WILEY

Growing uniform planar maps face by face

Alessandra Caraceni[1] | Alexandre Stauffer[2]

[1]Classe di Scienze, Scuola Normale Superiore, Pisa, Italy

[2]Department of Mathematical Sciences, University of Bath, Bath, UK

Correspondence
Alessandra Caraceni, Classe di Scienze, Scuola Normale Superiore, Pisa, Italy.
Email: alessandra.caraceni@sns.it

Abstract

We provide "growth schemes" for inductively generating uniform random $2p$-angulations of the sphere with n faces, as well as uniform random simple triangulations of the sphere with $2n$ faces. In the case of $2p$-angulations, we provide a way to insert a new face at a random location in a uniform $2p$-angulation with n faces in such a way that the new map is precisely a uniform $2p$-angulation with $n + 1$ faces. Similarly, given a uniform simple triangulation of the sphere with $2n$ faces, we describe a way to insert two new adjacent triangles so as to obtain a uniform simple triangulation of the sphere with $2n + 2$ faces. The latter is based on a new bijective presentation of simple triangulations that relies on a construction by Poulalhon and Schaeffer.

KEYWORDS

random generation, random planar maps, random trees, triangulations

1 | INTRODUCTION

The study of uniform random planar maps of given size conditioned to satisfy certain constraints is by now a sprawling subject rich with connections both internal and to other parts of mathematics and physics. Its combinatorial roots are firmly embedded in the works of Tutte, who in the 1960s lay the groundwork for an enumerative theory of planar maps [19,20]. Later works, whose original inspiration can in many cases be traced back to the celebrated bijection by Cori, Vanquelin and Schaeffer [11,18], have uncovered crucial links between certain classes of planar maps and classes of trees or tree-like structures. Armed with these bijections, probabilists have been able to embark on a deep investigation of the metric structure of large random planar maps and achieve very general results concerning scaling limits [13], local limits [4] as well as many other aspects.

The topic of *random generation* of planar maps is one that naturally inserts itself near the core of this subject, and one that we wish to further develop with this paper. More specifically, we are interested in generating a uniform map of size $n+1$ within a certain class by somehow *growing* it out of a uniform map of size n via a small local modification. Part of our motivation for considering this problem comes from investigating the mixing time of certain edge flip Markov chains on various classes of planar maps. A strategy for designing canonical paths using a uniform "growth" scheme has led us in previous papers [9,10] to consider this problem for the (rather special) case of quadrangulations; in a future paper, we shall leverage the results obtained here to achieve polynomial upper bounds for the mixing time of the edge flip Markov chain on $2p$-angulations and simple triangulations of the sphere.

Naturally, the problem of random generation within classes of planar maps has been considered before: many of the papers establishing links to branching structures do so with an eye towards efficient random generation (see e.g., [16]). However, few results have been obtained so far concerning iterative, increasing procedures, where a uniform map of size $n+1$ is coupled with a uniform map of size n in such a way that the larger map is always obtained from the smaller one by a small local (and random) modification that increases the number of faces. Other than the aforementioned paper [9], a notable exception is that of [1], which obtains such a growth procedure for irreducible quadrangulations of the hexagon and deduces a way to grow three–connected planar maps. In fact, the constructions of [9] and [1] are closely related, as they both stem from the same growth procedure for plane trees, or equivalently binary trees, which also appears as the simplest case in the paper [14]. Both [9] and [1] could be seen as very special cases of the present work, which, as we will later discuss, builds upon more sophisticated constructions from [14].

Another result in a similar direction is that of Bettinelli [7], who uses an explicit bijection to obtain a uniform quadrangulation with $n+1$ faces by performing a certain (random) surgery operation on a uniform quadrangulation with n faces, cutting it along a path, inserting a new face and re-glueing the edges of the cut in a slightly different way. Note that, though this operation may indeed count as a small random modification, it does affect a potentially large region of the quadrangulation, which makes it ill-suited to some applications (such as the one to edge flip chains mentioned above). Moreover, efforts of this kind have born fruits—to our knowledge—mainly in regards to quadrangulations, whose direct link to plane trees provides access to a more complete set of tools.

If one looks at random generation of uniform plane trees of fixed size, a rich panorama unfolds. From Rémy's algorithm for the generation of binary trees [17], to bijective approaches involving letter sequences and the cycle lemma [3], to the more probabilistic approach involving critical Galton–Watson trees, this topic continues to be investigated in its many facets, with very recent papers such as [15] providing further insights. Within this larger scope, one can consider the problem of generating uniform random trees in an iterative, increasing way, by grafting leaves onto smaller uniform trees; this is the approach of Lukzac and Winkler in the paper [14], which is one of the main sources of inspiration for this work. Note that the question of whether certain random trees can be coupled in an increasing way has been asked in various contexts and sometimes answered in the negative, see for example [12].

Our main objective within this paper is to bring increasing random generation in the style of [14] to more complicated classes of maps, and in particular to rooted $2p$-angulations of the sphere and rooted simple triangulations of the sphere. To this end, we better formalise the notion of a growth scheme (see Definitions 4.1 and 5.1), and we consolidate the existing setup of bijections between $2p$-angulations and labeled mobiles (see [8] and Section 2), as well as simple triangulations and blossoming trees (see [16] and Section 3), in order to obtain a direct correspondence between the notion of growing uniform $2p$-angulations (respectively, simple triangulations), and that of growing labeled mobiles (respectively,

blossoming trees). Finally, we explore the links between mobiles, blossoming trees and d-ary trees, relying on the main result of [14] in order to show

> **Theorem 1.** *For all $p \geq 2$, there exists a growth scheme for rooted $2p$-angulations of the sphere.*

as well as

> **Theorem 2.** *There exists a growth scheme for rooted simple triangulations of the sphere.*

Along the way to reaching the above results, we describe a construction proving the following proposition, highlighting new connections between combinatorial structures which might be of independent interest:

> **Proposition 1.1.** *There is an explicit $2n$-to-1 map (induced by Construction 3) from $\mathcal{T}_{n-1} \times \{\pm 1\}$ to STr_n, where $\mathcal{T}_{n-1}$ is the set of rooted plane trees with $n-1$ vertices of degree 5 and $3n-1$ leaves, and STr_n is the set of rooted simple triangulations of the sphere with $2n$ faces.*

Though Theorems 1 and 2 were obtained with a view of their applications to Markov chain mixing, we believe they might be a useful tool to add to the arsenal available in the investigation of planar maps. Moreover, they open up various avenues of further study: one might further their investigation in the direction of local limits similarly to [1], whose final questions can also apply to our growth schemes, or attempt to shed more light on the explicit underlying probabilities that are only characterised recursively. Growth schemes of this type could also be useful in developing graphical tools for the visualisation of large uniform triangulations or $2p$-angulations of the sphere, as they give a way to iteratively construct such maps "from the ground up", with each step being a small modification that does not strongly disrupt the overall structure.

The paper is organised as follows. Sections 2 and 3 contain the relevant notation, with compact reminders of known bijections and their most relevant features in our context. Section 4 explains how growing d-ary trees is equivalent to growing d-mobiles, which in turn lays the groundwork for a growth scheme for $2d$-angulations: the section culminates with the proof of Proposition 4.3, which yields Theorem 1 as a consequence. Finally, Section 5 proves Proposition 1.1 and Theorem 2.

2 | 2P-ANGULATIONS, MOBILES AND LABELED P-ARY TREES

In this section, we briefly define some of the combinatorial objects of interest and remind the reader of some relationships between them. The bijections we will be discussing are all known or easy to deduce from known facts, but we shall attempt a compact presentation of those parts of the constructions which will actually be involved in subsequent proofs. We refer readers to [8] for additional details.

> **Definition 2.1.** Given an integer $p \geq 2$, a p-angulation is a rooted planar map all of whose faces have a contour of length p. We denote the set of all p-angulations with n faces by M_n^p. We shall sometimes also consider the set of *pointed p-angulations* $\mathsf{M}_n^{p,\bullet}$, that is, the set of pairs (m, v), where $m \in \mathsf{M}_n^p$ and $v \in V(m)$ is a vertex of m.

Bijective approaches are especially successful in the case of p-angulations with p even, and more generally of classes of bipartite planar maps, which lend themselves well to a certain type of construction involving mobiles, introduced by Bouttier, Di Francesco and Guitter in [8]. (Though versions of the bijection and alternative constructions do exist in the case of odd p [2, 5, 6, 8], they do not in

general yield enumerative results of the same explicit nature as those one obtains in the case of even p.) We will discuss the special case of (simple) triangulations in the next section; this section will be mostly devoted to $2p$-angulations and to related branching structures.

In connection to $2p$-angulations, we will consider various types of plane trees.

> **Definition 2.2.** Let $d \geq 2$ and $n \geq 1$ be integers. A (complete) d-ary tree is a rooted plane tree each of whose vertices is of one of three types: the *origin*, which carries the root corner and has degree d; the *leaves*, which have degree 1; the internal vertices, which have degree $d + 1$. We say that a complete d-ary tree τ has *size* $|\tau| = n$, where $n \geq 1$, if it has $n - 1$ internal vertices.
>
> We denote by T_n^d the set of all complete d-ary trees of size n (see Figure 1), and conventionally define T_0^d as the set containing the tree with only one vertex, which we will sometimes call the *degenerate d-ary tree*, in spite of it not satisfying the definition above.

Remark that a tree in T_n^d has a total of $(d-1)n+1$ leaves, and therefore $dn+1$ vertices and dn edges. This is easy to see using the following construction, which we will repeatedly discuss throughout the paper: given a complete d-ary tree τ and a leaf v of τ, we denote by $\mathrm{grow}(\tau, v)$ the complete d-ary tree of size $|\tau| + 1$ obtained by turning v into an internal vertex, grafting d leaves onto it (right part of Figure 1); we say that $\mathrm{grow}(\tau, v)$ is obtained by *growing τ at v*. The fact that every tree in T_{n+1}^d can be obtained from some tree in T_n^d by growing it at some leaf immediately yields the number of leaves and vertices given above, by induction.

> *Remark* 2.1. It is worth remarking that a complete d-ary tree of size $n > 0$, as presented by Definition 2.2, can easily be reinterpreted as a pointed tree with n vertices, all labeled save for the origin: each vertex v has at most d children, each of which is assigned a different label in $\{1, 2, \ldots, d\}$. Indeed, this interpretation of a (complete) d-ary tree comes from numbering sibling vertices in left-to-right order and then erasing all the leaves (Figure 1). In this alternative interpretation, the operation of growing $\tau \in \mathsf{T}_n^d$ at a leaf v corresponds to making v into an internal vertex, thus adding it and its label to our labeled tree.
>
> It can be useful to have this alternative view in mind, particularly when referring to the paper [14], which predominantly uses it. The interpretation of T_0^d is a formal matter,
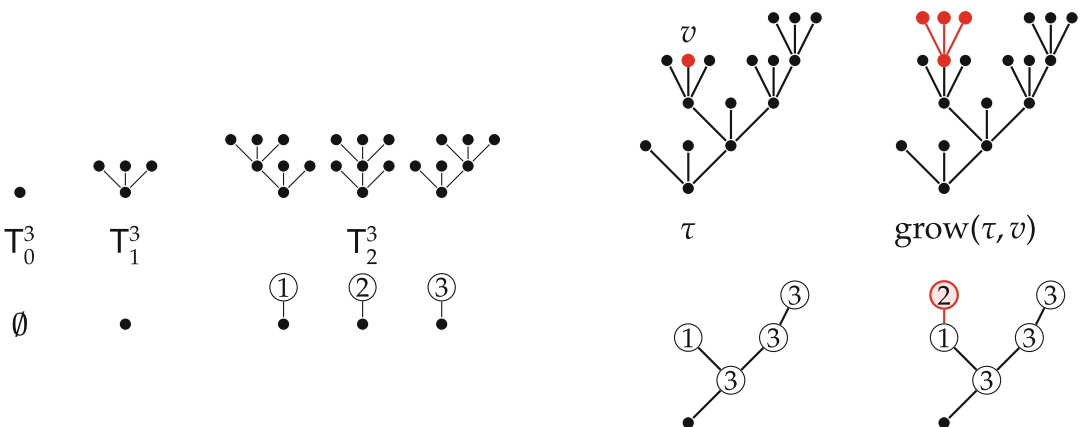

FIGURE 1 On the left, the set of complete three-ary trees of size 0, 1, and 2, reinterpreted below as pointed labeled trees. On the right, growing a three-ary tree at a leaf v.

but in the context of labeled trees we would interpret the degenerate complete d-ary tree as the empty tree with no vertices and no edges, which can only be grown into the tree with one (distinguished) vertex and no edges, corresponding to the single element of T_1^d.

Finally, it is time to introduce a branching structure that is closer to the original presentation of the BDFG bijection:

Definition 2.3. Let $p \geq 2$ and $n \geq 1$ be integers. An *unlabeled p-mobile of size n* is a rooted plane tree τ with the following properties:

1. vertices of $V(\tau)$ are colored either white or black and each edge of τ has a black endpoint and a white endpoint;
2. the origin is a white vertex;
3. each black vertex has degree p;
4. τ has n black vertices.

A *labeled p-mobile of size n* is an unlabeled p-mobile of size n endowed with a labeling l of its black vertices. Each black vertex is assigned an element in the set $S_p \subset \{\pm 1\}^{2p-1}$ of strings of ± 1's that have length $2p-1$ and sum 1 (or equivalently an integer between 1 and $\binom{2p-1}{p}$). We call Mob_n^p the set of all unlabeled p-mobiles of size n and LMob_n^p the set of all labeled p-mobiles of size n.

Conventionally, we let $\mathsf{Mob}_0^p = \mathsf{LMob}_0^p$ be the singleton of the *degenerate p-mobile* consisting of a single white vertex.

Given a corner c of an unlabeled p-mobile τ that is adjacent to a white vertex, we let $\mathrm{grow}(\tau, c)$ be the p-mobile obtained by grafting a size 1 p-mobile (which is a star with a central black vertex and p arms, rooted at one of the leaf corners) at the corner c (see Figure 2). If τ is a labeled p-mobile, for each $l \in S_p$ we let $\mathrm{grow}(\tau, c, l)$ be the labeled p-mobile obtained by grafting the star as above and labeling the new black vertex with the label l.

We now make some observations about unlabeled d-mobiles of size n, their enumeration and their relation to d-ary trees. We will then discuss how labeled d-mobiles can be interpreted as $2d$-angulations.

Remark 2.2. Any element $\tau \in \mathsf{Mob}_n^p$, where $n \geq 1$, can be represented as an ordered p-tuple $(\tau_1, \dots, \tau_p) \in \cup_{n_1 + \dots + n_p = n-1} \mathsf{Mob}_{n_1}^p \times \dots \times \mathsf{Mob}_{n_p}^p$ by considering the rightmost
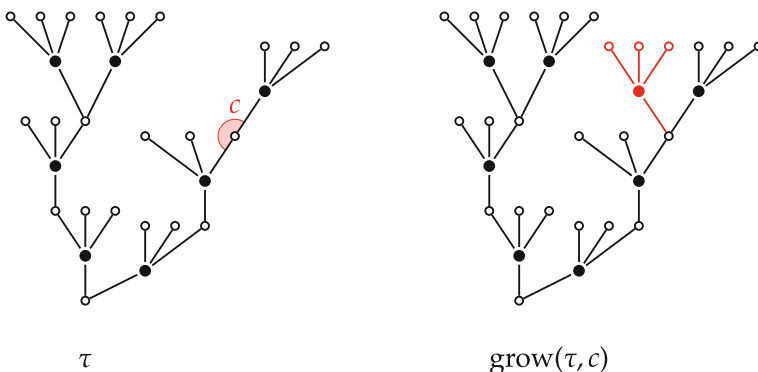

$$\tau \qquad\qquad \mathrm{grow}(\tau, c)$$

FIGURE 2 On the left, an unlabeled mobile $\tau \in \mathsf{Mob}_7^4$ with a marked corner c. On the right, the mobile $\mathrm{grow}(\tau, c) \in \mathsf{Mob}_8^4$ obtained by grafting the size one four-mobile onto the corner c.

black vertex of height 1 in τ, denoted by v, and letting τ_1 be the mobile obtained by erasing v and all of its descendants; we then let $\tau_2, \dots, \tau_p$, be the submobiles of descendants of the $p-1$ children of v, ordered from left to right. This immediately yields the recursion

$$|\mathsf{Mob}_n^p| = \sum_{n_1 + \dots + n_p = n-1} \prod_{i=1}^p |\mathsf{Mob}_{n_i}^p|.$$

Clearly, complete p-ary trees follow the same recursive structure (a tree in T_n^p corresponds to the p-tuple of subtrees of descendants of the p children of the origin), and we have $|\mathsf{Mob}_0^p| = |\mathsf{T}_0^p| = 1$, which yields that

$$|\mathsf{Mob}_n^p| = |\mathsf{T}_n^p| = \frac{1}{(p-1)n+1} \binom{pn}{n}. \tag{1}$$

Not only that, but the recursion yields an explicit bijection $\Phi : \mathsf{Mob}_n^p \to \mathsf{T}_n^p$ that one can define inductively: given $\tau \in \mathsf{Mob}_n^p$, $\Phi(\tau)$ is obtained by considering the complete p-ary trees $\Phi(\tau_1), \dots, \Phi(\tau_p)$ and grafting them in order (from left to right) onto the p leaves of a complete p-ary tree of size 1 (see Figure 3 for an example of this correspondence).

As for labeled p-mobiles, we have that $|\mathsf{LMob}_n^p| = \binom{2p-1}{p}^n |\mathsf{Mob}_n^p|$, and the bijection Φ can be used to interpret elements of LMob_n^p as pairs (t, l), where $t \in \mathsf{T}_n^p$ and l is a map from the set of vertices of t that are not leaves to the set $\left\{ 1, \dots, \binom{2p-1}{p} \right\}$.

The final link we need to describe is the one between $2p$-angulations and labeled p-mobiles, which has been thoroughly explored [8]. We shall outline the BDFG construction of a pointed, rooted $2p$-angulation $\Psi(\tau, \varepsilon) \in \mathsf{M}_n^{2p, \bullet}$ from a pair $(\tau, \varepsilon) \in \mathsf{LMob}_n^p \times \{+1, -1\}$ and summarize at the end of this section the main features of this correspondence Ψ that we will need.

Construction 1 (Bouttier, Di Francesco, Guitter). Given a pair $(\tau, \varepsilon) \in \mathsf{LMob}_n^p \times \{\pm 1\}$, construct the map $\Psi(\tau)$ as follows:
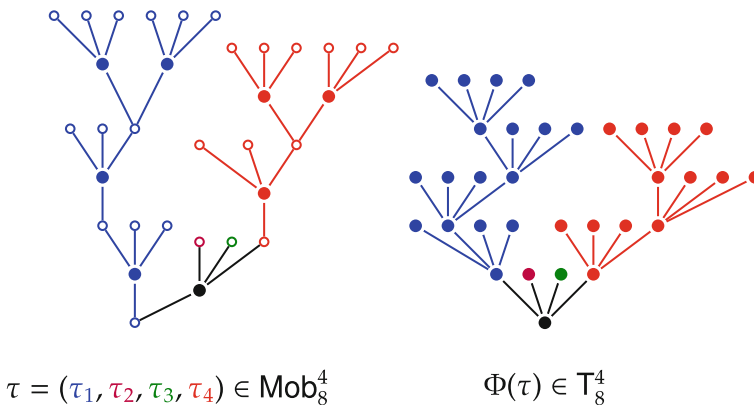

$$\tau = (\tau_1, \tau_2, \tau_3, \tau_4) \in \mathsf{Mob}_8^4 \qquad \Phi(\tau) \in \mathsf{T}_8^4$$

FIGURE 3 On the left, a mobile $\tau \in \mathsf{Mob}_8^4$ and its decomposition into the blue, purple, green and red submobiles $\tau_1 \in \mathsf{Mob}_4^4$, $\tau_2 \in \mathsf{Mob}_0^4$, $\tau_3 \in \mathsf{Mob}_0^4$, $\tau_4 \in \mathsf{Mob}_3^4$. On the right, the corresponding four-ary tree $\Phi(\tau)$ obtained by grafting $\Phi(\tau_1), \Phi(\tau_2), \Phi(\tau_3), \Phi(\tau_4)$ onto the four children of the origin.

1. first, deduce labels for the white vertices from the labels of the black vertices, as follows. Label the origin, which is a white vertex, with zero. Suppose v is a black vertex labeled $l = (l_1, l_2, \ldots, l_{2p-1}) \in S_p \subset \{\pm 1\}^{2p-1}$ such that the label of its (white) parent has already been determined and is equal to a. Let $i_1 < i_2 < \ldots < i_{p-1}$ be such that $l_{i_1} = \ldots = l_{i_{p-1}} = -1$. Label the $p-1$ (white) children of v from left to right as $a + i_j - 2j$ for $j = 1, \ldots, p-1$. Continue until all white vertices are labeled with an integer (see Figure 4).

2. Draw an additional white vertex δ.

3. For each corner c of a white vertex other than δ, draw a new "map edge" $e(c)$ issued from c. Assuming the white vertex carrying c is labeled a, let $t(c)$ (the "target corner" of c) be the first corner of a white vertex labeled $a - 1$ in a clockwise contour of the mobile started from the corner c, if such a corner exists. If the label $a - 1$ does not appear on white vertices, then set $t(c)$ to be the corner of δ. For each c, draw an edge $e(c)$ joining c to $t(c)$, taking care to perform this operation in such a way that edges do not cross (which is always possible).

4. Erase all edges of the original mobile and all black vertices. Forget all labels. Root the ensuing map in the edge drawn from the root corner of τ, oriented away from the corner if $\epsilon = 1$ and towards the corner if $\epsilon = -1$. Make δ the distinguished vertex.

For most applications, an important feature of Construction 1 is the significance of the labels in terms of graph distances on the $2p$-angulation. This will not be the case in this paper; the main features we need are given by

Proposition 2.1. *For each $n \geq 1$, the map Ψ of Construction 1 is a bijection from the set $\mathsf{LMob}_n^p \times \{\pm 1\}$ to the set $\mathsf{M}_n^{2p,\bullet}$ of pointed, rooted $2p$-angulations with n faces. Given $\tau \in \mathsf{LMob}_n^p$, this bijection induces a correspondence between black vertices of τ and faces of $\Psi(\tau)$ and between white vertices of τ and vertices of $\Psi(\tau)$ other than the distinguished vertex δ given by the pointing. Moreover, it induces a correspondence between corners of black vertices of τ and corners of $\Psi(\tau)$, and between corners of white vertices of τ and edges of $\Psi(\tau)$. The root corner of τ corresponds to the root edge of $\Psi(\tau)$.*
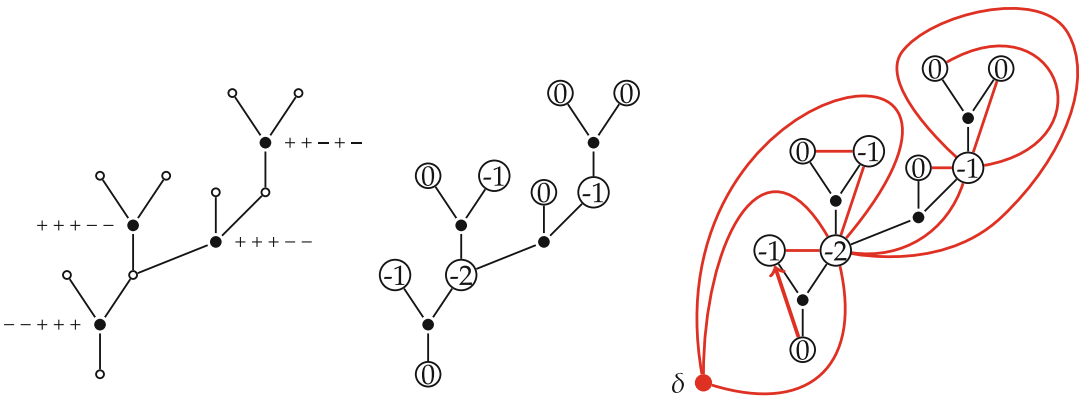

FIGURE 4 On the left, a labeled mobile $\tau \in \mathsf{LMob}_4^3$. In the middle, the labels of black vertices are turned into integer labels for white vertices. On the right, the full construction of the pointed six-angulation $\Psi(\tau, 1)$, consisting of the red edges, white vertices and distinguished vertex δ. Note how every face contains a single black vertex and has degree 6.
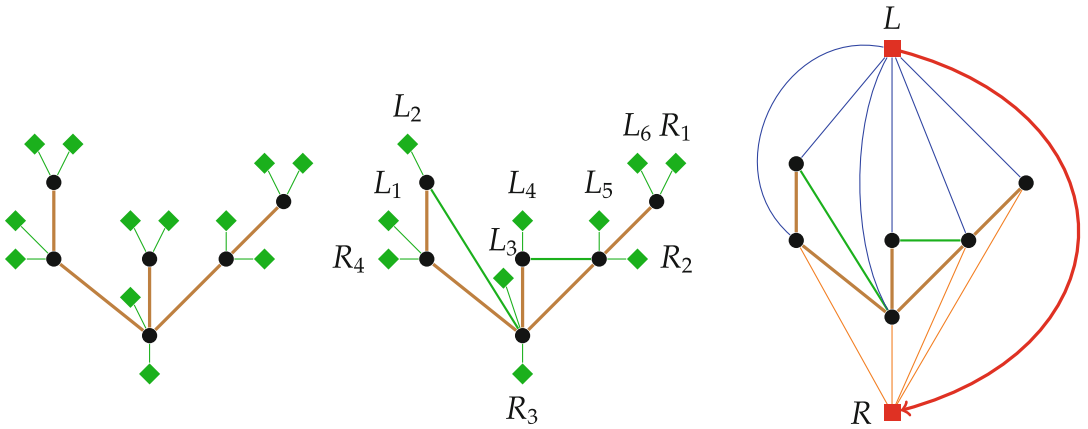
FIGURE 5 On the left, a blossoming tree $t \in \mathsf{BT}_6$. In the middle, two closure operations have been performed on t to create two new triangular faces. No other closure operations are possible and the remaining 10 blossoms have been labeled L_1 through L_6 and R_1 through R_4. Finally, on the right, the triangulation $\Xi(t, 1) \in \mathsf{STr}_6$. Its root edge is marked in red; the thicker brown edges are those involving non-blossoms in the original tree. The two green edges are the ones obtained from the two initial closures and the blue and orange edges are obtained from the remaining blossoms during the last steps of the construction.

3 | SIMPLE TRIANGULATIONS AND BLOSSOMING TREES

Definition 3.1. A simple triangulation of size n is a rooted map in M_n^3 with no loops or multiple edges.

Note that the number of faces in a triangulation has to be even; indeed, each edge is adjacent to two faces (in fact, two distinct faces in the case of a simple triangulation) and each face is adjacent to three edges. In particular, the number of edges of a triangulation with F faces is $E = \frac{3}{2}F$. Moreover, Euler's formula yields that the triangulation must have $V = \frac{1}{2}F + 2$ vertices.

It is therefore convenient to index simple triangulations by half their number of faces: we will denote the set of all simple triangulations of size $2n$ by STr_n; they have $2n$ faces, $n + 2$ vertices and $3n$ edges.

A convenient tree structure to define in conjunction to simple triangulation is blossoming trees, used in [16]:

Definition 3.2. A blossoming tree of size n (Figure 5, left) is a rooted plane tree t whose vertices are of two types, *blossoms* and non-blossoms, with the following properties:

1. every blossom has degree 1 (it is a leaf);
2. every non-blossom is adjacent to exactly two blossoms;
3. t has exactly n non-blossoms;
4. the origin of t is a blossom.

   We denote by $|t|$ the size of the blossoming tree t and by BT_n the set of all blossoming trees of size n. Note that any blossoming tree of size n has exactly $2n$ blossoms.

Poulalhon and Schaeffer provide in [16] an explicit bijection between the set STr_n and a special subset of BT_n containing *balanced blossoming trees*; for our purposes, a natural $2n$-to-1 map from $\mathsf{BT}_n \times \{\pm1\} \to \mathsf{STr}_n$, is actually enough. We briefly remind the reader of the construction of a triangulation in STr_n from a blossoming tree in BT_n and a spin $\varepsilon = \pm1$, highlighting those features that we will actually need.
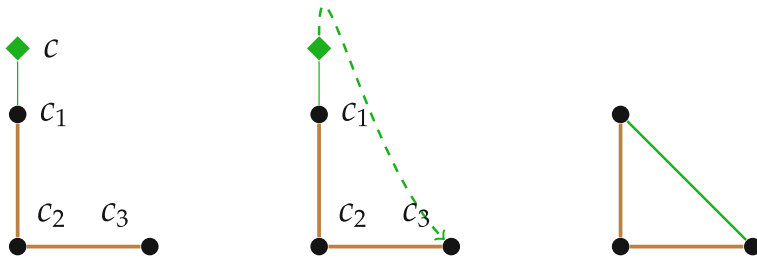
FIGURE 6 A closure operation performed within a blossoming tree. In the counterclockwise contour of the infinite face, after the corner c of a blossom, we encounter three corners in a row belonging to non-blossoms, namely, c_1, c_2, c_3. We bring the blossom to c_3, creating a new triangular face.

Construction 2 (Poulalhon, Schaeffer). Given $t \in \mathsf{BT}_n$ and $\varepsilon = \pm 1$, the construction of $\Xi(t, \varepsilon) \in \mathsf{STr}_n$ goes as follows:

1. First, perform a sequence of *closure* operations, each of which creates one new face. Consider the cyclic counterclockwise contour of the infinite face in the plane (which is initially the only face); if there is a corner c of a blossom that is immediately followed by corners c_1, c_2, c_3 adjacent to non-blossoms, we say that c is *closable*. One performs the closure of c by bringing its blossom to the vertex of c_3 as in Figure 6, thereby creating a new triangular face. After a closure is performed, re-compute the contour of the infinite face and keep performing closures until no more blossom corners are closeable. Note that the order in which closures are performed is not important.

2. When no more closures are possible, the resulting map has a number of triangular faces, and one infinite face that is not necessarily triangular. It can be shown that the boundary of the infinite face f_∞ is a cycle of non-blossoms, each of which is attached to one or two blossoms within f_∞ that have not been closed, and that there are exactly two non-blossoms that each have two blossoms attached to them (Figure 5, middle). As in the figure, label the blossoms in f_∞ as $L_1, \ldots, L_k, R_1, \ldots, R_h$, clockwise around f_∞, in such a way that L_1, R_h and L_k, R_1 are attached to the same non-blossom, and additionally that, if you consider the clockwise contour of the original blossoming tree started at the root corner, the corner of R_1 occurs before the corner of L_1.

3. Draw two new vertices L and R within f_∞. Identify $L_1, \ldots, L_k$ with L and $R_1, \ldots, R_h$ with R.

4. Draw an edge between L and R, which will be the root edge of the triangulation $\Xi(t, \varepsilon)$. If $\varepsilon = 1$, orient it from L to R; if $\varepsilon = -1$, orient it from R to L.

Note that the root blossom serves no purpose in the construction described above (other than determining how ε pairs with the orientation of the root edge), and is simply forgotten, which is why Ξ is a $2n$-to-one map (since $t \in \mathsf{BT}_n$ has $2n$ blossoms, each of which can potentially serve as the origin).

Proposition 3.1. *The map* $\Xi : \mathsf{BT}_n \times \{\pm 1\} \to \mathsf{STr}_n$ *of Construction 2 is $2n$-to-one. Non-blossom vertices of $t \in \mathsf{BT}_n$ naturally correspond to vertices of $\Xi(t, \varepsilon)$ that are not endpoints of the root edge. Edges between non-blossom vertices of t form a spanning tree of the map obtained from $\Xi(t, \varepsilon)$ by erasing the endpoints of its root edge (and any edges involving them). Blossoms of t correspond to edges of $\Xi(t)$ other than the root edge and the edges of the spanning tree just mentioned. The value of ε corresponds to the orientation of the root edge of $\Xi(t, \varepsilon)$.*
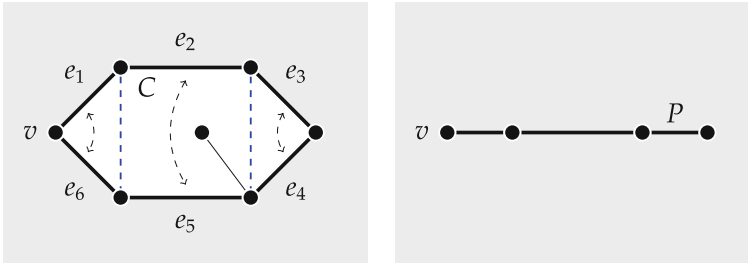
FIGURE 7 Collapsing a face with outer boundary C of length four in an eight-angulation M.

4 | GROWING 2P-ANGULATIONS

Our main concern in this article is constructing a uniform random map of size $n + 1$ in a certain class from a uniform random map of size n in the same class via a simple local (random) manipulation. What kind of manipulation we allow ourselves to make to "grow" our uniform maps is, to some extent, an arbitrary choice; in our case, we shall grow maps by adding faces, one at a time. We will presently give a more formal description of what we mean by adding a face: the fact that this is a very simple, natural, local operation will hopefully be apparent; as explained in the introduction, our motivation for considering this operation in particular comes from the potential applications to mixing time questions, which we shall explore in a subsequent paper.

Note that some previous work in the same direction does exist, but differs from ours in a few respects. Bettinelli [7] studies bijections that yield a procedure to grow uniform random quadrangulations; however, his allowed growth operations are more general than ours, as they involve "adding a face" but also performing certain types of surgery on the map, cutting it along a geodesic and re-gluing it in a different way. These surgeries make Bettinelli's growth operations non-local, which makes them ill suited to certain applications, such as the one of [9], where a different (but still explicit) growth scheme for quadrangulations was provided.

Given a 2p-angulation $M \in \mathsf{M}_{n+1}^{2p}$, let C be a simple cycle of M that encloses exactly one face and let v be a vertex along the cycle C. Since all 2p-angulations are bipartite, the cycle C has even length; since it encloses exactly one face, its length is $2k \leq 2p$. Thus, we can express C as a sequence of edges $e_1, \ldots, e_{2k}$ such that v is the common endpoint of e_1 and e_{2k}. Define the 2p-angulation $\mathrm{coll}(M, C, v) \in \mathsf{M}_n^{2p}$ as the one obtained by erasing all vertices and edges in the interior of C and pairwise identifying the edges of C: e_1 with e_{2k}, e_2 with e_{2k-1}, and so on, in such a way as to turn the cycle into a path of length P started at v (Figure 7). We say that $\mathrm{coll}(M, C, v)$ is obtained form M by *collapsing a face* (in this expression, we are referring to the face enclosed by C). Vice-versa, we say that M can be obtained by $\mathrm{coll}(M, C, v)$ by *growing a face*; more specifically, by *growing a face at P*.

Note that not all faces can be collapsed and that, given a face that can be, it can be collapsed in k ways, where $2k$ is the length of its outer boundary, depending on the choice of the vertex v (choosing the vertex opposite v in the cycle C yields the same result). Given a 2p-angulation $m \in \mathsf{M}_n^{2p}$ and a simple path P, there is a unique way to grow a face at P if the length of the path is p and it does not contain the root edge. If P contains the root edge, we might grow a face at P to the left or to the right of the root edge; additionally, if the length of the path is less than k, we get a choice of how to arrange edges and vertices in the interior of the newly grown face.

The notation introduced here is slightly different than the one used in [9] in the case of quadrangulations, where we allowed the collapse of faces whose outer boundary has other faces in its interior (see fig. 10 of [9]). Such operations were not actually used in the growth scheme we

constructed, and they are markedly easier to deal with in the case $p = 2$, so we shall avoid them altogether.

What we wish to do is show that it is possible to grow uniform $2p$-angulations by the successive growing of faces at random locations P. To this end, we define the notion of a growth scheme:

Definition 4.1. A growth scheme for $2p$-angulations is a collection of functions f_n : $\mathsf{M}_{n+1}^{2p} \times \mathsf{M}_n^{2p} \to [0, 1]$, for $n \geq 1$, such that

(i) if $f_n(M, m) > 0$, then $m = \text{coll}(M, C, v)$ for some C, v.
(ii) for all $M \in \mathsf{M}_{n+1}^{2p}$, $\sum_{m \in \mathsf{M}_n^{2p}} f_n(M, m) = 1$;
(iii) for all $m \in \mathsf{M}_n^{2p}$, $\sum_{M \in \mathsf{M}_{n+1}^{2p}} f_n(M, m) = \frac{|\mathsf{M}_{n+1}^{2p}|}{|\mathsf{M}_n^{2p}|}$.

Note that a growth scheme allows us to couple a uniform $2p$-angulation with n faces U_n with a uniform $2p$-angulation with $n + 1$ faces U_{n+1} in such a way that $U_n = \text{coll}(U_{n+1}, C, v)$, where C is a random simple cycle of U_{n+1} enclosing a single face and v a random vertex of U_{n+1} (whose distributions conditional on U_{n+1} can be determined from the function f_n). We can define a random variable X taking values in M_{n+1}^{2p} on the same probability space as U_n, as follows. Conditionally on $U_n = m$, set $X = M$ with probability proportional to $f_n(M, m)$. Property (i) ensures that U_n is of the form $\text{coll}(X, C, v)$, and properties (ii) and (iii) that the distribution of X is the same as the distribution of U_{n+1}.

In [9], a growth scheme was provided for quadrangulations, that is, for $p = 2$; note that the probabilities $f_n(\cdot, \cdot)$ were computed explicitly in terms of the structure of the tree underlying the quadrangulation; this was thanks to the fact that growth schemes for plane trees are easy to determine explicitly. In this paper, we will show the existence of a growth scheme for $2p$-angulations, but not give explicit expressions for the functions $f_n(\cdot, \cdot)$.

The rest of this section will be devoted to proving Theorem 1.

This can be done by using the fact, proven by Luczak and Winkler in [14], that there exists a growth scheme for d-ary trees, by which we mean a sequence of functions $g_n : \mathsf{T}_{n+1}^d \times \mathsf{T}_n^d \to [0, 1]$ satisfying properties (ii) and (iii) from Definition 4.1, where occurrences of M_k^{2p} are replaced by T_k^d, and the property (i') that, if $g_n(T, t) > 0$, then $T = \text{grow}(t, v)$ for some leaf v of t. For ease of reference, we include here the statement of Theorem 4.1 of [14]:

Theorem 3. (Luczak, Winkler). *For each $d \geq 1$, there exists a growth scheme $(g_n)_{n \geq 1}$ for complete d-ary trees.*

The ways in which our notation differs from the one of [14] are very minor: our explicit definition of a growth scheme, which we have adopted in order to maintain the same notation as our previous works [9,10], is equivalent to their notion of a *building scheme*; as per Remark 2.1, our use of the complete d-ary trees of Definition 2.2 and their "growth" operation does entirely correspond to the subtrees of $\mathbb{T}^d$ and their containment relation as defined in [14].

A first step in order to obtain Theorem 1 from Theorem 3 is to recast Theorem 3 in terms of p-mobiles. Thanks to Remark 2.2 and the bijection Φ, this is rather immediate:

Corollary 4.1. Given a growth scheme $(g_n)_{n \geq 1}$ for p-ary trees, the functions $(\tau, \tau') \mapsto g_n(\Phi(\tau), \Phi(\tau'))$ yield a growth scheme for unlabeled p-mobiles.

Proof. This follows immediately from the fact that Φ is a bijection and that τ is obtained by growing τ' at a corner c if and only if $\Phi(\tau)$ is obtained by growing $\Phi(\tau')$ at a leaf. ∎

What remains to do is to establish a link between growing p-mobiles and growing $2p$-angulations via the BDFG construction:

Lemma 4.2. *Let τ be a labeled p-mobile in LMob_n^p, let c be a corner of τ adjacent to a white vertex and let $s \in S_p$; let $\varepsilon = \pm 1$. There exists a simple cycle K and a vertex w of the pointed 2p-angulation $\Psi(grow(\tau, c, s), \varepsilon)$ such that $coll(\Psi(grow(\tau, c, s), \varepsilon), K, w) = \Psi(\tau, \varepsilon)$, where Ψ is the bijection from Construction 1.*

Proof. Let v be the new black vertex in $grow(\tau, c, s)$ and $v_1, \dots, v_p$ be its neighbours, ordered clockwise starting with the parent of v.

In $\Psi(grow(\tau, c, s), \varepsilon)$, a map which we call m, the black vertex v corresponds to a face f_v whose external boundary is a simple cycle C.

We shall show that C contains the vertex v_1, that it has only the face f_v in its interior, and that the statement of the theorem is satisfied if we take $K = C$, $w = v_1$.

The cycle C is obtained as follows in the context of the construction Ψ: let a be the label of v_1 after the first step of Construction 1; start with the corner c_l of v_1 immediately to the left of the edge (v_1, v), and follow the map edge drawn towards the corner of the first vertex w_1 you meet in the contour that is labeled $a - 1$; if $w_1 \notin \{v_2, \dots, v_p\}$, stop; otherwise, follow the map edge drawn from that corner to the next corner labeled $a - 2$, and so on. This procedure yields a simple path P_l of map edges which starts at v_1 and ends at a vertex $w \notin \{v_1, \dots, v_p\}$, labeled $a - k$ for some $k > 0$. Consider then the simple path P_r obtained by following successive map edges drawn from the corner c_r of v immediately to the right of the edge (v_1, v) until they hit w. The fact that w will be hit is clear, as the corner of w hit by the first path is the first one labeled $a - k$ coming after c_r.

The cycle C obtained by concatenating the two simple paths P_l and P_r does not enclose any black vertices other than v because $v_2, \dots, v_p$ are leaves of the mobile, so it is the external boundary of the face f_v. The 2p-angulation $coll(m, C, v_1)$ is obtained by collapsing the face f_v, gluing together the simple paths P_l and P_r starting from their common origin v_1.

The same 2p-angulation is obtained as $\Psi(\tau, \varepsilon)$, where the construction of path P_r is unaffected and map edges that were hitting corners of P_l in the previous construction now hit the corners with the same label hit by P_r (Figure 8). ∎

Remark 4.1. Note that, given 2p-angulations m, M such that M is obtained from m by growing a face, it is not necessarily true that $\Psi^{-1}(M)$ can be obtained by growing $\Psi^{-1}(m)$, in the sense of grafting a p-mobile of size 1 somewhere. Since our construction of growth schemes for 2p-angulations will rely on growth schemes for mobiles, what this entails is that there are some allowed growth operations on 2p-angulations that will never be used.

Finally, we have

Proposition 4.3. *Suppose $(g_n)_{n \geq 1}$ is a growth scheme for unlabeled p-mobiles. Define*

$$f_n(m, m') = \frac{1}{|V(m)| \binom{2p-1}{p}^n} \sum_{v \in V(m)} \sum_{v' \in V(m')} 1_{\varepsilon = \varepsilon'} g_n(\tilde{\tau}, \tilde{\tau}'),$$

where $(m, v) = \Psi(\tau, \varepsilon)$ and $(m', v') = \Psi(\tau', \varepsilon')$, and $\tilde{\tau}, \tilde{\tau}'$ are obtained by forgetting the labels on the black vertices of τ, τ'.

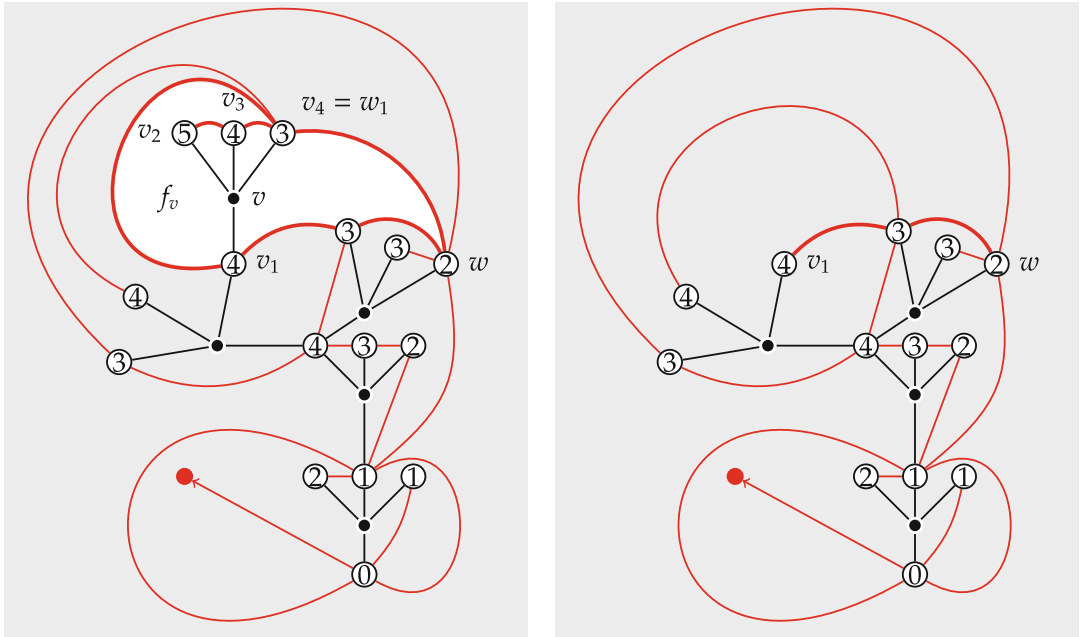*The sequence of functions $(f_n)_{n \geq 1}$ is a growth scheme for 2p-angulations.*

FIGURE 8 On the left, an eight-angulation $m = \Psi(\text{grow}(\tau, c, s), \varepsilon)$ for some $(\tau, \varepsilon) \in \mathsf{LMob}_4^4 \times \{\pm 1\}$, a corner c of τ and $s = + + - - - + +$. On the right, the eight-angulation $\Psi(\tau, \varepsilon)$, which is obtained from m by collapsing the face f_v, whose outer boundary becomes the thick red path of length two from v_1 to w in the picture.

Proof. The fact that the functions $(f_n)_{n \geq 1}$ satisfy condition (i) of Definition 4.1 is due to Lemma 4.2. Indeed, if $f_n(m, m') > 0$, then there are $v \in m$ and $v' \in m'$ such that the mobile τ corresponding to (m, v) is obtained by growing the mobile τ' corresponding to (m', v'). By Lemma 4.2, it follows that m' differs from m by collapsing a face.

As for property (ii), given $m \in \mathsf{M}_{n+1}^p$, set $C(n) = \frac{1}{|V(m)|\binom{2p-1}{p}^n}$ and consider the sum

$$\sum_{m' \in \mathsf{M}_n^{2p}} f_n(m, m') = \sum_{m' \in \mathsf{M}_n^{2p}} C(n) \sum_{v \in V(m)} \sum_{v' \in V(m')} 1_{\varepsilon = \varepsilon'} g_n(\tilde{\tau}, \tilde{\tau}') = C(n) \sum_{v \in V(m)} \sum_{(m', v') \in \mathsf{M}_n^{2p, \bullet}} 1_{\varepsilon = \varepsilon'} g_n(\tilde{\tau}, \tilde{\tau}');$$

using the BDFG correspondence, we can rewrite the internal sum to obtain

$$C(n) \sum_{v \in V(m)} \sum_{(\tau', \varepsilon') \in \mathsf{LMob}_n^p \times \{\pm 1\}} 1_{\varepsilon = \varepsilon'} g_n(\tilde{\tau}, \tilde{\tau}') = C(n) \sum_{v \in V(m)} \sum_{\tau' \in \mathsf{LMob}_n^p} g_n(\tilde{\tau}, \tilde{\tau}') =$$

$$= C(n) \sum_{v \in V(m)} \sum_{\tilde{\tau}' \in \mathsf{Mob}_n^p} \binom{2p-1}{p}^n g_n(\tilde{\tau}, \tilde{\tau}') = C(n) \binom{2p-1}{p}^n \sum_{v \in V(m)} 1 = 1.$$

Similarly, for (iii) one has

$$\sum_{m \in \mathsf{M}_{n+1}^{2p}} f_n(m, m') = C(n) \sum_{v' \in V(m')} \sum_{\tau \in \mathsf{LMob}_{n+1}^p} g_n(\tilde{\tau}, \tilde{\tau}') =$$

$$= C(n) \binom{2p-1}{p}^{n+1} |V(m')| \frac{|\mathsf{Mob}_{n+1}^p|}{|\mathsf{Mob}_n^p|} = \frac{|\mathsf{LMob}_{n+1}^p||V(m')|}{|\mathsf{LMob}_n^p||V(m)|} = \frac{|\mathsf{M}_{n+1}^{2p}|}{|\mathsf{M}_n^{2p}|}. \qquad \blacksquare$$

5 │ GROWING SIMPLE TRIANGULATIONS

In this section, we show the existence of a growth scheme for simple triangulations; we intend to use a similar growth operation to the one introduced for $2p$-angulations in the previous section, but note that some small tweak is necessary: it is of course not possible to collapse a single face of a simple triangulation to obtain a smaller simple triangulation. The operations we shall allow in this case consist in growing (or collapsing) *two faces* at a time.

Given a triangulation $t \in \mathsf{STr}_n$ and an edge e of t that is not the root edge, we let $\mathrm{coll}(t, e)$ be the triangulation obtained as follows: the edge e has endpoints v, w and is adjacent to two triangular faces; we collapse these two faces by eliminating e, identifying v with w, and "gluing together" the remaining pairs of edges that now form two cycles of length two (Figure 9). If the root edge is on the boundary of one of the two collapsed faces, we simply maintain its status and orientation in the new triangulation.

We say that $\mathrm{coll}(t, e)$ is obtained from t by *collapsing a pair of (adjacent) faces* and that t is obtained from $\mathrm{coll}(t, e)$ by *growing a pair of (adjacent) faces*.

In complete analogy with Definition 4.1, we give

> **Definition 5.1.** A growth scheme for simple triangulations is a collection of functions $f_n : \mathsf{STr}_{n+1} \times \mathsf{STr}_n \to [0, 1]$, for $n \geq 1$, such that
>
> (i) if $f_n(T, t) > 0$, then $t = \mathrm{coll}(T, e)$ for some edge e of T;
> (ii) for all $T \in \mathsf{STr}_{n+1}$, $\sum_{t \in \mathsf{STr}_n} f_n(T, t) = 1$;
> (iii) for all $t \in \mathsf{STr}_n$, $\sum_{T \in \mathsf{STr}_{n+1}} f_n(T, t) = \frac{|\mathsf{STr}_{n+1}|}{|\mathsf{STr}_n|}$.

In order to show the existence of a growth scheme for simple triangulations, we intend to relate the operation of growing/collapsing pairs of faces to the growth operations defined on complete d-ary trees in Section 2, and then adapt a growth scheme for d-ary trees obtained by Theorem 3.

In order to do this, we first need to bridge the gap between triangulations and d-ary trees, which we do via the Poulalhon-Schaeffer construction we recalled in Section 3. But first, we express a general blossoming tree τ as a pair (L_τ, R_τ) of complete four-ary trees, which will enable us to introduce a new bijection relating such pairs to simple triangulations.

> **Construction 3** (Of a pair of complete four-ary trees from a blossoming tree). Given a blossoming tree $t \in \mathsf{BT}_n$, consider the two (multi-type) plane trees $L(t), R(t)$ obtained as follows. The origin ρ of t has one child, which is a non-blossom vertex v; v has two blossom neighbours, ρ and ρ'. Build $L(t)$ by erasing ρ, ρ' and all descendants of v that come after ρ' in the clockwise contour of t; root the ensuing tree in the corner that used to contain ρ.
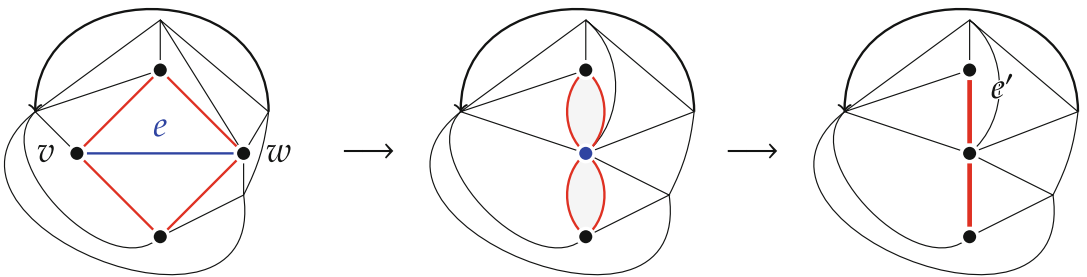

FIGURE 9 On the left, a simple triangulation $T \in \mathsf{STr}_{12}$; collapsing the two faces adjacent to the edge e identifies vertices v and w an yields the triangulation $\mathrm{coll}(T, e) \in \mathsf{STr}_{10}$ depicted on the right. Note that collapsing a pair of faces does not necessarily yield a simple triangulation: for example, the triangulation $\mathrm{coll}(\mathrm{coll}(T, e), e')$ is no longer simple.

Similarly, $R(t)$ is built by erasing ρ, ρ' and all descendants of v that come before ρ' in the clockwise contour of t, and rooting the ensuing tree in the corner that used to contain ρ. The objects $L(t)$ and $R(t)$ (see Figure 10) are plane trees whose vertices are either blossoms or non-blossoms; their origin is a non-blossom and has no blossom neighbors; all other non-blossoms have exactly two blossom neighbors and all blossoms are leaves. We shall call the set of such plane trees having n non-blossoms T_n. Naturally, given $L \in T_a$ and $R \in T_b$, by gluing the two trees together and separating them with two blossoms (one of which becomes the root), one can recover the corresponding blossoming tree in BT_{a+b-1}. The mapping $t \mapsto (L(t), R(t))$ is a natural bijection between BT_n and $\cup_{k=1}^n T_k \times T_{n+1-k}$.

We now show how $L(t) \in T_a$ can be interpreted as an element of T_{a-1}^4 (similarly for $R(t)$). Indeed, the sets $(T_k)_{k\geq 1}$ and $(\mathsf{T}_k^4)_{k\geq 0}$ essentially satisfy the same recursive structure; we will therefore build the desired correspondence inductively.

Suppose $a = 1$, in which case $L(t)$ is the tree with a single non-blossom vertex and no blossoms, which is the only element of T_1; this corresponds to the "degenerate" complete four-ary tree with only an origin, which is the only element of T_0^4. Now assume that we have a way to interpret elements of T_i, for $1 \leq i < k$, as elements of T_{i-1}^4, and that $L(t)$ is an element of T_k, where $k > 1$. We can decompose $L(t)$ into four multi-type plane trees t_1, t_2, t_3, t_4, each of them in T_i for some $i < k$, as follows.

Since $k > 1$, $L(t)$ has an origin w and a rightmost child of the origin z (both of them non-blossoms, since w has no blossom neighbours). Let t_1 be the multi-type tree obtained by erasing z and all its descendants, which is in T_{k_1} for some $k_1 < k$. Now consider the multi-type plane tree t' obtained by erasing t_1 from $L(t)$ and rooting the ensuing tree in the corner that used to contain the edge from w to z. The tree t' is not quite a blossoming tree, as it's not rooted in a blossom; however, it can be split similarly to how we split t at the very beginning, by considering the multi-type tree to the left of the leftmost blossom child of z (rooted in z), the one between the two blossom children of z, and the one to the right of the right blossom child of z. These three trees are t_2, t_3, t_4, and they belong to $T_{k_2}, T_{k_3}, T_{k_4}$ for some $k_2, k_3, k_4 < k$ such that $k_1 + k_2 + k_3 + k_4 = k + 2$.

By assumption, t_i corresponds to a complete four-ary tree $\tilde{t}_i$ in $\mathsf{T}_{k_i-1}^4$; we will simply have $L(t)$ correspond to the complete four-ary tree obtained by taking the four-ary tree $\tilde{t}_0 \in \mathsf{T}_1^4$ with leaves v_1, v_2, v_3, v_4 (in order from left to right) and grafting onto the leaf v_i the four-ary tree $\tilde{t}_i$, which yields a four-ary tree in $\mathsf{T}_{1+k_1+k_2+k_3+k_4-4}^4 = \mathsf{T}_{1+k+2-4}^4 = \mathsf{T}_{k-1}^4$.

Proof. (Proof of Proposition 1.1) The proposition now comes as an immediate consequence of the construction above: by joining the four-ary trees l and r with an edge between their origins, oriented from l to r, we obtain a tree all of whose vertices have degree 5 or 1, endowed with a distinguished oriented edge. This mapping is a bijection between $\cup_{a=0}^{n-1} \mathsf{T}_a^4 \times \mathsf{T}_{n-1-a}^4$ and $\mathcal{T}_{n-1}$. Construction 3 is in turn a bijection between the set BT_n and the set $\cup_{a=0}^{n-1} \mathsf{T}_a^4 \times \mathsf{T}_{n-1-a}^4$. Construction 2 gives a 2n-to-1 map from $\mathsf{BT}_n \times \{\pm 1\}$ to STr_n, hence a 2n-to-1 map from $\mathcal{T}_{n-1} \times \{\pm 1\}$ to STr_n, proving Proposition 1.1. ∎

Remark 5.1. The map of Proposition 1.1, which essentially induces a bijection between simple triangulation and a certain set of trees with vertices of degree 5 or 1, will prove very useful in our context, but it is not without some drawbacks. On the one hand, we no longer have to distinguish between blossoms and non-blossoms, and tasks related to enumeration tend to become more straightforward. Some important properties are preserved, and we
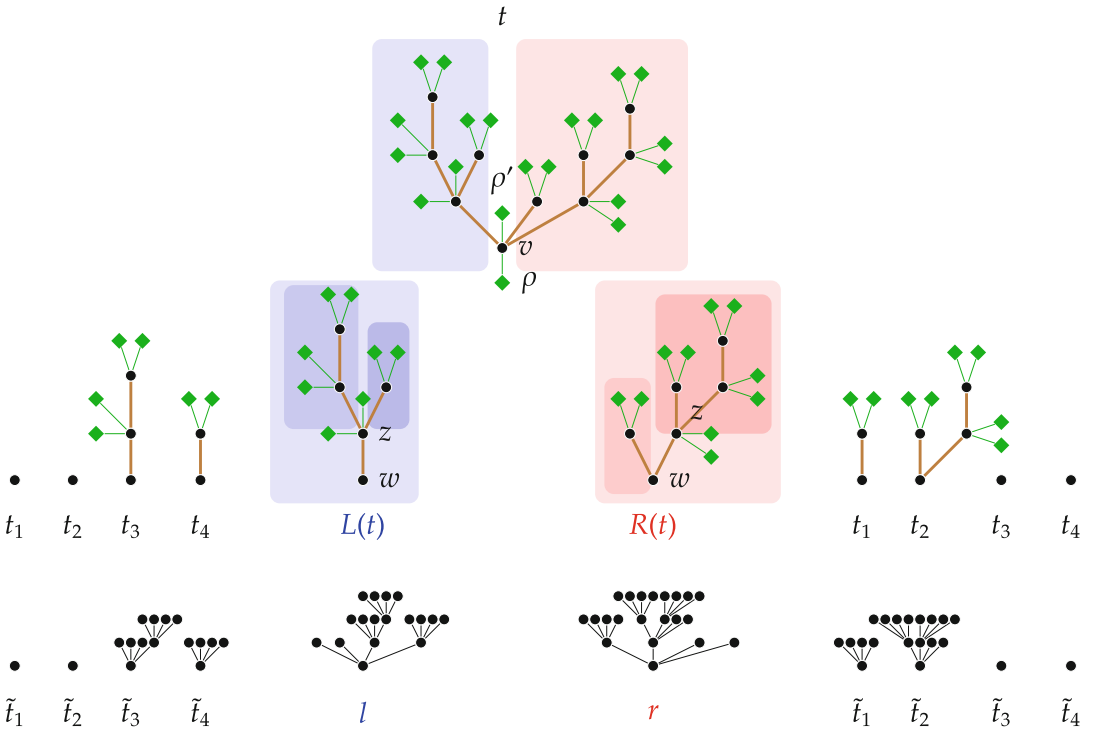
FIGURE 10 Construction 3 performed on $t \in \mathsf{BT}_{10}$. Below t, we see its left part $L(t) \in T_5$ and its right part $R(t) \in T_6$. Below these, the corresponding four-ary trees $l \in \mathsf{T}_4^4$ and $r \in \mathsf{T}_5^4$, constructed by grafting the four trees $\tilde{t}_1, \tilde{t}_2, \tilde{t}_3, \tilde{t}_4$ onto the unique element of T_0^4.

will see later that growing the tree at a leaf does correspond to growing a pair of faces in the triangulation.

On the other hand, metric information about the triangulation, and even just the blossoming tree, becomes very "distorted" when seen through the lens of this four-ary tree. Note that, because of the special role of the first child of every vertex in the construction, exploring "on the right" corresponds to moving faster in the blossoming tree than exploring "on the left". Moreover, the construction is non-canonical in many respects, as our choice of having the leftmost child represent the subtree to the left of the rightmost branch is one of several possible.

Lemma 5.1. *Construction 3 gives a bijection between* BT_n *and* $\cup_{a=0}^{n-1} \mathsf{T}_a^4 \times \mathsf{T}_{n-1-a}^4$. *Moreover, given* $t \in \mathsf{BT}_n$ *and its pair of corresponding four-ary trees* (l, r), *we can naturally identify the set of internal vertices of l and r with the set of non-blossom vertices of t other than the one connected to the origin blossom. Finally, given a leaf v of l or r, let t^v be the tree in* BT_{n+1} *that corresponds to the pair* $(grow(l, v), r)$ *or* $(l, grow(r, v))$, *respectively. There is a corner of a non-blossom vertex w of t such that t^v is obtained by grafting onto that corner a single edge from w to a non-blossom w', in turn connected to two blossom children.*

Proof. The fact that the given construction is a bijection immediately follows from the fact that, given any pair $(l, r) \in \cup_{a=0}^{n-1} \mathsf{T}_a^4 \times \mathsf{T}_{n-1-a}^4$, we can easily obtain a unique tree $t \in \mathsf{BT}_n$ by reversing it. Its inductive structure also makes it clear that internal vertices of l and r

correspond to non-blossoms of $L(t)$ and $R(t)$, with the exception of their identified origin which is the non-blossom neighbor of the root of t.

Let us now consider what the multi-type tree $\text{grow}(l, v)$ corresponds to. Consider the parent $p(v)$ of v in l; this corresponds to a non-blossom vertex w of $L(t)$. If v is the left-most child of $p(v)$, the fact that it is a leaf implies that w has only one non-blossom child. Changing l into $\text{grow}(l, v)$ corresponds to adding a leftmost non-blossom child to w, with no left non-blossom siblings or children (and therefore only two blossom children). If v is the second, third or fourth child of $p(v)$, the fact that it is a leaf tells us that, letting w' be the rightmost non-blossom child of w, w' has no non-blossom children to the left of its left blossom child, or no non-blossom children between its two blossom children, or no non-blossom children to the right of the right blossom child. Replacing l by $\text{grow}(l, v)$ corresponds to adding a non-blossom child to w' in the appropriate position with respect to the blossoms; the child added must have two blossom children and no non-blossom neighbours other than w'.

The construction for $v \in r$ is completely analogous. ∎

Remark 5.2. Note that we again have a situation along the lines of Remark 4.1, namely, not all ways of growing a new non-blossom vertex with two blossom children actually correspond to a growth operation on the side of four-ary trees. Our growth scheme, being based on four-ary tree growth, will only use a set of "special" growth moves, and in particular (e.g.) it will never grow a non-blossom at a corner where both its immediate left sibling and its immediate right sibling would exist and be non-blossoms.

At this point, given a simple triangulation $t \in \text{STr}_n$ and a sign $\varepsilon = \pm 1$, we have $2n$ blossoming trees $\tau \in \text{BT}_n$ such that $\Xi(\tau, \varepsilon) = t$; each of these can be expressed as a pair consisting of its left and right part, which we can interpret as a pair of four-ary trees with a total of $n - 1$ internal vertices between them. Note that, indeed, we have the following identity involving the number of simple triangulations and the number of four-ary trees:

$$\frac{2(4n-3)!}{n!(3n-1)!} = |\text{STr}_n| = \frac{2}{2n} \sum_{k=0}^{n-1} |\text{T}_k^4||\text{T}_{n-1-k}^4| = \frac{1}{n} \sum_{k=0}^{n-1} \frac{1}{(3k+1)(3n-3k-2)} \binom{4k}{k} \binom{4(n-1-k)}{n-1-k}.$$
(2)

What we now wish to show is that the operation of growing one of these two four-ary trees at a leaf yields a new blossoming tree $\tau' \in \text{BT}_{n+1}$ such that $\Xi(\tau', \varepsilon)$ is obtained from t by growing a pair of adjacent faces. More specifically,

Lemma 5.2. *Let τ be a blossoming tree in BT_n and let $(l, r) \in \cup_{a=0}^{n-1} \text{T}_a^4 + \text{T}_{n-1-a}^4$ be obtained from τ via Construction 3; let v be a leaf of l or r and let τ' be the blossoming tree in BT_{n+1} such that Construction 3 on τ' yields $(\text{grow}(l, v), r)$ or $(l, \text{grow}(r, v))$ (depending on which tree v is a leaf of). Given $\varepsilon = \pm 1$, there exists an edge e of $T = \Xi(\tau', \varepsilon)$ such that $\Xi(\tau, \varepsilon) = \text{coll}(T, e)$.*

Proof. By Lemma 5.1, the blossoming tree τ' is obtained from τ by grafting a non-blossom w with two blossom children onto some corner c. Let α, β be the left and right blossom child of w, respectively, and let e be the edge of the triangulation $T = \Xi(\tau', \varepsilon)$ that corresponds to β according to Construction 2. We intend to show that $\Xi(\tau, \varepsilon) = \text{coll}(T, e)$.

Let us consider the following separate cases:

a w is not a neighbour of an endpoint of the root edge in T;

b w is a neighbour of a single endpoint of the root edge in T;

c w is a neighbour of both endpoints of the root edge in T.

Suppose we are in case (a). In the triangulation T, let f_1 and f_2 the two triangles adjacent to the edge e; we may assume that the edges of f_1 are $(w, p(w))$, by which we mean the edge of τ' joining w to its parent $p(w)$, the edge e and some edge e_1, and that the edges of f_2 are the edge created by closing the blossom α, e and some edge e_2. Let z be the common vertex of e_1 and e_2.

We know that we can interpret τ' as a spanning tree of T, excluding the endpoints of the root edge. As in Figure 11, we can therefore consider the set of faces of T that lie between e_1 and the shortest path in τ' between $p(w)$ and z (call this set S_1), as well as the set of faces between e_2 and the shortest path in τ' joining the two endpoints of e_2 (call this set S_2). Note that if e_i belongs to τ' then S_i is empty.

Assuming S_1 is non-empty, it is a triangulation with a boundary, all of whose vertices actually lie on the boundary; the "internal" edges are obtained by closing blossoms and the "boundary" edges are all originally in τ', save for e_1. By starting with triangular faces that have two boundary edges (i.e., so-called "ears" of the triangulation S_1) and iterating, one can perform successive closures on τ' in order to obtain all faces in S_1 (and no other faces of T). Similarly, one can perform closures of blossoms to create all faces of S_2. After e_1 and e_2 have been created in this way if they were not originally present in τ', one can close β and then α. This sequence of closures yields a partial closure C' of τ' whose faces are the triangles in $S_1 \cup S_2 \cup \{f_1, f_2\}$, plus the infinite face (see the upper central part of Figure 11). It is possible to perform the exact same closures (other than the last two) in the same order in τ, yielding a partial closure C that can be obtained from C' by collapsing f_1 and f_2, identifying w with z. Note that unclosed blossoms in C and C' correspond to each other and that the same successive closures can now be performed, so that completing Construction 2 yields exactly the same triangulation but for f_1 and f_2, or in other words $\Xi(\tau, \varepsilon) = \text{coll}(T, e)$.

Now consider case (b). We can again define z, f_1, e_1, e_2 as before; note that z is *not* an endpoint of the root edge of T, as if only one blossom among α, β becomes closable before creating the root edge endpoints, this must be the right blossom β. We can therefore define the set of faces S_1 as before, since $p(w)$ and z are both vertices of τ'. Again, close the necessary blossoms of τ' to obtain the faces of S_1, and then close β, thus yielding a partial closure C' of τ'. Perform the exact same closures on τ to obtain a map C that differs from C' by the absence of the face f_1 (see Figure 12).

Consider now the following set S_3 of faces of T. Order the edges incident to w clockwise starting with the edge obtained by closing α and let η be the *last* edge according to this ordering, which is either $(w, p(w))$ or is obtained by closing a blossom. Let u be the endpoint of η that is not w, which cannot be an endpoint of the root edge. Let the faces in S_3 be those that lie within the cycle formed by η and the shortest path in τ' leading from u to w.

As before, a sequence of closures can be performed on C' to create precisely the faces in S_3, thus yielding a map C_1'. Such closures can be performed on C in the same order; what they yield is a map C_1 that can be constructed from C_1' by collapsing the face f_1, identifying w with z and the blossom α with a blossom γ issued from z, which must exist because otherwise α would be closable in C'. It follows that subsequent closures can be
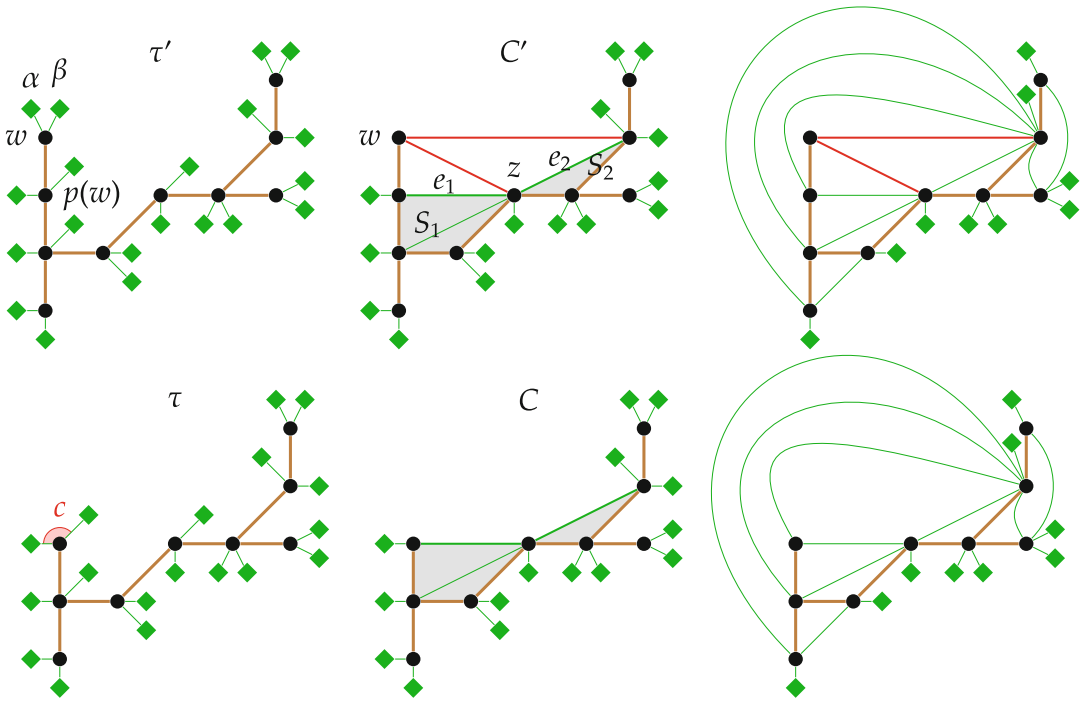
FIGURE 11   Case (a): the vertex w is grafted onto the corner c of τ to obtain τ'. One can perform three closures of blossoms to obtain the three triangular faces of $S_1 \cup S_2$ (shaded); when performed on τ, such closures yield C (below); when performed on τ' and followed by the closures of α an β, they yield C' (above). Since C differs from C' by the collapse of the two faces adjacent to e (and the identification of w with z), subsequent closures can be performed in the same way above and below, yielding the result on the right, where the contour of the infinite face is exactly the same.

performed in C_1 and C_1' at the same time, preserving the property that intermediate maps are obtained one from the other by collapsing f_1 and eliminating α. The final steps of the construction are also the same, with an additional face involving e being formed in T (see the rightmost part of Figure 12).

Case (c) is quite similar. We can define u, η, S_3 exactly as above; letting S be the set of faces of T containing w, set $S_4 = S_3 \setminus S$. There is a sequence of closures of blossoms in τ' that builds exactly the faces in S_4, and such a sequence can be mirrored on τ to yield a map which only differs by the absence of w and its two blossoms α, β. Note that the fact that w is adjacent to both root edge endpoints implies the presence of a blossom γ immediately to the right of the edge $(w, p(w))$, which does not become closable until the very last phase of the construction. After the sequence of closures performed on τ, there is a series of unclosed blossoms $b_1, \ldots, b_k$ (lower central part of Figure 13) adjacent to the infinite face, starting with two blossoms attached to u and ending with γ, which are not closable and will not become so until the root edge endpoints have been created. Any further closure performed on τ does not involve these blossoms and can be replicated on the partial closure of τ'. After performing all possible closures on τ (and mirroring them on τ') we are faced with two partial closures C and C' such that

1. C is a triangulation with a (simple) boundary with one or two blossoms hanging off every boundary vertex; the vertex u is adjacent to two unclosed blossoms;
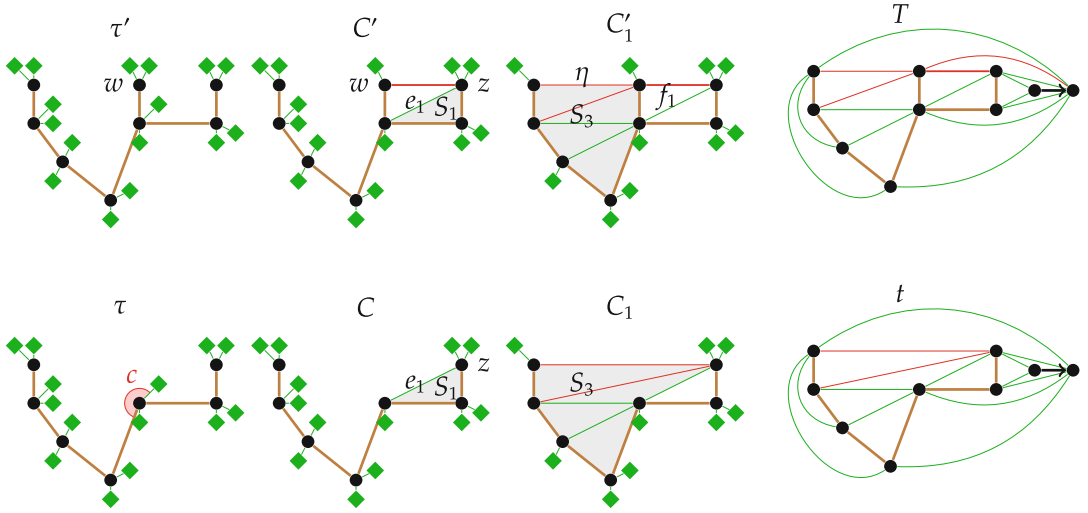
FIGURE 12   Case (b): the blossoming tree τ' is obtained by grafting w onto the corner c, and the vertex w is adjacent to a single endpoint of the root edge in $T = \Xi(\tau', \varepsilon)$. The intermediate maps C' and C'_1 are obtained by first closing the blossom that creates e_1, then the right blossom β of w, and then four blossoms, two of which become edges with w as an endpoint. The corresponding closure operations performed on τ (when possible) yield the intermediate maps C and C_1. Finally, we obtain T and $t = \text{coll}(T, e)$.
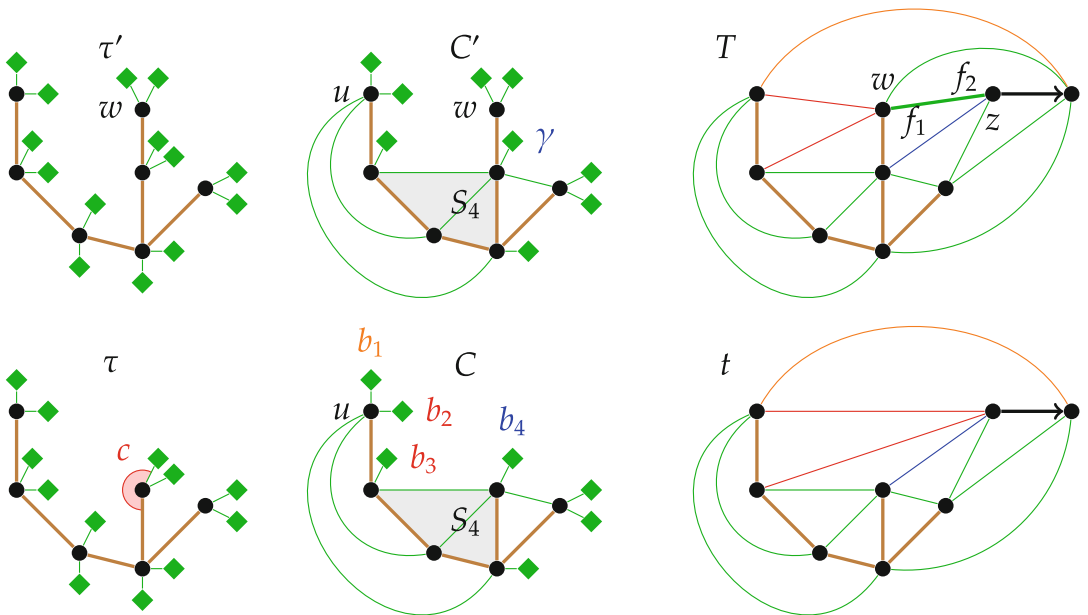
FIGURE 13   Case (c): τ' is obtained from τ by grafting w and with its two blossom children onto the corner c; in the triangulation $T = \Xi(\tau', \varepsilon)$, the vertex w is a neighbor of both endpoints of the root edge. The intermediate map C is built by first performing the two closures that create the faces of S_4 and then performing the other three possible closures. Above, the map C' is the one obtained by closing the corresponding blossoms of τ' in the same order. Note how the blossoms $b_1, \dots, b_4$ of C end up joined to the endpoints of the root edge in t; above, the corresponding blossoms b_2, b_3 are joined to w, while $\gamma = b_4$ and b_1 are joined to the endpoints of the root edge. The blossoms α, β and the edge $(w, p(w))$ create the faces f_1, f_2 of T.

2. C' is exactly the map C, with an added edge $(w, p(w))$ hanging off the boundary and blossoms α, β attached to w.

The triangulation T is obtained from C' by first closing $b_2, \ldots, b_{k-1}$, thus creating the faces in $S_3 \cap S$ and then adding the endpoints of the root edge and performing the final closures. The triangulation $t = \Xi(\tau, \varepsilon)$ is obtained by drawing the endpoints of the root edge and joining them to unclosed blossoms, including $u_1, \ldots, u_k$. The difference between the two is given purely by the presence of the two faces f_1, f_2 adjacent to e in T, and by the fact that the blossoms $u_2, \ldots, u_k$ create edges to w rather than the appropriate root edge endpoint. Collapsing f_1, f_2 in T by identifying the two endpoints of e yields precisely the triangulation t obtained by performing the last construction phase on the partial closure C. ∎

What we now wish to do is obtain the existence of a growth scheme for triangulations from the existence of a growth scheme for four-ary trees, which is proven in [14]. We will do this in two steps.

Lemma 5.3. *There exists a function* $h : \mathbb{N}^2 \to [0, 1]$ *such that for all* $a, b \in \mathbb{N}^2$

1. *if* $a + b \geq 1$, *then* $h(a, b) + h(b, a) = 1$;
2. $h(a + 1, b) \frac{|T_{a+1}^4|}{|T_a^4|} + h(b + 1, a) \frac{|T_{b+1}^4|}{|T_b^4|} = \frac{|STr_{a+b+2}|}{|STr_{a+b+1}|} \frac{a+b+2}{a+b+1}$.

Proof. In order to prove the statement, we rephrase it in terms of flow networks. For each $n \in \mathbb{N}$, consider the bipartite directed graph whose vertices are indexed by pairs $(a, b) \in \mathbb{N}^2$ such that $a + b = n$ or $a + b = n + 1$, with an edge from each pair (a, b) such that $a + b = n$ to the two pairs $(a + 1, b)$ and $(a, b + 1)$. Add a source s connected to all vertices (a, b) such that $a + b = n$ and a sink t connected to all vertices (a, b) such that $a + b = n + 1$. Let G_n be the flow network given by the directed graph just described (see Figure 14), with capacities assigned as follows:
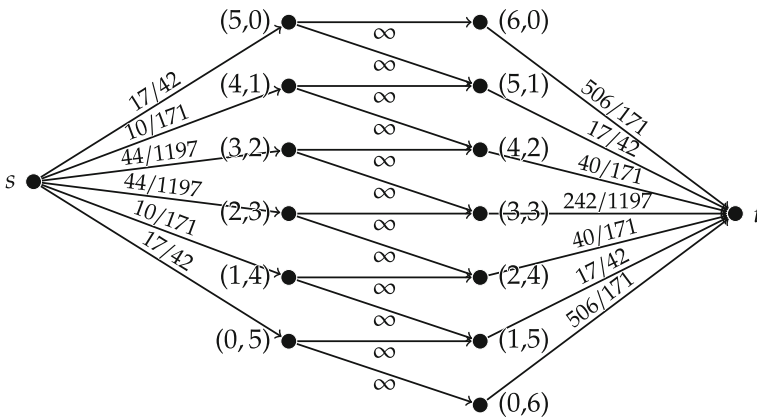

FIGURE 14 The flow network G_n for $n = 5$. Edges are labeled with their capacities: for example, $\frac{|T_5^4||T_0^4|}{6|STr_6|} = \frac{17}{42}$ and $\frac{|T_3^4||T_3^4|}{7|STr_7|} = \frac{242}{1197}$.

1. the edge from s to a pair (a, b) such that $a + b = n$ has capacity $C_{s \to (a,b)} = \frac{|T_a^4||T_b^4|}{(n+1)|STr_{n+1}|}$;

2. the edge from a pair (a, b) such that $a + b = n + 1$ to t has capacity $C_{(a,b) \to t} = \frac{|T_a^4||T_b^4|}{(n+2)|STr_{n+2}|}$;

3. edges not involving the source or sink have infinite capacity.

Note that we have $\sum_{a=0}^{N} |T_a^4||T_{N-a}^4| = (N + 1)|STr_{N+1}|$ by (2), so the total capacity of edges issued from the source is one, and the same is true for the total capacity of edges involving the sink.

The existence of a function h satisfying the requirements of the lemma is implied by the existence of a flow of value 1 on each network G_n. Indeed, given a unit flow $(f_e)_{e \in E(G_n)}$ on the flow network G_n, we may assume without loss of generality that $f_{(a,b) \to (a+1,b)} = f_{(b,a) \to (b,a+1)}$, as replacing $f_{(a,b) \to (a+1,b)}$ and $f_{(b,a) \to (b,a+1)}$ by $\frac{f_{(a,b) \to (a+1,b)} + f_{(b,a) \to (b,a+1)}}{2}$ for all a, b still yields a unit flow on G_n, thanks to the fact that $C_{s \to (a,b)} = C_{s \to (b,a)}$ and $C_{(a,b) \to t} = C_{(b,a) \to t}$.

We will now set $h(a+1, b) = \frac{f_{(a,b) \to (a+1,b)}}{C_{(a+1,b) \to t}}$ for all $(a, b) \in \mathbb{N}^2$ and $h(0, b) = 0$ for all $b \in \mathbb{N}$. When $a, b \geq 1$, we have

$$h(a, b)C_{(a,b) \to t} = f_{(a-1,b) \to (a,b)} = f_{(b,a-1) \to (b,a)} = C_{(a,b) \to t} - f_{(b-1,a) \to (b,a)} = (1 - h(a, b))C_{(a,b) \to t};$$

moreover, $h(0, b) = 0$ and $h(a, 0) = 1$ for all $a, b \geq 1$.

We also have

$$h(a + 1, b)C_{(a+1,b) \to t} + h(b + 1, a)C_{(b+1,a) \to t} = f_{(a,b) \to (a+1,b)} + f_{(b,a) \to (b+1,a)} =$$

$$= f_{(a,b) \to (a+1,b)} + f_{(a,b) \to (a,b+1)} = C_{s \to (a,b)},$$

which is exactly the second required expression for h.

We now claim the existence of a unit flow on G_n, and will proceed to show this claim by the max-flow-min-cut theorem via a series of quick lemmas, culminating in Corollary 5.6, which completes the proof of the existence of h. ∎

Let G_n be the flow network described within the proof of Lemma 5.3 and depicted by Figure 14. An s-t-cut of G_n is a partition of the vertices of G_n as $S \cup T$, where $s \in S$ and $t \in T$, and its capacity is the sum of the capacities of edges of G_n that join a vertex in S to a vertex in T. For $n \geq 1$, let $S_n = \{(a, b) \in \mathbb{N}^2 | a + b = n\}$, so that we can think of the vertices of G_n as $S_n \cup S_{n+1} \cup \{s, t\}$. We can represent an s-t-cut of G_n by the set $K = (S_n \cap T) \cup (S_{n+1} \cap S)$; this is a handy representation, as the capacity of the cut is infinite if there exists an edge (v, w), where $v \in S_n$ and $w \in S_{n+1}$, such that $v, w \notin K$, and it is equal to $\sum_{v \in K} 1_{v \in S_n} C_{s \to v} + 1_{v \in S_{n+1}} C_{v \to t}$ otherwise.

In the following series of Lemmas, we refer to the flow network G_n without redefining it and denote cuts by the corresponding set K as constructed above; we write $|K|$ for the capacity of the cut.

Lemma 5.4. *The cuts of G_n having smallest capacity are of the form $K_i = \{(n - j, j)|0 \leq j \leq i\} \cup \{(n + 1 - j, j)|i + 1 \leq j \leq n + 1\}$ for some $-1 \leq i < n$, or $K = S_n$.*

Proof. Note that all cuts listed in the statement have finite capacity.

Consider a cut K of G_n that is not of the form in the statement and suppose it has finite capacity (otherwise it is not minimal). If K strictly contains S_n or some K_i, then it cannot be minimal, so we may assume this is not the case.

Under these assumptions, there is a vertex $(n - i, i) \in K$ such that $i > 0$ and $(n - i + 1, i - 1) \notin K$. Indeed, if this were not the case then $K \cap S_n$ would be of the form $\{(n - j, j) | 0 \leq j \leq i\}$ for some $i > 0$, which, in order for K to have finite capacity, would imply that it contains K_i or (if $i = n$) S_n.

Consider the cut $K' = K \cup \{(n - i, i + 1)\} \setminus \{(n - i, i)\}$; we claim that its capacity is strictly smaller than that of K. Indeed, the cut K' has finite capacity: if there is an edge of G_n of infinite capacity that has no endpoint in K', then it must involve the vertex $(n - i, i)$ and cannot be the one joining $(n - i, i)$ to $(n - i, i + 1)$; it must therefore be the edge between $(n - i, i)$ and $(n - i + 1, i)$. However, this cannot be the case, because $(n - i + 1, i) \in K$ (otherwise K would not contain either endpoint of the edge from $(n - i + 1, i - 1)$ to $(n - i + 1, i)$ and therefore have infinite capacity) and so $(n - i + 1, i) \in K'$.

The difference in capacities is given by

$$|K'| - |K| = C_{(n-i,i+1)\to t} - C_{s\to(n-i,i)} = \frac{|T^4_{i+1}||T^4_{n-i}|}{(n+2)|STr_{n+2}|} - \frac{|T^4_i||T^4_{n-i}|}{(n+1)|STr_{n+1}|};$$

this quantity is strictly negative, as

$$\frac{T^4_{i+1}}{T^4_i} = \frac{4(4i+1)(4i+2)(4i+3)}{(3i+2)(3i+3)(3i+4)} = 4\prod_{j=1}^{3} \frac{4i+j}{3i+j+1} \tag{3}$$

is strictly increasing in i (each term in the product is), which implies that

$$\frac{(n+1)|STr_{n+1}||T^4_{i+1}||T^4_{n-i}|}{(n+2)|STr_{n+2}||T^4_i||T^4_{n-i}|} \leq \frac{(n+1)|STr_{n+1}||T^4_{n+1}|}{(n+2)|STr_{n+2}||T^4_n|} = \frac{4(4n+1)(4n+2)(4n+3)}{(3n+2)(3n+3)(3n+4)} \cdot \frac{(n+1)|STr_{n+1}|}{(n+2)|STr_{n+2}|}.$$

Using the expression from (2) for the number of simple triangulations, we get that the ratio above equals

$$\frac{4(4n+1)(4n+2)(4n+3)(n+1)(4n+1)!(n+2)!(3n+5)!}{(3n+2)(3n+3)(3n+4)(n+2)(4n+5)!(n+1)!(3n+2)!} = \frac{(4n+1)(3n+5)}{(3n+2)(4n+5)} < 1. \tag{4}$$

∎

Lemma 5.5. *Among the cuts of G_n of the form $K_i = \{(n - j, j) | 0 \leq j \leq i\} \cup \{(n + 1 - j, j) | i + 1 \leq j \leq n + 1\}$, where $-1 \leq i < n$, or S_n, the ones with smallest capacity are K_{-1} (which is S_{n+1}) and S_n, whose capacity is 1.*

Proof. We have the inequality $|K_{i+1}| > |K_i|$ for $i = -1, \ldots, n - 1$. Indeed, we have $K_{i+1} = K_i \cup \{(n - i - 1, i + 1)\} \setminus \{(n - i, i + 1)\}$, so

$$|K_{i+1}| - |K_i| = C_{s\to(n-i-1,i+1)} - C_{(n-i,i+1)\to t} = \frac{|T^4_{n-i-1}||T^4_{i+1}|}{(n+1)|STr_{n+1}|} - \frac{|T^4_{n-i}||T^4_{i+1}|}{(n+2)|STr_{n+2}|},$$

and this quantity is strictly positive by the argument of the previous lemma. Indeed, since $\frac{T^4_{a+1}}{T^4_a}$ is increasing in a by (3), we have

$$\frac{|K_{i+1}|}{|K_i|} = \frac{(n+2)|STr_{n+2}||T^4_{n-i-1}|}{(n+1)|STr_{n+1}||T^4_{n-i}|} \geq \frac{(n+2)|STr_{n+2}||T^4_n|}{(n+1)|STr_{n+1}||T^4_{n+1}|} > 1,$$

as it is the inverse of (4). It follows that $|K_{-1}| = 1 < |K_i|$ for $i \geq 0$; the capacity of the cut S_n is also 1. ∎

Finally, by the max-flow-min-cut theorem, we have obtained the following, which concludes the proof of Lemma 5.3:

Corollary 5.6. *For all $n \geq 1$, there exists a unit flow on the network G_n.*

We are now ready to prove the main result, that is, Theorem 2.

Proof. (Proof of Theorem 2) Let $(g_n)_{n \geq 0}$ be a growth scheme for complete four-ary trees, whose existence is guaranteed by Theorem 3. Define $f_n : \mathsf{STr}_{n+1} \times \mathsf{STr}_n \to [0, 1]$ as follows. Given a triangulation $\Delta \in \mathsf{STr}_k$ and $\varepsilon = \pm 1$, let $S^\varepsilon(\Delta)$ be the set of all pairs (L, R) of four-ary trees for which the corresponding blossoming tree τ is such that $\Xi(\tau, \varepsilon) = \Delta$. Note that $|S^1(\Delta) \cup S^{-1}(\Delta)| = 2k$, as the correspondence between pairs (L, R) and simple triangulations with $2k$ faces is $2k$-to-2 by Construction 2.

For $T \in \mathsf{STr}_{n+1}$ and $t \in \mathsf{STr}_n$, set

$$f_n(T, t) = \frac{1}{2n+2} \sum_{\varepsilon = \pm 1} \sum_{(L,R) \in S^\varepsilon(T)} \sum_{(l,r) \in S^\varepsilon(t)} h(|L|, |R|) 1_{R=r} g_{|l|}(L, l) + (1 - h(|L|, |R|)) 1_{L=l} g_{|r|}(R, r).$$

We now check the three properties that $(f_n)_{n \geq 1}$ needs to satisfy in order to be a growth scheme for simple triangulations.

First of all, in order for $f_n(T, t)$ to be nonzero, we must have $g_{|r|}(R, r) > 0$ or $g_{|l|}(L, l) > 0$ for some R, r or L, l such that $(R, L) \in S^\varepsilon(T)$ and $(r, l) \in S^\varepsilon(t)$ for some $\varepsilon = \pm 1$. This implies that $R = \mathrm{grow}(r, v)$ or $L = \mathrm{grow}(l, v)$ for some leaf v by definition of a growth scheme for four-ary trees. Since we have $T = \Xi(\tau, \varepsilon)$, where τ is the blossoming tree corresponding to (L, R), and $t = \Xi(\tau', \varepsilon)$, where τ' is the blossoming tree corresponding to (l, r), Lemma 5.2 implies the existence of an edge e of T such that $t = \mathrm{coll}(T, e)$, as wanted.

Now consider $\sum_{t \in \mathsf{STr}_n} f_n(T, t)$ for any fixed $T \in \mathsf{STr}_{n+1}$, that is,

$$\frac{1}{2n+2} \sum_{t \in \mathsf{STr}_n} \sum_{\varepsilon = \pm 1} \sum_{(L,R) \in S^\varepsilon(T)} \sum_{(l,r) \in S^\varepsilon(t)} h(|L|, |R|) 1_{R=r} g_{|l|}(L, l) + (1 - h(|L|, |R|)) 1_{L=l} g_{|r|}(R, r).$$

For ease of notation, denote by $F(L, R, l, r)$ the expression appearing within the repeated sums (which does not depend on ε). By changing the order of summation, we obtain

$$\frac{1}{2n+2} \sum_{\varepsilon = \pm 1} \sum_{(L,R) \in S^\varepsilon(T)} \sum_{t \in \mathsf{STr}_n} \sum_{(l,r) \in S^\varepsilon(t)} F(L, R, l, r),$$

where the two internal sums yield one term for each pair (l, r) in the set $\cup_{k=0}^{n-1} \mathsf{T}_k^4 \times \mathsf{T}_{n-1-k}^4$, of which the collection of sets $\{S^\varepsilon(t)\}_{t \in \mathsf{STr}_n}$ (where ε is fixed) forms a partition. The above sum is therefore equal to

$$\frac{1}{2n+2} \sum_{\varepsilon = \pm 1} \sum_{(L,R) \in S^\varepsilon(T)} \sum_{k=0}^{n-1} \sum_{l \in \mathsf{T}_k^4, r \in \mathsf{T}_{n-1-k}^4} F(L, R, l, r).$$

Using the fact that the expression for $F(L, R, l, r)$ is zero unless $R = r$ or $L = l$, one can further rewrite the above as

$$\frac{1}{2n+2} \sum_{\varepsilon=\pm 1} \sum_{(L,R)\in S^\varepsilon(T)} \left(\sum_{l\in \mathsf{T}^4_{|L|-1}} F(L,R,l,R) + \sum_{r\in \mathsf{T}^4_{|R|-1}} F(L,R,L,r) \right) =$$

$$= \frac{1}{2n+2} \sum_{\varepsilon=\pm 1} \sum_{(L,R)\in S^\varepsilon(T)} h(|L|,|R|) \sum_{l\in \mathsf{T}^4_{|L|-1}} g_{|l|}(L,l) + (1-h(|L|,|R|)) \sum_{r\in \mathsf{T}^4_{|R|-1}} g_{|r|}(R,r).$$

Finally, using the fact that $(g_n)_{n\geq 0}$ is a growth scheme for four-ary trees yields

$$\frac{1}{2n+2} \sum_{\varepsilon=\pm 1} \sum_{(L,R)\in S^\varepsilon(T)} h(|L|,|R|) + (1-h(|L|,|R|)) = \frac{1}{2n+2} \sum_{\varepsilon=\pm 1} \sum_{(L,R)\in S^\varepsilon(T)} 1 = \frac{2n+2}{2n+2} = 1,$$

where we have used the fact that $|S^1(T) \cup S^{-1}(T)| = 2(n+1)$.

Note that, by combining it with the obvious fact that $f_n(T,t) > 0$, the computation above also proves that the image of f_n is contained in $[0,1]$.

Finally, consider the sum $\sum_{T\in \mathsf{STr}_{n+1}} f_n(T,t)$ for a fixed $t \in \mathsf{STr}_n$. This time, we are computing

$$\frac{1}{2n+2} \sum_{\varepsilon=\pm 1} \sum_{(l,r)\in S^\varepsilon(t)} \sum_{T\in \mathsf{STr}_{n+1}} \sum_{(L,R)\in S^\varepsilon(T)} F(L,R,l,r) = \frac{1}{2n+2} \sum_{\varepsilon=\pm 1} \sum_{(l,r)\in S^\varepsilon(t)} \sum_{k=0}^{n} \sum_{L\in \mathsf{T}^4_k, R\in \mathsf{T}^4_{n-k}} F(L,R,l,r) =$$

$$= \frac{1}{2n+2} \sum_{\varepsilon=\pm 1} \sum_{(l,r)\in S^\varepsilon(t)} \sum_{L\in \mathsf{T}^4_{|l|+1}} h(|l|+1,|r|)g_{|l|}(L,l) + \sum_{R\in \mathsf{T}^4_{|r|+1}} (1-h(|l|,|r|+1))g_{|r|}(R,r),$$

where we have performed the same manipulations as for the previous sum. We can now use the fact that $(g_n)_{n\geq 0}$ is a growth scheme for four-ary trees to compute the value of $\sum_{L\in \mathsf{T}^4_{|l|+1}} g_{|l|}(L,l)$ and $\sum_{R\in \mathsf{T}^4_{|r|+1}} g_{|r|}(R,r)$, which gives

$$\frac{1}{2n+2} \sum_{\varepsilon=\pm 1} \sum_{(l,r)\in S^\varepsilon(t)} \left(h(|l|+1,|r|)\frac{|\mathsf{T}^4_{|l|+1}|}{|\mathsf{T}^4_{|l|}|} + (1-h(|l|,|r|+1))\frac{|\mathsf{T}^4_{|r|+1}|}{|\mathsf{T}^4_{|r|}|} \right) =$$

$$= \frac{1}{2n+2} \sum_{\varepsilon=\pm 1} \sum_{(l,r)\in S^\varepsilon(t)} \frac{(|l|+|r|+2)|\mathsf{STr}_{|l|+|r|+2}|}{(|l|+|r|+1)|\mathsf{STr}_{|l|+|r|+1}|} = \frac{2n}{2n+2} \cdot \frac{(n+1)|\mathsf{STr}_{n+1}|}{n|\mathsf{STr}_n|} = \frac{|\mathsf{STr}_{n+1}|}{|\mathsf{STr}_n|},$$

as wanted, thanks to the properties of h expressed in Lemma 5.3. ∎

ACKNOWLEDGMENTS

REFERENCES

1. L. Addario-Berry, *Growing random 3-connected maps or comment s'enfuir de l'hexagone*, Electron. Commun. Probab. **19** (2014), no. 54, 1–12.
2. M. Albenque and D. Poulalhon, *A generic method for bijections between blossoming trees and planar maps*, Electron J Combin **22** (2015), no. 2, 1–44.
3. L. Alonso, J. L. Rémy, and R. Schott, *A linear-time algorithm for the generation of trees*, Algorithmica **17** (1997), 162–182.
4. O. Angel and O. Schramm, *Uniform infinite planar triangulation*, Commun. Math. Phys. **241** (2003), 191–213.
5. O. Bernardi and É. Fusy, *A bijection for triangulations, quadrangulations, pentagulations, etc*, J Combin Theory Ser A **119** (2012), 218–244.
6. O. Bernardi and É. Fusy, *Unified bijections for maps with prescribed degrees and girth*, J Combin Theory Ser A **119** (2012), 1351–1387.
7. J. Bettinelli, *Increasing forests and quadrangulations via a bijective approach*, J Combin Theory Ser A **122** (2014), 107–125.
8. J. Bouttier, P. Di Francesco, and E. Guitter, *Planar maps as labeled mobiles*, Electron. J. Combin. **11** (2004), 27.
9. A. Caraceni, *A polynomial upper bound for the mixing time of edge rotations on planar maps*, Electron. J. Probab. **25** (2020), 1–30.
10. A. Caraceni and A. Stauffer, *Polynomial mixing time of edge flips on quadrangulations*, Probab. Theory Relat. Fields **176** (2020), 35–76.
11. R. Cori and B. Vauquelin, *Planar maps are well labeled trees*, Can. J. Math. **33** (1981), 1023–1042.
12. B. Haas and R. Stephenson, *Scaling limits of multi-type Markov branching trees*, Probab. Theory Relat. Fields **180** (2021), 727–797.
13. J.-F. Le Gall and G. Miermont. Scaling limits of random trees and planar maps, in Probability and statistical physics in two and more dimensions, vol. 15 of clay math. Proc., Amer. Math. Soc., Providence, RI. 2012 155–211.
14. M. Luczak and P. Winkler, *Building uniformly random subtrees*, Random Struct. Algoritm. **24** (2004), 420–443.
15. J.-F. Marckert. Growing random uniform d-ary trees. 2021.
16. D. Poulalhon and G. Schaeffer, *Optimal coding and sampling of triangulations*, Algorithmica **46** (2006), 505–527.
17. J.-L. Rémy, *Un procédé itératif de dénombrement d'arbres binaires et son application à leur génération aléatoire*, RAIRO Inform. Théor. **19** (1985), 179–195.
18. G. Schaeffer. Conjugaison d'arbres et cartes combinatoires aléatoires. PhD thesis. 1998.
19. W. T. Tutte, *A census of planar triangulations*, Can. J. Math. **14** (1962), 21–38.
20. W. T. Tutte, *A census of planar maps*, Can. J. Math. **15** (1963), 249–271.