number of a's (number of $\underline{b}$'s) in W_{t} on the left of the kth occurrence of $\underline{a}$ (kth occurrence of b, resp.). Let p be the prefix of W_{t} ending just before

the kth occurrence of $\underline{a}$. Notice that p ends at a letter in $\{b, \underline{b}\}$, so p has even length $2m$ with m letters in $\{a, \underline{a}\}$ and m letters in $\{b, \underline{b}\}$. Let m_a, $m_{\underline{a}}$, m_b, $m_{\underline{b}}$ be respectively the numbers of a's, $\underline{a}$'s, b's, and $\underline{b}$'s in p (notice that $m_{\underline{a}} = k-1$ and $m_a = N_a(k)$). Since $W_{\rm t}$ is a Dyck word and since p is followed by an underlined letter, we have $m_a + m_b > m_{\underline{a}} + m_{\underline{b}}$. But $m_{\underline{a}} = m - m_a$ and $m_b = m - m_{\underline{b}}$, so we obtain both (i): $m_{\underline{b}} < m_a = N_a(k)$ and (ii): $m_b > m_{\underline{a}} = k-1$. From (ii) the kth occurrence of b in $W_{\rm t}$ belongs to p, and from (i) the number $N_{\underline{b}}(k)$ of $\underline{b}$'s on its left is strictly smaller than $N_a(k)$. This concludes the proof. □

The words W_a and W_b have the additional property that two letters are redundant in each word. Indeed, the first and the last letter of W_a are a's and the last two letters of W_b are $\underline{b}$'s, because of the rightmost branch of $T_{\rm blue}$ being reduced to an edge, see Figure 4.

3.1.2. *The matching word encodes the red edges.* Let us now focus on W_c and on the matching word $W_{\rm m}$. Clearly, any occurrence of a letter c ($\underline{c}$) in W_Q corresponds to a red edge with white (black, *resp.*) origin, see Figure 4. Hence $W_c \in \mathfrak{S}(c^{j+2}\underline{c}^i)$. Moreover W_c starts and ends with a letter c, corresponding to the two outer red edges.

Observe also that any occurrence of a in $W_{\rm m}$, which corresponds to the first visit to a white vertex v, is immediately followed by a pattern $\underline{c}^\ell c$, with ℓ the number of incoming red edges at v. Hence $W_{\rm m}$ satisfies the regular expression:

$$W_{\rm m} \in ac(\underline{a}^* a \underline{c}^* c)^*, \tag{3}$$

where E^* denotes the set of all (possibly empty) sequences of elements from E. Notice that this property uniquely defines $W_{\rm m}$ as a shuffle of W_a and W_c.

Lemma 4. *Let S be a separating decomposition, with $T_{\rm blue}$ the tree induced by the blue edges. Consider a red edge e of S not incident to t, with b (w) the black (white,* resp.*) extremity of e. Then the last visit to b occurs before the first visit to w during a cw traversal around $T_{\rm blue}$ starting at s.*

Proof. First, the local conditions of separating decompositions ensure that e is connected to b (w) in the corner corresponding to the last visit to b (first visit to w, *resp.*). Hence we just have to prove that, if C denotes the unique simple cycle formed by e and edges of the blue tree, then the edge e is traversed from b to w when walking cw around C. Assume a contrario that e is traversed from w to b during a cw walk around C. If e is directed from b to w (the case of e directed from w to b can be treated similarly), then the local conditions of separating decompositions ensure that the red outgoing path $P(w)$ of w (*i.e.,* the unique oriented red path that goes from w to t) starts going into the interior of C. According to the local conditions, no oriented red path can cross the blue tree, hence $P(w)$ has to go out of C at b or at w: going out at w is impossible as it would induce a red circuit, going out at b contradicts the local conditions; hence either case yields a contradiction. □

Let us now consider a red edge $e = (b, w)$ with a black origin. The outgoing half-edge of e is in the corner of the last visit to b, encoded by a letter $\underline{a}$, while the incoming half-edge of e, which is encoded by a letter $\underline{c}$, is in the corner of the first visit to w. Hence, according to Lemma 4, the $\underline{a}$ occurs before the $\underline{c}$. In other words, the restriction of $W_{\rm m}$ to the alphabet $\{\underline{a}, \underline{c}\}$ is a parenthesis word (interpreting each $\underline{a}$ as an opening parenthesis and each $\underline{c}$ as a closing parenthesis), and each parenthesis matching corresponds to a red edge with a black origin, see Figure 4. According to the correspondence between the a's and the c's (see the regular expression (3) of $W_{\rm m}$), this parenthesis property of $W_{\rm m}$ is translated as follows:

Property 2. *For $1 \le k \le j+2$, the number of $\underline{a}$'s on the left of the kth occurrence of a in W_a is at least as large as the number of $\underline{c}$'s on the left of the kth occurrence of c in W_c.*

Definition. A triple of words (W_a, W_b, W_c) in $\mathfrak{S}(a^{j+2}\underline{a}^i) \times \mathfrak{S}(b^i\underline{b}^{j+2}) \times \mathfrak{S}(c^{j+2}\underline{c}^i)$ is said to be *admissible of type* (i, j) if W_a (W_c, *resp.*) ends with a letter a (c, *resp.*) and if Property 1 and Property 2 are satisfied.

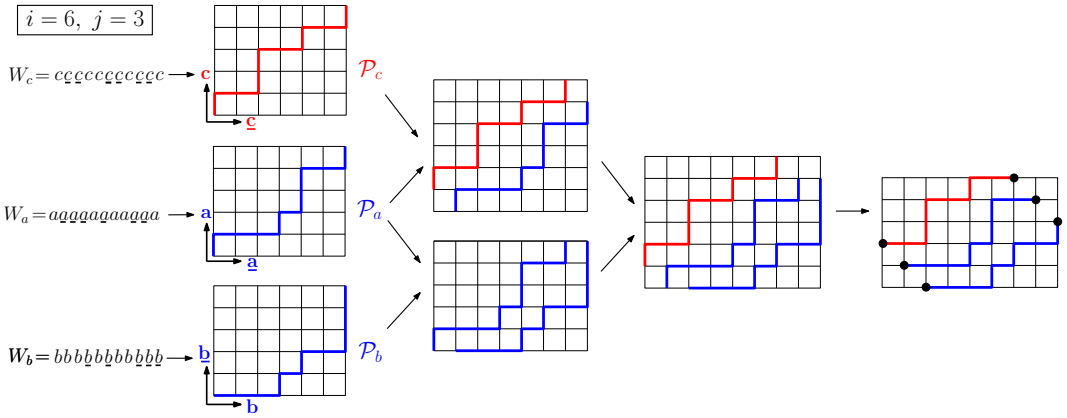

FIGURE 5. An admissible triple of words is naturally represented as a non-intersecting triple of paths.

Observe that this definition yields other redundant letters, namely, W_a has to start with a letter a, W_c has to start with a letter c, and W_b has to end with two letters $\underline{b}$.

3.2. **From an admissible triple of words to a triple of non-intersecting paths.** The properties of an admissible triple of words are formulated in a more convenient way on lattice paths. This section describes the correspondence, illustrated in Figure 5.

Consider an admissible triple of words (W_a, W_b, W_c) of type (i, j), and represent each word as an upright lattice path starting at the origin, the binary word being read from left to right, and the associated path going up or right depending on the letter. The letters associated to up steps are a, $\underline{b}$ and c. Clearly, as $(W_a, W_b, W_c) \in \mathfrak{S}(a^{j+2}\underline{a}^i) \times \mathfrak{S}(b^i\underline{b}^{j+2}) \times \mathfrak{S}(c^{j+2}\underline{c}^i)$, the three paths end at $(i, j + 2)$.

Property 1 is translated into:

> "for $1 \leq k \leq i$, the kth horizontal step of $\mathcal{P}_a$ (ending at abscissa k) is strictly above the kth horizontal step of $\mathcal{P}_b$."

Hence, Property 1 is equivalent to the fact that $\mathcal{P}_a$ and the shift of $\mathcal{P}_b$ one step to the right are non-intersecting.

Similarly, Property 2 is translated into:

> "for $1 \leq k \leq j + 2$, the kth vertical step of $\mathcal{P}_a$ is weakly on the right of the kth vertical step of $\mathcal{P}_c$."

In other words, $\mathcal{P}_c$ is weakly topleft of $\mathcal{P}_a$. Hence, Property 2 is equivalent to the fact that $\mathcal{P}_a$ and the shift of $\mathcal{P}_c$ one step up-left are non-intersecting. Let us now consider the redundant letters; they correspond to two vertical steps in each path, and removing them leads to a triple $(\mathcal{P}'_b, \mathcal{P}'_a, \mathcal{P}'_c)$ of non intersecting upright lattice paths with origins $(-1, 1)$, $(0, 0)$, $(1, -1)$ and endpoints $(i-1, j+1)$, (i, j), $(i + 1, j - 1)$. Such a triple of paths is called a *non-intersecting triple of paths of type* (i, j).

To sum up, we have described a mapping Φ from separating decompositions with $(i + 2)$ black and $(j + 2)$ white vertices to non-intersecting triples of paths of type (i, j).

4. The inverse mapping

As we show in this section, the mapping Φ is easily checked to be a bijection, as all steps (taken in reverse order) are invertible. Start from a non-intersecting triple of paths $(\mathcal{P}'_b, \mathcal{P}'_a, \mathcal{P}'_c)$ of type (i, j), where $\mathcal{P}'_b$ goes from $(-1, 1)$ to $(i - 1, j + 1)$, $\mathcal{P}'_a$ goes from $(0, 0)$ to (i, j), and $\mathcal{P}'_c$ goes from $(1, -1)$ to $(i + 1, j - 1)$. Append two up-steps in each of the 3 paths: $\mathcal{P}_b = \mathcal{P}'_b \uparrow\uparrow$, $\mathcal{P}_a = \uparrow \mathcal{P}'_a \uparrow$, $\mathcal{P}_c = \uparrow \mathcal{P}'_c \uparrow$.

4.1. **Associate an admissible triple of words to the triple of paths.** Each of the three paths $(\mathcal{P}_b, \mathcal{P}_a, \mathcal{P}_c)$ is equivalent to a binary word on the alphabet $\{u, r\}$, corresponding to the sequence of up and right steps when traversing the path. Let (W_a, W_b, W_c) be the three binary words associated respectively to $(\mathcal{P}_a, \mathcal{P}_b, \mathcal{P}_c)$. In order to have different alphabets for the three words, we substitute the alphabet (u, r) by $(a, \underline{a})$ for the word W_a, by $(\underline{b}, b)$ for the word W_b, and by $(c, \underline{c})$ for the word W_c. As the triple $(\mathcal{P}_b, \mathcal{P}_a, \mathcal{P}_c)$ is non-intersecting, the triple of words $W_a \in \mathfrak{S}(a^{j+2}\underline{a}^i), W_b \in \mathfrak{S}(b^i\underline{b}^{j+2}), W_c \in \mathfrak{S}(c^{j+2}\underline{c}^i)$ is readily checked to be an admissible triple of words of type (i, j).

4.2. **Construct the blue tree.** Define the *tree-word* W_{t} as the word obtained by interlacing W_a and W_b starting with a.

Claim 5. *The word W_{t} is a Dyck word (when seeing each letter in $\{a, b\}$ as opening parenthesis and each letter in $\{\underline{a}, \underline{b}\}$ as closing parenthesis).*

Proof. Clearly W_{t} has the same number of underlined as non-underlined letters. Assume that W_{t} is not a Dyck word, and consider the shortest prefix of W_{t} having more underlined letters than non-underlined letters. By minimality, the last letter of the prefix has to be underlined and is at an odd position $2m+1$, so that this letter is an $\underline{a}$. By minimality also, the prefix w_{2m} of length $2m$ has the same number of non-underlined letters as underlined letters. Moreover, w_{2m} has m letters in $\{a, \underline{a}\}$ and m letters in $\{b, \underline{b}\}$, because the letters of type $\{a, \underline{a}\}$ alternate with letters of type $\{b, \underline{b}\}$. Hence, if we denote by k the number of $\underline{a}$'s in w_{2m}, then w_{2m} has $m-k$ occurrences of a, k occurrences of b, and $m-k$ occurrences of $\underline{b}$. In particular, the number of occurrences of a on the left of the $(k+1)$th occurrence of $\underline{a}$ in W_a is $(m-k)$, and the number of occurrences of $\underline{b}$ on the left of the $(k+1)$th occurrence of b is at least $(m-k)$. This contradicts Property 1. □

Denote by T_{blue} the plane tree whose Dyck word is W_{t}. Actually, as we have seen in Section 3.1.1, W_{t} is a refined Dyck encoding of T_{blue} that also takes account of the number of vertices at even depth, colored black, and the number of vertices at odd length, colored white. Precisely, T_{blue} has $i+1$ black vertices and $j+2$ white vertices. Denote by s the (black) root of T_{blue}, and orient all the edges of T_{blue} toward the root.

4.3. **Insert the red half-edges.** The next step is to insert the red edges. Precisely we first insert the red *half-edges* (to be merged into complete red edges). Define the *matching word* W_{m} as the unique shuffle of W_a and W_c that satisfies the regular expression $ac(\underline{a}^*a\underline{c}^*c)^*$. For $1 \le k \le j+2$, consider the kth white vertex w in T_{blue}, the vertices being ordered *w.r.t.* the first visit during a cw traversal of T_{blue} starting at s. Let $\ell \ge 0$ be the number of consecutive $\underline{c}$'s that follow the kth occurrence of a in W_{m}. Insert ℓ incoming and one outgoing red half-edges (in clockwise order) in the corner of T_{blue} traversed during the first visit to w. Then, add an outgoing red half-edge to each black vertex b in the corner traversed during the last visit to b. The red half-edges are called *stems* as long as they are not completed into complete red edges, which is the next step. Observe that the local conditions of a separating decomposition are already satisfied around each vertex (the pole t is not added yet).

4.4. **Merge the red stems into red edges.** Next, we match the outgoing red stems at black vertices and the incoming red stems (which are always at white vertices). Property 2 ensures that the restriction of W_{m} to the alphabet $\{\underline{a}, \underline{c}\}$ is a parenthesis word, viewing each $\underline{a}$ as an opening parenthesis and each $\underline{c}$ as a closing parenthesis. By construction, this word corresponds to walking around T_{blue} and writing a $\underline{a}$ for each last visit to a black vertex and a $\underline{c}$ for each incoming red stem.

This yields a matching of the red half-edges; the red outgoing half-edge inserted in the corner corresponding to the kth black vertex (black vertices are ordered *w.r.t.* the last visit in T_{blue}) is merged with the incoming red half-edge associated with the letter $\underline{c}$ matched with the kth occurrence of $\underline{a}$ in W_{m}, see Figure 6(a). Such an operation is called a *closure*, as it "closes" a bounded face f on the right of the new red edge e. The origin of e is called the *left-vertex* of f.

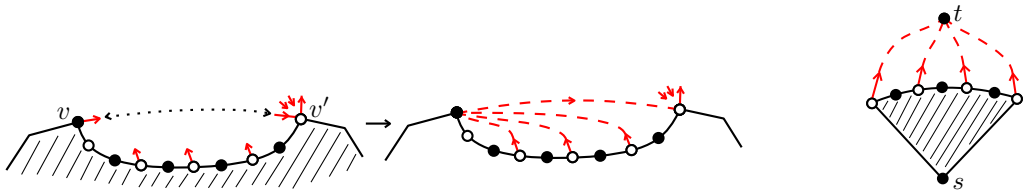

FIGURE 6. Completing the red stems going out of white vertices.

We perform the closures one by one, following an order consistent with the $\underline{a}$'s being matched inductively with the $\underline{c}$'s in $W_{\rm m}$. In Figure 4, this means that the red edges with a black origin are processed "from bottom to top". Observe that the planarity is preserved throughout the closures: the red edges that are completed are nested in the same way as the corresponding arches in the parenthesis word.

4.5. **Insert the remaining half-edges.** The last step is to complete the stems going out of white vertices into complete red edges going into black vertices, so as to obtain a quadrangulation endowed with a separating decomposition.

Lemma 6. *For each $k \in [0..i]$ consider the planar map F_k formed by the blue edges and the completed red edges after k closures have been performed. The following invariant holds.*

> *(I): "Consider any pair c_w, c_b of consecutive corners of F_k during a ccw traversal of the outer face of F_k (i.e., with the outer face on the right), such that c_w is incident to a white vertex (thus c_b is incident to a black vertex). Then exactly one of the two corners contains an outgoing (unmatched) stem."*

Proof. Induction on k. At the initial step, F_0 is the tree $T_{\rm blue}$. The red stems are inserted in the corners of $T_{\rm blue}$—as described in Section 4.4—in a way that satisfies the local conditions of separating decompositions. Hence it is an easy exercise to check that F_0 satisfies (I). Now assume that, for $k \in [0..i-1]$, F_k satisfies (I), and let us show that the same holds for F_{k+1}. Consider the closure that is performed from F_k to F_{k+1}. This closure completes a red edge $e = (b, w)$, where e starts from the corner c_b at the last visit to b and ends at the corner c_w at the first visit to w. As we see in Figure 6(a), the closure expels all the corners strictly between c_w and c_b from the outer face, and it makes c_b the new follower of c_w. According to the local conditions of separating decompositions, c_w contains an outgoing stem in the outer face of F_{k+1}. In addition, c_b contains no outgoing stem in F_{k+1}, because the outgoing stem of b is matched by the closure. Hence, F_{k+1} satisfies (I). □

Denote by $F = F_i$ the figure that is obtained after all closures have been performed (there are i closures, as each closure is associated with one of the i non-root black vertices of $T_{\rm blue}$). Note that each bounded face f of F has been "closed" by matching a red half-edge going out of a black vertex b with a red half-edge going into a white vertex w. The vertex b is called the *left-vertex* of f.

Let us now describe how to complete F into a separating decomposition on a quadrangulation. Add an isolated vertex t in the outer face of F. Taking advantage of Invariant (I), it is easy to complete suitably each red stem h going out of a white vertex:

- if h is in a bounded face f of F we complete h into an edge connected to the left-vertex of f; completing all the half-edges inside the face f splits f into quadrangular faces, as shown in Figure 6(a).

- if h is in the outer face of F we complete h into an edge connected to the vertex t; completing all such half-edges splits the outer face of F into quadrangular faces all incident to t, and t is incident to red incoming edges only, see Figure 6(b).

The planar map we obtain is thus a quadrangulation. In addition it is easy to check that the orientations and colors of the edges satisfy the local conditions of a separating decomposition. Indeed, the local conditions are satisfied in F. Afterwards the (black) left-vertex of each bounded face of F receives new incoming red edges in cw order after the red outgoing edge, and the vertex t receives red incoming edges only. Hence the local conditions remain satisfied after inserting the last red half-edges.

To sum up, we have described a mapping Ψ from non-intersecting triples of paths of type (i, j) to separating decompositions with $i + 2$ black vertices and $j + 2$ white vertices. It is easy to check step by step that the mapping Φ described in Section 3 and the mapping Ψ are mutually inverse. Together with Proposition 3, this yields our main bijective result announced in Theorem 1.

5. Specialization into a bijection for Schnyder woods

A *triangulation* is a planar map with no loop nor multiple edge such that each face is triangular. Given a triangulation T, let s, t, u be its outer vertices in cw order. A Schnyder wood on T is an orientation and coloration—in blue, red, or green—of the inner edges of T such that the following local conditions are satisfied (in the figures, blue edges are solid, red edges are dashed, and green edges are dotted):

- Each inner vertex v of T has exactly one outgoing edge in each color. The edges leaving v in color blue, green, and red, occur in cw order around v. In addition, the incoming edges of one color appear between the outgoing edges of the two other colors, see Figure 7(a).
- All the inner edges incident to the outer vertices are incoming, and such edges are colored blue, green, or red, whether the outer vertex is s, t, or u, respectively.

Definition, properties, and applications of Schnyder woods are given in Felsner's monograph [11, Chapter 2]. Among the many properties of Schnyder woods, it is well known that the subgraphs of T in each color are trees that span all the inner vertices and one outer vertex (each of the 3 outer vertices is the root of one of the trees).

We show here that Schnyder woods are in bijection with specific separating decompositions, and that such separating decompositions have one of the 3 encoding paths that is redundant, and the two other ones are Dyck paths. Afterward we show that this bijection is exactly the one recently described by Bernardi and Bonichon in [2] (which itself reformulates Bonichon's original construction [4]).

Starting from a Schnyder wood S with n inner vertices, we construct a separating decomposition $D = \alpha(S)$ as follows, see Figure 7:

- Split each inner vertex v of T into a white vertex w and a black vertex b that are connected by a blue edge going from b to w. In addition w receives the outgoing green edge, the outgoing blue edge and the incoming red edges of v, and b receives the outgoing red edge, the incoming blue edges, and the incoming green edges of v.
- Add a white vertex in the middle of the edge (s, t), and change the color of u from black to white.
- Recolor the green edges into red edges.
- Color red the two outer edges incident to t and orient these edges toward t. Color blue the two outer edges incident to s and orient these edges toward s.

Clearly we obtain from this construction a bipartite planar map Q with no multiple edge. The map Q has a quadrangular outer face, $2n$ inner vertices, and $4n$ inner edges (the $3n$ inner edges of the original triangulation plus the n new edges), hence Q has to be a maximal bipartite planar map, *i.e.,* Q is a quadrangulation. In addition, it is easily checked that Q is endowed with a separating decomposition $D = \alpha(S)$ via the construction, as shown in Figure 7. A separating

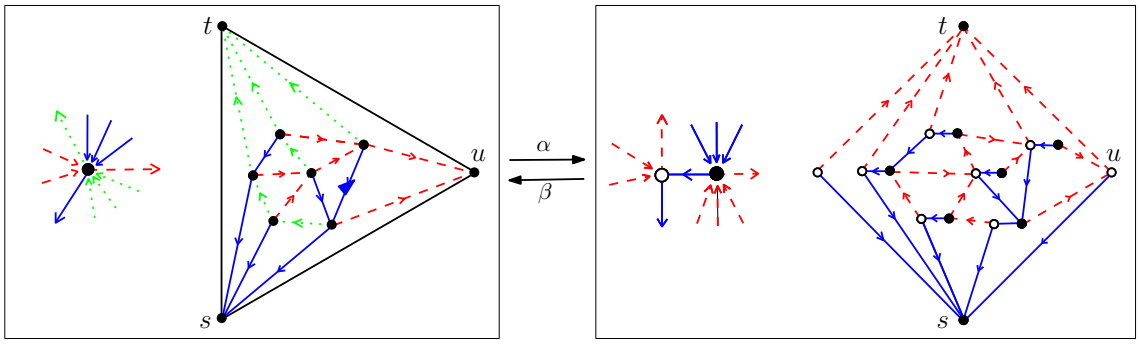

(a) (b)

FIGURE 7. From a Schnyder wood to a contractible separating decomposition.

decomposition is called *contractible* if each inner white vertex has blue indegree equal to 1 and the two outer white vertices have blue indegree 0. Clearly $D = \alpha(S)$ is contractible, see Figure 7(b).

Conversely, starting from a contractible separating decomposition D, we construct the associated Schnyder wood $S = \beta(D)$ as follows:

- recolor the red edges of D going out of white vertices into green edges.
- contract the blue edges going from a black to a white vertex.
- remove the colors and directions of the outer edges of D; contract into a single edge the path of length 2 going from s to t with the outer face on its left.

Clearly, the local conditions of Schnyder woods are satisfies by S. Hence, proving that S is a Schnyder wood comes down to proving that the planar map we obtain is a triangulation. In fact, it is enough to show that all faces are triangular (it is well known that a map with all faces of degree 3 and endowed with a Schnyder wood has no loop nor multiple edges), which clearly relies on the following lemma.

Lemma 7. *Take a contractible separating decomposition D and remove the path of length 2 going from s to t with the outer face on its left (which yields a separating decomposition with one inner face less). Then around each inner face there is exactly one blue edge going from a black vertex to a white vertex.*

Proof. Let O be the plane bipolar orientation associated to D. Observe that s and t are adjacent in O, the edge (s, t) having the outer face of O on its left. To each inner face f of O corresponds the unique edge e of O that is in the interior of f. For each edge e of O, except for (s, t), let ℓ_e (r_e) be the face of O on the left (right, *resp.*) of e, and let w_ℓ (w_r, *resp.*) be the corresponding white vertex on D. Notice that the inner face of D associated with e is the face f incident to the extremities of e and to the white vertices w_ℓ, w_r. As w_ℓ has blue indegree 1, ℓ_e has two edges on its right side. Hence, one extremity v of e is extremal for ℓ_e, and the other extremity v' of e is in the middle of the right side of ℓ_e. Hence the edge (v', w_ℓ), which is on the contour of f, is a blue edge with a black origin. In addition, the edge (w_ℓ, v) goes into v (as v is extremal for ℓ_e), and each of the other two edges of f is either red or is blue with w_r as origin, by the rules to translate a plane bipolar orientation into a separating decomposition. Hence any inner face of D, except the one corresponding to (s, t), has on its contour exactly one blue edge with a black origin. □

Clearly the mappings α and β are mutually inverse, so that we obtain the following result (which to our knowledge is new):

Proposition 8. *Schnyder woods with n inner vertices are in bijection with contractible separating decompositions with n black inner vertices.*

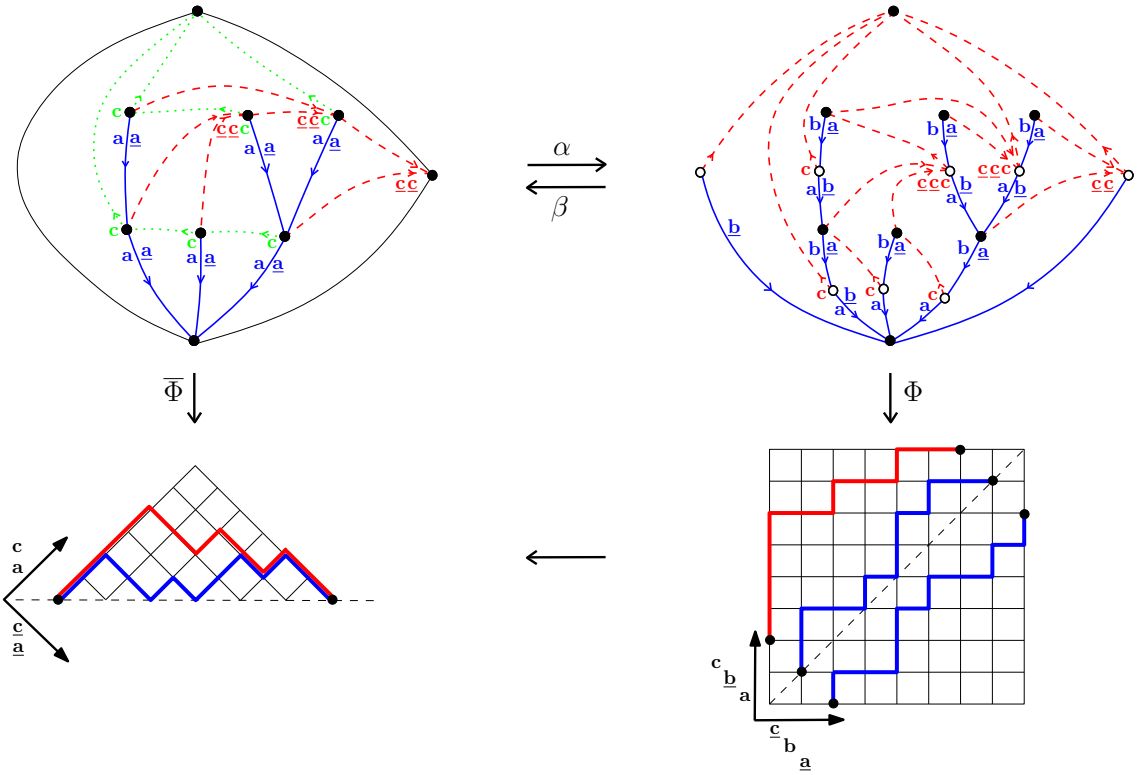

FIGURE 8. Encoding a Schnyder wood by two non-crossing Dyck paths via the associated contractible separating decomposition.

Let us now describe the nonintersecting triples of paths associated with contractible separating decompositions. Let D be a contractible separating decomposition with $2n$ inner vertices, and let $(P'_a, P'_b, P'_c) = \Phi(D)$ be the associated non-intersecting triple of paths, which has type (n, n). Let T_{blue} be the blue tree of D. Observe that T_{blue} has one 1-leg on the left and on the right side and all the other white vertices have exactly one child. Let T be the tree obtained from T_{blue} by deleting the 1-legs on each side and by merging each white vertex with its unique black child. Then it is easily checked that P'_a is the Dyck path encoding T. In addition, P'_b is redundant: it is obtained as the mirror of P'_a *w.r.t.* the diagonal $x = y$, shifted one step to the right, and with the last (up) step moved so as to prepend the path, see Figure 8 (right part). Finally the path P'_c is also a Dyck path, since it does not intersect P'_a and its respective endpoints are one step up-left of the corresponding endpoints of P'_a. To have a more classical representation, one rotates cw by 45 degrees the two paths P'_c and P'_a and shifts them to have the same starting point (and same endpoint), see Figure 8 (lower part). After doing this, the pair (P'_a, P'_c) is a noncrossing pair of Dyck paths (each of length $2n$) that is enough to encode the separating decomposition.

Conversely, starting from a pair (P'_a, P'_c) of non-crossing Dyck paths, we rotate the two paths ccw by 45 degrees and shift P'_c one step up-left, so that P'_c now does not intersect P'_a. Then we construct the path P'_b as the mirror of P'_a according to the diagonal $x = y$, with the last step moved to the start of the path, and we place P'_b so as to have its starting point one step bottom-right of the starting point of P'_a. As P'_a is a Dyck word (*i.e.,* stays weakly above the diagonal $x = y$), the path P'_b does not intersect P'_a. Furthermore it is easily checked that the blue tree T_{blue} of the separating decomposition $D = \Psi(P'_b, P'_a, P'_c)$ has one 1-leg on each side and all other white vertices have one child in T_{blue}. (Proof: by definition of Ψ, the Dyck path P for T_{blue} is obtained as a shuffle at even and odd positions of the path $P_a := \uparrow P'_a \uparrow$ and of the path $P_b := P'_b \uparrow\uparrow$. By construction of P'_b from P'_a, it is easily checked that there is a $\wedge$ at the beginning—starting at

position 0—and at the end of P and that all the other peaks and valleys of P start at odd position, hence the corresponding leaves and forks of T are at black vertices only.)

To conclude, we have proved that contractible separating decompositions with n inner vertices are encoded (via the bijection Φ) by noncrossing pairs of Dyck paths each having $2n$ steps. Given Proposition 8, we recover Bonichon's result [4]:

Theorem 9. *Schnyder woods on triangulations with n inner vertices are in bijection with noncrossing pairs of Dyck paths that have both $2n$ steps.*

As shown in Figure 8, the bijection can be formulated as a mapping $\overline{\Phi}$ operating directly on the Schnyder wood S. Indeed, let $D = \alpha(S)$. The blue tree T of S is equal to the tree T_{blue} of D where the 1-legs on each side are deleted and where each white node is merged with its unique black child. Hence the path P'_a (the lower Dyck path) associated with D is the Dyck path encoding the blue tree T of S. And the upper Dyck path P'_c can be read directly on S: P'_c is obtained by walking cw around T, drawing an up-step (down-step) each time an outgoing green edge (incoming red edge, *resp.*) is crossed, and completing the end of the path by down-steps. This mapping is exactly the bijection that has been recently described by Bernardi and Bonichon [2] for counting Schnyder woods (and more generally for counting some intervals of Dyck paths), which itself is a reformulation of Bonichon's original construction [4].

References

[1] R. J. Baxter. Dichromatic polynomials and potts models summed over rooted maps. *Annals of Combinatorics*, 5:17, 2001.

[2] O. Bernardi and N. Bonichon. Intervals in catalan lattices and realizers of triangulations. *Journal of Combinatorial Theory. Series A*, 116(1):55–75, 2009.

[3] T. Biedl and F. J. Brandenburg. Partitions of graphs into trees. In *Proceedings of Graph Drawing'06 (Karlsruhe)*, volume 4372 of *LNCS*, pages 430–439, 2007.

[4] N. Bonichon. A bijection between realizers of maximal plane graphs and pairs of non-crossing dyck paths. *Discrete Math.*, 298(1-3):104–114, 2005.

[5] N. Bonichon, M. Bousquet-Mélou, and É. Fusy. Baxter permutations and plane bipolar orientations. arXiv:0805.4180, 2008.

[6] M. Bousquet-Mélou. Four classes of pattern-avoiding permutations under one roof: Generating trees with two labels. *Elect. J. Comb.*, 2002.

[7] W. G. Brown and W.T. Tutte. On the enumeration of rooted nonseparable maps. *Canad. J. Math.*, 16:572–577, 1964.

[8] N. Chiba, T. Nishizeki, S. Abe, and T. Ozawa. A linear algorithm for embedding planar graphs using pq-trees. *J. Comput. Syst. Sci.*, 30(1):54–76, 1985.

[9] H. De Fraysseix, P. Ossona de Mendez, and P. Rosenstiehl. Bipolar orientations revisited. *Discrete Appl. Math.*, 56(2-3):157–179, 1995.

[10] S. Dulucq and O. Guibert. Baxter permutations. *Discr. Math.*, 180:143–156, 1998.

[11] S. Felsner. *Geometric graphs and arrangements*. Vieweg Verlag, 2004.

[12] S. Felsner, É. Fusy, M. Noy, and D. Orden. Bijections for baxter families and related objects. arXiv:0803.1546, 2008.

[13] É. Fusy. Straight-line drawing of quadrangulations. In *Proceedings of Graph Drawing'06*, volume 4372 of *LNCS*, pages 234–239. Springer, 2006.

[14] I. Gessel and X. Viennot. Binomial determinants, paths, and hook formulae. *Adv. Math*, 58:300–321, 1985.

[15] I. Gessel and X. Viennot. Determinants, paths, and plane partitions, 1989. Preprint.

[16] C. Huemer and S. Kappes. A binary labelling for plane laman graphs and quadrangulations. In *Proceedings of EWCG, Delphi*, pages 83–86, 2006.

[17] G. Kant and X. He. Regular edge labeling of 4-connected plane graphs and its applications in graph drawing problems. *Theoretical Computer Science*, 172(1-2):175–193, 1997.

[18] A. Lempel, S. Even, and I. Cederbaum. An algorithm for planarity testing of graphs. In *Theory of Graphs, Int. Symp (New York)*, pages 215–232, 1967.

[19] P. Ossona de Mendez. *Orientations bipolaires*. PhD thesis, Ecole des Hautes Etudes en Sciences Sociales, Paris, 1994.

[20] R. Tamassia and I. G. Tollis. A unified approach to visibility representations of planar graphs. *Discrete Comput. Geom.*, 1(4):321–341, 1986.

[21] R. Tamassia and I. G. Tollis. Planar grid embedding in linear time. *IEEE Trans. on Circuits and Systems*, CAS-36(9):1230–1234, 1989.

É. Fusy, G. Schaeffer: LIX, École Polytechnique, 91128 Palaiseau Cedex, France.
E-mail address: fusy@lix.polytechnique.fr, schaeffe@lix.polytechnique.fr

D. Poulalhon: LIAFA, Université Paris Diderot, case 7014, 75205 Paris Cedex 13, France.
E-mail address: Dominique.Poulalhon@liafa.jussieu.fr