



Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

**Journal of Combinatorial Theory,
Series A**

www.elsevier.com/locate/jcta



Increasing forests and quadrangulations via a bijective approach [☆]



Jérémie Bettinelli

Institut Élie Cartan de Lorraine, B.P. 239, F-54506 Vandœuvre-lès-Nancy Cedex, France

ARTICLE INFO

Article history:

Received 6 December 2012

Available online 29 October 2013

Keywords:

Map

Tree

Forest

Bijection

Polygons gluing

Graph on surface

Random discrete surface

ABSTRACT

In this work, we expose four bijections each allowing to increase (or decrease) one parameter in either uniform random forests with a fixed number of edges and trees, or quadrangulations with a boundary having a fixed number of faces and a fixed boundary length. In particular, this gives a way to sample a uniform quadrangulation with $n + 1$ faces from a uniform quadrangulation with n faces or a uniform forest with $n + 1$ edges and p trees from a uniform forest with n edges and p trees.

© 2013 Elsevier Inc. All rights reserved.

Contents

1.	Introduction	108
2.	Setting and presentation of the results	108
3.	Increasing forests	110
3.1.	Adding an edge	111
3.2.	Adding a tree	112
4.	Modifying the length of the boundary in a quadrangulation	113
4.1.	Adding two half-edges to the boundary	114
4.2.	Removing two half-edges from the boundary	115
4.3.	These mappings are reverse one from another	116
5.	Changing the number of faces in a quadrangulation	118
5.1.	Removing a face	119
5.1.1.	Confluent face	119
5.1.2.	Simple face	120

[☆] This work is partially supported by ANR-08-BLAN-0190.

E-mail address: jeremie.bettinelli@normalesup.org.

URL: <http://www.normalesup.org/~bettinel>.

5.2.	Adding a face	121
5.2.1.	Confluent case	121
5.2.2.	Simple case	122
5.3.	These mappings are reverse one from another	123
5.4.	Interpretation through the Bouttier–Di Francesco–Guitter bijection	124
References		125

1. Introduction

Maps are known to have a lot of applications in different fields of mathematics, computer science and physics. These applications strongly rely on their combinatorial structures, which have been widely investigated during the last few decades. A particularly interesting class of maps is the class of quadrangulations with or without a boundary, which are for example natural candidates for discretizing surfaces.

Our interest in this paper, inspired from Rémy’s algorithm [7] on growing trees, is in finding a natural way to grow a planar quadrangulation that preserves the uniform measure. For instance, we will present a bijection between quadrangulations with a boundary having n faces and $2p$ half-edges on the boundary carrying some distinguished elements and quadrangulations with a boundary having $n + 1$ faces and $2p$ half-edges on the boundary also carrying distinguished elements. Our bijection is designed in such a way that the number of possibilities for distinguishing the necessary elements only depends on the size and boundary length of the quadrangulations. Forgetting these distinguished elements, we obtain a way to sample a uniform quadrangulation with a boundary having a prescribed number of faces and boundary length from a uniform quadrangulation with a boundary having one less face. In other words, there exist some integer constants $c_{n,p}$ and $c'_{n,p}$ such that our construction provides a $c_{n,p}$ -to- $c'_{n,p}$ mapping between the set of quadrangulations with a boundary having n faces and $2p$ half-edges on the boundary and the set of quadrangulations with a boundary having $n + 1$ faces and $2p$ half-edges on the boundary.

In addition to the probabilistic point of view, such bijections also present a combinatorial interest as they provide an interpretation to some combinatorial identities. The bijections we present in this work interpret already known identities so that they actually provide alternate proofs for these identities. In the future, we hope that other similar bijections will allow to solve some open enumeration problems.

Our method of “cut and glue” bijections, which could informally be pictured as unbuttoning a shirt and buttoning it back incorrectly by putting each button into the hole that immediately follows the correct one, possesses a certain robustness and can be declined in many ways. We present here in detail four such bijections by focusing on forests and quadrangulations with a boundary. In an upcoming work, we plan to present more bijections relying on the same idea. In particular, one of these bijections will allow to recover Tutte’s formula [9] counting the number of planar maps with n faces having prescribed degrees a_1, \dots, a_n where at most two a_i ’s are odd numbers.

It has also been pointed to us that our method somehow recalls a work by Cori [5] where he used a so-called transfer bijection roughly consisting in transferring one degree from a face to a neighboring face. Using a properly defined chain, this allows to transfer one degree from a face to any other face, step by step. We do not believe that his results are related to the present work but we think they are worth mentioning at this point.

2. Setting and presentation of the results

Recall that a planar map is an embedding of a finite connected graph (possibly with loops and multiple edges) into the two-dimensional sphere, considered up to direct homeomorphisms of the sphere. The faces of the map are the connected components of the complement of edges. We will call *half-edge* an edge carrying one of its two possible orientations. We say that a half-edge h is *incident* to a face f (or that f is incident to h) if h belongs to the boundary of f and is oriented in such a way




Fig. 1. Example of quadrangulation with a boundary having 19 internal faces and 46 half-edges on the boundary. The root corner is represented by the rake-like symbol in the middle part of the picture.

that f lies to its left. A *corner* is an angular sector of a face delimited by two consecutive half-edges incident to the face. We will implicitly consider our maps to be rooted, which means that one corner of one face is distinguished. This distinguished corner will be called the *root corner*.

A *quadrangulation with a boundary* is a particular instance of planar map whose faces are all incident to exactly 4 half-edges, with the exception of the face containing the root corner, which may be of arbitrary even degree. The latter face will be called the *external face*, whereas the other ones will be called *internal faces*. To match this terminology, we will as often as possible draw the external face as the infinite component of the plane on our figures. As a result, note that, since the cyclic ordering of edges around a vertex is prescribed, it is unambiguous to speak of clockwise and counterclockwise order. The half-edges incident to the external face will constitute the *boundary* of the map. Beware that we do not require here the boundary to be a simple curve. See Fig. 1.

Definition 1. For $n \geq 0$ and $p \geq 1$, we will denote by $\mathcal{Q}_{n,p}$ the set of all quadrangulations with a boundary having n internal faces and $2p$ half-edges on the boundary.

Let us make a few observations:

- ◊ As we do not require the boundary of a quadrangulation with a boundary to be a simple curve, the set $\mathcal{Q}_{n,p}$ is never empty.
- ◊ There is a trivial bijection between $\mathcal{Q}_{n+1,1}$ and $\mathcal{Q}_{n,2}$: in the direct sense, just remove the one edge of the boundary that is directly to the right of the root. Conversely, double the edge directly to the left of the root in such a way that the degree-2 face so created contains the root.
- ◊ The two previous sets correspond to the set of quadrangulations (without boundary) with $n + 1$ faces, so our approach actually contains this case.
- ◊ The set $\mathcal{Q}_{0,p}$ is the set of rooted plane trees having p edges.
- ◊ All the cycles in a quadrangulation with a boundary have an even length. Indeed, if there were cycles with an odd length, let us consider one of these cycles. By the Jordan curve theorem, it splits the map into two maps. Let us consider any of these maps. All its faces except one have an even degree. This is in contradiction with the fact that the number of half-edges is necessarily even.
- ◊ Using Euler's characteristic formula, it is easy to see that an element of $\mathcal{Q}_{n,p}$ has exactly
 - n internal faces,
 - $2p$ half-edges on the boundary,
 - $n + p + 1$ vertices,
 - $2n + p$ edges.

The cardinality of the set $\mathcal{Q}_{n,p}$ may be computed by several methods. We may use Tutte's formula [9], a generating series method [4] or a bijective approach using the Bouttier–Di Francesco–Guitter bijection [3] (see for example [2]):

$$|\mathcal{Q}_{n,p}| = \frac{3^n (2p)! (2n + p - 1)!}{p! (p - 1)! n! (n + p + 1)!}. \tag{1}$$

In this work, we present two bijections accounting for the following combinatorial identities:

$$(2p + 1)(2p + 2)(2n + p)|\mathcal{Q}_{n,p}| = p(p + 1)(n + p + 2)|\mathcal{Q}_{n,p+1}|, \tag{2}$$

$$3(2n + p)(2n + p + 1)|\mathcal{Q}_{n,p}| = (n + 1)(n + p + 2)|\mathcal{Q}_{n+1,p}|. \tag{3}$$

In particular, using the initial condition $\mathcal{Q}_{0,1} = 1$, our approach also provides a new proof for (1).

The strategy we use is roughly the following. Thanks to some distinguished elements of a map (faces, edges, vertices, corners), we construct a path in the map. Then, we “cut” along this path and “glue” back after shifting a little bit. This operation mildly modifies the map along the path and changes its structure around the extremities of the path, creating new distinguished elements. In order for this operation to work, the path we construct has to be totally recoverable from the new distinguished elements. In general, the notion of left-most geodesics or right-most geodesics will allow this.

As a warm-up, we present in Section 3 two bijections between forests. Although not directly related to our bijections on quadrangulations, they use the same kind of ideas and are easier to handle. In addition, some definitions and notation that will be used throughout this paper are also provided at this point. We then present in Section 4 our bijection interpreting identity (2) and Section 5 is devoted to our last bijection, corresponding to (3).

Let us also mention that our bijections on maps are not coming from bijections on forests that are transferred through the Bouttier–Di Francesco–Guitter bijection [3]; our bijections work directly on the maps. See Section 5.4 for more details regarding this matter. For more simplicity, we will from now on use the term quadrangulation to mean rooted quadrangulation with a boundary.

3. Increasing forests

This section presents two bijections allowing to grow uniform forests with a fixed number of trees and edges. These bijections are simpler than our main bijections for quadrangulations and, as they are not the main point of this work and as there are no real difficulties in handling them, we leave the proofs to the reader.

We call *tree* a rooted planar map with only one face. Beware that in particular, the trees we consider are embedded in the plane and that the vertex tree (the only tree with no edges) is allowed here. For $n \geq 0$ and $p \geq 1$, we call *forest* with p trees and n edges a p -uple of trees where the total number of edges is n . We write $\mathcal{F}_{n,p}$ the set of forests with p trees and n edges. A simple application of the so-called cycle lemma (see for example [6, Section 6.1] or [1, Lemma 3]) yields that

$$|\mathcal{F}_{n,p}| = \frac{p}{2n + p} \binom{2n + p}{n}.$$

It will be convenient to add extra half-edges between the roots of successive trees: if $f = (t_1, \dots, t_p)$ is a forest, we add, for $1 \leq i \leq p$, a half-edge from ρ_i to ρ_{i+1} , where ρ_i is the root of t_i and $\rho_{p+1} := \rho_1$. With this convention, a forest corresponds to a map with two faces, one of which being of degree p and having a simple curve as a boundary (see Fig. 3). The *corners* of the forest are defined as the corners of the other face of this map. It is easy to see that an element of $\mathcal{F}_{n,p}$ has

- ◊ n edges,
- ◊ $n + p$ vertices,
- ◊ $2n + p$ corners.

Before we begin, we need to introduce some vocabulary. For a half-edge h , we write h^- its origin, h^+ its end, and $\text{rev}(h)$ its reverse. It will be convenient to consider corners as half-edges having no origin, only an end. In particular, if c is a corner, we will write c^+ the vertex corresponding to it, that is, if c is the corner delimited by the consecutive half-edges h and h' , then $c^+ := h^+ = h'^-$.

Definition 2. A *path* from a vertex v to a vertex v' is a finite sequence $\mathbb{p} = (\mathbb{p}_1, \mathbb{p}_2, \dots, \mathbb{p}_\ell)$ of half-edges such that $\mathbb{p}_1^- = v$, for $1 \leq k \leq \ell - 1$, $\mathbb{p}_k^+ = \mathbb{p}_{k+1}^-$, and $\mathbb{p}_\ell^+ = v'$. Its *length* is the integer $[\mathbb{p}] := \ell$. We will use the convention that an empty path has length 0.




Fig. 2. The two different ways of distinguishing twice the same corner.

A path \mathbb{p} is called *self-avoiding* if it does not meet twice the same vertex, that is,

$$|\{\mathbb{p}_1^-, \dots, \mathbb{p}_{[\mathbb{p}]}, \mathbb{p}_{[\mathbb{p}]}^+\}| = [\mathbb{p}] + 1.$$

The *reverse* of a path $\mathbb{p} = (\mathbb{p}_1, \mathbb{p}_2, \dots, \mathbb{p}_\ell)$ is the path $\text{rev}(\mathbb{p}) := (\text{rev}(\mathbb{p}_\ell), \text{rev}(\mathbb{p}_{\ell-1}), \dots, \text{rev}(\mathbb{p}_1))$.

Let \mathbb{p} be a path. We denote by $\mathbb{p}_{i \rightarrow j}$ the path $(\mathbb{p}_i, \dots, \mathbb{p}_j)$ if $1 \leq i \leq j \leq [\mathbb{p}]$, or the empty path otherwise. If q is another path satisfying $q_1^- = \mathbb{p}_{[\mathbb{p}]}^+$, we set

$$\mathbb{p} \bullet q := (\mathbb{p}_1, \dots, \mathbb{p}_{[\mathbb{p}]}, q_1, \dots, q_{[q]})$$

the concatenation of \mathbb{p} and q . Throughout this paper, the notion of metric we use is the graph metric: if m is a map, the distance $d_m(v, v')$ between two vertices v and v' is the smaller ℓ for which there exists a path of length ℓ from v to v' . A *geodesic* from v to v' is such a path.

We will also say that a half-edge is *directed toward* a set if its end is strictly closer to the set than its origin. A half-edge is *directed away* from a set if its reverse is directed toward the set. In what follows, we will always use the convention that distinguishing a corner “splits” it into two new corners. In other words, when we distinguish the same corner for the second time, we have to specify which of its two sides is distinguished (see Fig. 2).

3.1. Adding an edge

Our first bijection is between the set $\mathcal{F}_{n,p}^{\text{cc}}$ of forests from $\mathcal{F}_{n,p}$ carrying two distinguished corners and the set $\mathcal{F}_{n+1,p}^{\text{ev}}$ of forests from $\mathcal{F}_{n+1,p}$ having a distinguished edge and a distinguished vertex. It provides a combinatorial interpretation to the following identity:

$$(2n + p)(2n + p + 1)|\mathcal{F}_{n,p}| = (n + 1)(n + p + 1)|\mathcal{F}_{n+1,p}|.$$

Let $(f; c, c') \in \mathcal{F}_{n,p}^{\text{cc}}$. We let \mathbb{p} be the only self-avoiding path from c^+ to c'^+ (recall that we added half-edges between the roots of the trees). Starting from the first tree of f and following the contour of its trees one by one gives a natural ordering of its corners. A classical way of picturing this order is to imagine an animal flying around the first tree, then jumping to the second one and so on. We “cut” along \mathbb{p} between c and c' in the following sense: we consider, on the one hand, the part of f made of all the elements encountered between c and c' in its contour and, on the other hand, the part made of all the elements encountered between c' and c , together with the path \mathbb{p} . We call these parts respectively *left part* and *right part*. The path \mathbb{p} has a copy in each part: in the left part, we consider c as an extra half-edge and set $\mathbb{l} := c \bullet \mathbb{p}$; in the right part, we consider c' as an extra half-edge and set $\mathbb{r} := \mathbb{p} \bullet \text{rev}(c')$. Then, we define the forest f' by gluing back together both parts while matching \mathbb{l}_k with \mathbb{r}_k for $1 \leq k \leq [\mathbb{l}] = [\mathbb{r}]$. In f' , we define e as the edge corresponding to $\mathbb{l}_{[0]}$ and $v := \mathbb{l}_1^-$. We set $\Psi_{n \uparrow, p}(f; c, c') := (f'; e, v)$. See Fig. 3.

Conversely, let $(f'; e, v) \in \mathcal{F}_{n+1,p}^{\text{ev}}$. We denote by \bar{e} the half-edge corresponding to e directed away from v , and \mathbb{p} the only self-avoiding path from v to \bar{e}^- . We will cut along $\mathbb{p} \bullet \bar{e}$. To this end, we need to specify between which corners we cut. We cut along $\mathbb{p} \bullet \bar{e}$ between the corner delimited by $(\mathbb{p} \bullet \bar{e})_1$ and its predecessor in the contour and the corner delimited by $\text{rev}(\bar{e})$ and its predecessor. In other words, we cut along $\mathbb{p} \bullet \bar{e}$, starting from the first corner to the left before $\mathbb{p} \bullet \bar{e}$ and stopping at the first corner to the right after $\mathbb{p} \bullet \bar{e}$. As above, this defines a left part and a right part. We denote by \mathbb{l} the image of $\mathbb{p} \bullet \bar{e}$ in the left part and \mathbb{r} its image in the right part. We define the forest f by gluing




Fig. 3. Adding an edge to a forest. The rake-like symbol represents the first tree of the forest.

back these parts while matching \mathbb{Q}_{k+1} with σ_k for $1 \leq k \leq [\mathbb{Q}] - 1$. Because of how we chose the corners between which we cut, in the left part, nothing except \mathbb{Q}_1 is attached to \mathbb{Q}_1^- . As a result, \mathbb{Q}_1 , which is not glued with anything, gives birth to a corner, which we denote by c . Similarly, $\text{rev}(\sigma_{[\sigma]})$ defines a corner c' and we set $\Psi_{n+1\downarrow,p}(\mathfrak{f}'; e, v) := (\mathfrak{f}; c, c')$.

Theorem 1. *The mappings $\Psi_{n\uparrow,p} : \mathcal{F}_{n,p}^{\text{cc}} \rightarrow \mathcal{F}_{n+1,p}^{\text{ev}}$ and $\Psi_{n+1\downarrow,p} : \mathcal{F}_{n+1,p}^{\text{ev}} \rightarrow \mathcal{F}_{n,p}^{\text{cc}}$ are one-to-one and reverse one from another.*

3.2. Adding a tree

Our second bijection on forests is between the set $\mathcal{F}_{n,p}^{\text{ci}}$ of forests from $\mathcal{F}_{n,p}$ carrying a distinguished corner together with an integer in $\{1, \dots, p + 1\}$ and the set $\mathcal{F}_{n,p+1}^{\text{vi}}$ of forests from $\mathcal{F}_{n,p+1}$ carrying a distinguished vertex together with an integer in $\{1, \dots, p\}$. It accounts for the following combinatorial identity:

$$(2n + p)(p + 1)|\mathcal{F}_{n,p}| = (n + p + 1)p|\mathcal{F}_{n,p+1}|.$$

Let us take $(\mathfrak{f}; c, i) \in \mathcal{F}_{n,p}^{\text{ci}}$. The vertex c^+ belongs to some tree of \mathfrak{f} : let $j \in \{1, \dots, p\}$ be its index and c' its last corner. With $(\mathfrak{f}; c, c')$, we define the same path \mathbb{p} as at the beginning of Section 3.1 and we perform the same operation to define a left and a right part. In the left part, we still consider c as an extra half-edge and set $\mathbb{Q} := c \bullet \mathbb{p}$. In the right part, we add an extra tree consisting in one vertex between the j -th and the $j + 1$ -th tree of \mathfrak{f} . We denote by h the half-edge linking the root of the j -th tree to this extra tree and define $\sigma := \mathbb{p} \bullet h$. Then, as above, we glue back together the left part and the right part while matching \mathbb{Q}_k with σ_k for $1 \leq k \leq [\mathbb{Q}] = [\sigma]$ and we define $v := \mathbb{Q}_1^-$. Finally, we define \mathfrak{f}' by re-rooting this forest in such a way that v belongs to the i -th tree and we set $\Psi_{n,p\uparrow}(\mathfrak{f}'; c, i) := (\mathfrak{f}'; v, j)$. See Fig. 4.

Note that the somehow surprising re-rooting accounts for the factors $p + 1$ in the left-hand side and p in the right-hand side of the combinatorial identity. Without this re-rooting, we would miss the forests carrying the vertex v in their last tree.

Conversely, let us take $(\mathfrak{f}'; v, j) \in \mathcal{F}_{n,p+1}^{\text{vi}}$. Let i be the index of the tree to which v belongs and consider the path \mathbb{p} from v to the root of the $i + 1$ -th tree (with the convention that the $i + 1$ -th tree is the first tree is $i = p + 1$). We cut along this path between the corner delimited by \mathbb{p}_1 and its predecessor in the contour and the last corner of the $i + 1$ -th tree. This again defines a left part and a right part; we denote respectively by \mathbb{Q} and σ the images of \mathbb{p} in these parts. We define a new forest by gluing them back while matching \mathbb{Q}_{k+1} with σ_k for $1 \leq k \leq [\mathbb{Q}] - 1$. As before, \mathbb{Q}_1 gives birth to a corner c . Moreover, the $i + 1$ -th tree of this new forest consists in only one vertex; we remove it. Finally, we define \mathfrak{f} by re-rooting the latter forest in such a way that c^+ belongs to the j -th tree, and we set $\Psi_{n,p+1\downarrow}(\mathfrak{f}'; v, j) := (\mathfrak{f}; c, i)$.




Fig. 4. Adding a tree to forest. In this example, $i = 4$ and $j = 6$.

Theorem 2. The mappings $\Psi_{n,p\uparrow} : \mathcal{F}_{n,p}^{ci} \rightarrow \mathcal{F}_{n,p+1}^{vi}$ and $\Psi_{n,p+1\downarrow} : \mathcal{F}_{n,p+1}^{vi} \rightarrow \mathcal{F}_{n,p}^{ci}$ are one-to-one and reverse one from another.

4. Modifying the length of the boundary in a quadrangulation

We leave forests and concentrate on quadrangulations again. In this section, we will use the corners of the external face. Recall the convention (stated right before Section 3.1) we use for distinguishing corners. As the external face of a quadrangulation from $\mathcal{Q}_{n,p}$ already has the root corner distinguished, we will consider that it has $2p + 1$ corners. If we distinguish one of these $2p + 1$ corners, it will then have $2p + 2$ corners. On the one hand, we consider the set $\mathcal{Q}_{n,p}^{ecc}$ of quadrangulations from $\mathcal{Q}_{n,p}$ carrying

- ◇ one distinguished edge e ,
- ◇ one first distinguished corner c of the external face,
- ◇ one second distinguished corner c' of the external face.

On the other hand, we consider the set $\mathcal{Q}_{n,p+1}^{vhh}$ of quadrangulations from $\mathcal{Q}_{n,p+1}$ carrying

- ◇ one distinguished vertex v ,
- ◇ one first half-edge h of the boundary directed toward v ,
- ◇ one second half-edge $h' \neq h$ of the boundary directed toward v .

Note that there are exactly $p + 1$ half-edges of the boundary that are directed toward v . To see this, label the vertices of the boundary with their distance to v . As there are no cycles of odd length in a quadrangulation, the labels encountered when traveling along the boundary form a $2(p + 1)$ -step bridge whose steps are either $+1$ or -1 . As a result, exactly half of them are -1 steps.

We will present an explicit bijection between the previous two sets; this will provide a combinatorial interpretation to (2).

We will often use the notion of *left-most geodesic* from a half-edge h (or a corner) to some set S of vertices (typically a vertex, an edge or a face). It is constructed as follows. First, we consider all the geodesics from h^+ to the closest elements of S . We take the set of all the first steps of these geodesics. Starting from h , we select the first half-edge to its left that belongs to this set. In other words, we turn clockwise around h^+ and select the first half-edge of this set that we meet. Note that this half-edge may be $\text{rev}(h)$ if this is the only half-edge in the set. Then we iterate the process from this half-edge until we reach S . Note that this path may be empty if $h^+ \in S$ and that it is a geodesic. The *right-most geodesic* from a half-edge or a corner to some set is defined in a similar way, by changing left to right in the previous definition.




Fig. 5. Example where the right-most geodesic from \bar{e} to c (on the bottom) is not the reverse of the left-most geodesic from c to e (on the top).




Fig. 6. The mapping from $Q_{n,p}^{ecc}$ to $Q_{n,p+1}^{vhh}$ in the simple case. We define the path p , then cut along it and glue back after shifting the parts one notch. This creates two new half-edges on the boundary, h and h' . On this example, $[c] = 4$ and $[c'] = 3$.

4.1. Adding two half-edges to the boundary

We begin with the mapping from $Q_{n,p}^{ecc}$ to $Q_{n,p+1}^{vhh}$. Let $(q; e, c, c') \in Q_{n,p}^{ecc}$. Note that, as q is a quadrangulation, one of the extremities of e is strictly closer to c^+ than the other. We denote by \bar{e} the half-edge corresponding to e directed toward c^+ , and c the right-most geodesic from \bar{e} to c^+ . Beware that $rev(c)$ is not necessarily the left-most geodesic from c to e : there may exist a path of same length leaving c to the left and meeting it again from the right, as shown in Fig. 5.

Until further notice, we suppose that $rev(\bar{e})$ is directed toward c'^+ : in this case, we say that the quadruple $(q; e, c, c')$ is simple. We denote by c' the right-most geodesic from $rev(\bar{e})$ to c'^+ . It is not hard to see that, in this case, $rev(c)$ is actually the left-most geodesic from c to e : indeed, a path leaving $rev(c)$ to the left and meeting it from the right has to intersect c' ; as a result, this path cannot be a geodesic as $rev(\bar{e})$ is directed toward c'^+ . We split the map q in two parts between c and c' along the self-avoiding path

$$p := rev(c) \bullet rev(\bar{e}) \bullet c'.$$

Let us call *left part* (resp. *right part*) the part consisting of the path p together with all the elements located to its left (resp. its right). In the left part, we replace c with an extra half-edge h and set $l := h \bullet p$. In the right part, we replace c' with an extra half-edge h' and set $r := p \bullet rev(h')$. Then, we glue back the two parts together in such a way that l_k coincides with r_k for $1 \leq k \leq [l] = [r]$. We denote by q' the map we obtain and we set $v := [c]_{[c]+1}^+ = (rev(r))_{[c']+1}^+$. See Fig. 6.




Fig. 7. The shifting operation in the left-pinned case. We shift along the path p and circumvent the part where p is pinched. On this picture, $[e] = 3$ and $[c'] = 2$.

Remark. Note that, when sliding one notch along the path, the marked edge is duplicated and the intersection between the two copies of this edge will become the new marked vertex. Conversely, a marked vertex will be duplicated into two vertices delimiting a new marked edge. This general principle will often be used in what follows.

Now, let us suppose that \vec{e} is directed toward c'^+ . We denote by c' the right-most geodesic from \vec{e} to c'^+ . Let $i \geq 1$ be the smallest integer such that $c_i \neq c'_i$. Note that c and c' do not meet again after time i , in the sense that $\{c_i^+, \dots, c_{[e]}^+\}$ and $\{c'_i^+, \dots, c'_{[c']}\}$ are disjoint sets. If c'_i is to the left of c_i after c_{i-1} (with the convention $c_0 = c'_0 := \vec{e}$), we say that the quadruple $(q; e, c, c')$ is *left-pinned*; we say that it is *right-pinned* in the other case. This terminology comes from the fact that the path

$$p := \text{rev}(c) \bullet \text{rev}(\vec{e}) \bullet \vec{e} \bullet c'$$

is “pinned,” the pinched part being to the left or right of the path $\text{rev}(c_{i \rightarrow [e]}) \bullet c'_{i \rightarrow [c']}$. Let us suppose here that $(q; e, c, c')$ is left-pinned. Similarly as before, we split the map in two parts. The right part is defined as previously. The definition of the left part, however, is slightly different: the difference is that we cut along the pinched part $\vec{e} \bullet c_{1 \rightarrow i-1}$ (see the middle part of Fig. 7). We define h, h' and the paths l and r by the same method as above. Note that here l is self-avoiding whereas r is not. We then proceed as before: we glue the parts while matching l_k with r_k for $1 \leq k \leq [l]$. Let q' be the map we obtain and let us define $v := l_{[e]+1}^+ = (\text{rev}(r))_{[c']\downarrow+2}^+$. See Fig. 7.

To treat the right-pinned case, we use a symmetric argument, by exchanging the roles of c and c' . Note that $(q; e, c, c')$ is right-pinned if and only if $(q; e, c', c)$ is left-pinned. We define the quadruple $(q'; v, h', h)$ corresponding to $(q; e, c', c)$ by the previous construction. We then associate with $(q; e, c, c')$ the quadruple $(q'; v, h, h')$. In all cases, we set $\Phi_{n,p\uparrow}(q; e, c, c') := (q'; v, h, h')$.

Remark. As a forest may be seen as a particular 2-face map, one might wonder whether $\Phi_{n,p\uparrow}$ is a simple extension of the mapping presented in Section 3.1. This is not the case as its restriction $\Phi_{0,p\uparrow}$ to trees is not the same mapping as $\Psi_{p\uparrow,1}$ from Section 3.1. Interestingly, this provides yet another mapping on trees, leaving us with three different mappings: Rémy’s algorithm, $\Phi_{0,p\uparrow}$ and $\Psi_{p\uparrow,1}$.

4.2. Removing two half-edges from the boundary

Let us now describe the reverse mapping $\Phi_{n,p+1\downarrow}$. We will also have to distinguish between three cases, and we try to keep our notation coherent with the previous section. Let $(q'; v, h, h') \in \mathcal{Q}_{n,p+1}^{\text{vhh}}$. Let h_0 denote the corner delimited by h and its predecessor in the contour of the external face, and h the left-most geodesic from this corner to v . Note that, as h is directed toward v and is the

first half-edge to the left after \mathfrak{h}_0 , we necessarily have $[\mathfrak{h}] \geq 1$ and $\mathfrak{h}_1 = h$. We define \mathfrak{h}'_0 and \mathfrak{h}' in a similar way with h' instead of h . As $h \neq h'$, there is no ambiguity in the definition of \mathfrak{h}_0 and \mathfrak{h}'_0 .

If \mathfrak{h} and \mathfrak{h}' do not meet before reaching v , we say that $(q'; v, h, h')$ is *simple*. In this case, we set $\mathfrak{p} := \mathfrak{h} \bullet \text{rev}(\mathfrak{h}')$. Notice that this path is self-avoiding. As before, we define a left part and a right part by cutting along \mathfrak{p} between the corners \mathfrak{h}_0 and \mathfrak{h}'_0 . We denote by \mathfrak{l} and \mathfrak{r} the path \mathfrak{p} in these respective parts. We define q by gluing back the two parts while matching \mathfrak{l}_{k+1} with \mathfrak{r}_k for $1 \leq k \leq [\mathfrak{l}] - 1$. As in Section 3.1, \mathfrak{l}_1 and $(\text{rev}(\mathfrak{r}))_1$ define two corners c and c' in q . Finally, we let e be the edge corresponding to $\mathfrak{l}_{[\mathfrak{h}]+1}$ and $(\text{rev}(\mathfrak{r}))_{[\mathfrak{h}'+1]}$.

If \mathfrak{h} and \mathfrak{h}' meet before reaching v , we let i and j be the smallest integers such that $\mathfrak{h}_i^+ = \mathfrak{h}'_j^+$. Note that \mathfrak{h} and \mathfrak{h}' merge after times i and j , in the sense that $\mathfrak{h}_{i+k} = \mathfrak{h}'_{j+k}$ for all $k \geq 1$ such that these quantities are defined. To see this, observe that the “path” $\mathfrak{h}_0 \bullet \mathfrak{h}_{1 \rightarrow i} \bullet \text{rev}(\mathfrak{h}'_{1 \rightarrow j}) \bullet \text{rev}(\mathfrak{h}'_0)$ separates the map into two disjoint components and that v belongs to only one of them. We use the same vocabulary as in the previous section, left-pinchd or right-pinchd, depending on whether v belongs to the component located to the left or to the right of this path.

Let us suppose that $(q'; v, h, h')$ is left-pinchd. In this case, it is not hard to see that $j \geq 1$. We define $\mathfrak{p} := \mathfrak{h} \bullet \text{rev}(\mathfrak{h}')$ and, as in the previous section, a left part with a self-avoiding path \mathfrak{l} and a right part with a path \mathfrak{r} . The map q is then obtained by gluing \mathfrak{l}_{k+1} to \mathfrak{r}_k for $1 \leq k \leq [\mathfrak{l}] - 1$. The fact that $j \geq 1$ shows that only $(\text{rev}(\mathfrak{r}))_1$ is attached to $(\text{rev}(\mathfrak{r}))_1^-$ in the right part, so that $(\text{rev}(\mathfrak{r}))_1$ creates a corner c' . As \mathfrak{l} is self-avoiding, \mathfrak{l}_1 also defines a corner c . The edge e is the one corresponding to $\mathfrak{r}_{[\mathfrak{h}]}$.

The right-pinchd case is treated by the same method as in the previous section, by exchanging the roles of h and h' . We define $\Phi_{n,p+1\downarrow}(q'; v, h, h') := (q; e, c, c')$.

4.3. These mappings are reverse one from another

Now that we have described in detail our mappings, we may show the main result of this section:

Theorem 3. *The mappings $\Phi_{n,p\uparrow} : \mathcal{Q}_{n,p}^{\text{ecc}} \rightarrow \mathcal{Q}_{n,p+1}^{\text{vhh}}$ and $\Phi_{n,p+1\downarrow} : \mathcal{Q}_{n,p+1}^{\text{vhh}} \rightarrow \mathcal{Q}_{n,p}^{\text{ecc}}$ are one-to-one and reverse one from another.*

Proof. We will see that simple quadruples of $\mathcal{Q}_{n,p}^{\text{ecc}}$ correspond to simple quadruples of $\mathcal{Q}_{n,p+1}^{\text{vhh}}$ through both mappings, and that the same goes for left-pinchd quadruples and right-pinchd quadruples. We treat these cases separately.

Simple case. Let $(q; e, c, c') \in \mathcal{Q}_{n,p}^{\text{ecc}}$ be a simple quadruple and $(q'; v, h, h') := \Phi_{n,p\uparrow}(q; e, c, c')$. We claim that it is a simple quadruple of $\mathcal{Q}_{n,p+1}^{\text{vhh}}$. The fact that $q' \in \mathcal{Q}_{n,p+1}$ is clear, as $\Phi_{n,p\uparrow}$ only alters the external face by adding the two half-edges h and h' . For now, we use the notation of Section 4.1. We claim that, in q' , the path $\mathfrak{l}_{1 \rightarrow [c]+1}$ is the left-most geodesic from the last corner before h in the contour toward v . To show this, we use an argument that we will often use throughout this paper. We argue by contradiction and suppose that this does not hold. Let \mathfrak{h} (with possibly $[\mathfrak{h}] < [c] + 1$) be the left-most geodesic in question. We will see that \mathfrak{h} has to “wind” an infinite number of times around v . Let $k \leq [c] + 1$ be the smallest integer such that $\mathfrak{h}_k \neq \mathfrak{l}_k$. Let us first suppose that \mathfrak{h}_k belongs to the left part. The path \mathfrak{h} cannot stay in the left part as otherwise it would hit the set $\{\mathfrak{l}_1^+, \dots, \mathfrak{l}_{[c]+1}^+\}$ (to which v belongs) and, in q , this would either provide a better alternative to c or contradict the fact that \vec{e} is directed toward c^+ (this case could happen if \mathfrak{h} ends with $\text{rev}(\mathfrak{l}_{[c]+2})$). Let $l > k$ be the smallest integer such that \mathfrak{h}_l does not lie in the left part. By the preceding argument, we see that $\mathfrak{h}_l^- \in \{\mathfrak{l}_{[c]+2}^+, \dots, \mathfrak{l}_0^+\}$. A symmetric argument shows that \mathfrak{h} has to leave the right part: we let m be the smallest integer greater than l such that \mathfrak{h}_m does not lie in the right part. Then, we have $\mathfrak{h}_m^- \in \{\mathfrak{l}_k^+, \dots, \mathfrak{l}_{[c]}^+\}$. As a result, $\mathfrak{h}_{k \rightarrow m-1}$ creates a wind around v and we are back to the same situation as before with \mathfrak{h}_m instead of \mathfrak{h}_k . Reiterating the argument, we obtain that \mathfrak{h} has to infinitely wind around v , which is a contradiction. By similar arguments, we obtain that \mathfrak{h}_k does not lie in the right part either and our claim follows. See Fig. 8.




Fig. 8. The first wind of h around v .

By symmetry, the path $\text{rev}(\tau)_{1 \rightarrow [c'] + 1}$ is the left-most geodesic from the last corner before h' to v . In particular, h and h' are directed toward v , so that $(q'; v, h, h') \in \mathcal{Q}_{n,p+1}^{\text{vhh}}$. With the notation of Section 4.2, we see that these two paths are h and h' , so that the notation \mathbb{I} and τ from both sections is coherent. It becomes clear that $(q'; v, h, h')$ is simple and that $\Phi_{n,p+1 \downarrow}(q'; v, h, h') = (q; e, c, c')$.

Conversely, let $(q'; v, h, h') \in \mathcal{Q}_{n,p+1}^{\text{vhh}}$ be simple, and $(q; e, c, c') := \Phi_{n,p+1 \downarrow}(q'; v, h, h')$. A similar method shows that, with the notation of Section 4.2, the path $\text{rev}(\mathbb{I}_{2 \rightarrow [h]})$ is the right-most geodesic in q from \vec{e} to c^+ . Using a symmetry argument, this is sufficient to conclude that $(q; e, c, c')$ is simple and that $(q'; v, h, h') = \Phi_{n,p \uparrow}(q; e, c, c')$.

Pinched case. As we pass from a left-pinched quadruple to a right-pinched quadruple by exchanging the last two coordinates in both $\mathcal{Q}_{n,p}^{\text{ecc}}$ and $\mathcal{Q}_{n,p+1}^{\text{vhh}}$, we may restrict our attention to the left-pinched case. Let $(q; e, c, c') \in \mathcal{Q}_{n,p}^{\text{ecc}}$ be left-pinched and $(q'; v, h, h') := \Phi_{n,p \uparrow}(q; e, c, c')$. Let h_0 and h'_0 be the corners before h and h' in the contour of the external face. To conclude that $(q'; v, h, h')$ belongs to $\mathcal{Q}_{n,p+1}^{\text{vhh}}$, is left-pinched and that $\Phi_{n,p+1 \downarrow}(q'; v, h, h') = (q; e, c, c')$, it is sufficient to show that, with the notation of Section 4.1, $\mathbb{I}_{1 \rightarrow [c] + 1}$ and $\text{rev}(\tau)_{1 \rightarrow [c'] + 2}$ are the left-most geodesics from h_0 and h'_0 to v in q' .

We use the same winding argument as in the simple case but some extra care is needed in this case. See Fig. 9. Let h be the left-most geodesic from h_0 to v . We argue by contradiction and suppose that $h \neq \mathbb{I}_{1 \rightarrow [c] + 1}$. Let $k \leq [c] + 1$ be the smallest integer such that $h_k \neq \mathbb{I}_k$. Let us first suppose that h_k belongs to the left part. There are several different ways in which h may hit \mathbb{I} after time k : in q' , let us define

$$i := \inf \{ r : \mathbb{I}_r^+ \in \{ \mathbb{I}_1^+, \dots, \mathbb{I}_{r-1}^+, \mathbb{I}_{r+1}^+, \dots, \mathbb{I}_{[0]}^+ \} \} \quad \text{and}$$

$$j := \sup \{ r : \mathbb{I}_r^+ \in \{ \mathbb{I}_1^+, \dots, \mathbb{I}_{r-1}^+, \mathbb{I}_{r+1}^+, \dots, \mathbb{I}_{[0]}^+ \} \}.$$

These are the smallest and largest integers such that $\mathbb{I}_i^+ = \mathbb{I}_j^+$ in q' ; they separate the path \mathbb{I} in four parts, $\mathbb{I}_{1 \rightarrow i}$, $\mathbb{I}_{i+1 \rightarrow [c] + 1}$, $\mathbb{I}_{[c] + 2 \rightarrow j}$ and $\mathbb{I}_{j+1 \rightarrow [0]}$, the second and third being reverse one from another. Beware that in q , this is no longer true. We consider the smallest $l \geq k$ such that either h_l does not lie in the left part or $h_l^+ \in \{ \mathbb{I}_1^+, \dots, \mathbb{I}_j^+ \}$. Such an integer exists as v lies in the latter set. We claim that h_l does not lie in the left part. Indeed, let us argue by contradiction and suppose that $h_{k \rightarrow l}$ entirely lies in the left part. Let s be such that $h_l^+ = \mathbb{I}_s^+$ and $\mathbb{I}_s, h_l, \mathbb{I}_{s+1}$ are arranged according to the clockwise order around h_l^+ . By the following arguments, we obtain a contradiction.

- ◊ If $s \leq [c] + 1$, in q , the path $\text{rev}(h_{k \rightarrow l})$ is to the right of $\text{rev}(\mathbb{I}_{k \rightarrow s})$ and is of the same length or shorter, so that c is not the right-most geodesic from \vec{e} to c^+ .
- ◊ If $s = [c] + 2$, then $[h_{k \rightarrow l}] \leq [k \rightarrow s] - 2$. As a result, \vec{e} is not directed toward c^+ .
- ◊ If $[c] + 3 \leq s \leq j$, $h_{k \rightarrow l}$ creates a shortcut in q and c is not a geodesic.




Fig. 9. The winding argument in the pinched case.




Fig. 10. The two different ways of distinguishing twice the same edge.

As a result, we obtain that h_l does not belong to the left part and that $h_l^- \in \{h_{j+1}^+, \dots, h_{[q]}^+\}$. Let m be the smallest integer greater than l such that either h_m does not lie in the right part or $h_m^+ \in \{h_j^+, \dots, h_{[q]}^+\}$. It is easy to see that h_m does not lie in the right part, otherwise this would contradict the fact that, in q , c' is the right-most geodesic from \bar{e} to c'^+ . We then conclude as in the simple case that h has to infinitely wind around v , which is a contradiction.

By similar arguments, we obtain that h_k does not lie in the right part either, so that $h_{1 \rightarrow [c]+1}$ is the left-most geodesic from h_0 to v in q' . We also obtain that $\text{rev}(\tau)_{1 \rightarrow [c'+2]}$ is the left-most geodesic from h'_0 to v . This allows us to conclude that $\Phi_{n,p \uparrow}$ maps left-pinned quadruples of $\mathcal{Q}_{n,p}^{\text{ecc}}$ to left-pinned quadruples of $\mathcal{Q}_{n,p+1}^{\text{vhh}}$ and that $\Phi_{n,p+1 \downarrow} \circ \Phi_{n,p \uparrow}$ is the identity on the set of left-pinned quadruples of $\mathcal{Q}_{n,p}^{\text{ecc}}$. A very similar technique also shows that $\Phi_{n,p+1 \downarrow}$ maps left-pinned quadruples of $\mathcal{Q}_{n,p+1}^{\text{vhh}}$ to left-pinned quadruples of $\mathcal{Q}_{n,p}^{\text{ecc}}$ and that $\Phi_{n,p \uparrow} \circ \Phi_{n,p+1 \downarrow}$ is the identity on the set of left-pinned quadruples of $\mathcal{Q}_{n,p+1}^{\text{vhh}}$. We leave the details to the reader. \square

5. Changing the number of faces in a quadrangulation

We now present the second main bijection of this work. When distinguishing two edges, we use a convention similar to the one we used for corners. The second time we distinguish an edge, we have to specify on which side it is distinguished (see Fig. 10). On the one hand, we consider the set $\mathcal{Q}_{n,p}^{\text{eem}}$ of quadruples $(q; e, e', m)$ where $q \in \mathcal{Q}_{n,p}$ and

- ◊ e is a distinguished edge of q ,
- ◊ e' is a second distinguished edge of q with the convention mentioned above,
- ◊ m is a mark in $\{c, s_1, s_2\}$.

On the other hand, the set $\mathcal{Q}_{n+1,p}^{fv}$ is the set of triples $(q'; f, v)$ where $q' \in \mathcal{Q}_{n+1,p}$ and

- ◊ f is a distinguished face of q' ,
- ◊ v is a distinguished vertex of q' .

Once we will have presented our bijection between these two sets, we will obtain a proof and combinatorial interpretation of (3).

5.1. Removing a face

It seems more pedagogical to start with the mapping $\Phi_{n+1,p}$ from $\mathcal{Q}_{n+1,p}^{fv}$ to $\mathcal{Q}_{n,p}^{em}$. Let $(q'; f, v) \in \mathcal{Q}_{n+1,p}^{fv}$. Using the face f and the vertex v , we will define a path and shift a part of the map along it in order to suppress f . We will consider two cases: as there are no cycles of odd length in q' , the distances between the ends of the four corners of f and v may only be of two types, either $d, d + 1, d, d + 1$ or $d, d + 1, d + 2, d + 1$, for some $d \geq 0$. In the first case, we will say that f is *confluent* with respect to v and, in the second case, we will say that it is *simple*. This terminology is borrowed from [8].

5.1.1. Confluent face

We suppose here that f is confluent with respect to v and we denote by d the distance between f and v (as in the previous definition). Let h_1, h_2, h_3, h_4 be the four half-edges incident to f read in the counterclockwise order, with h_2 and h_4 directed toward v (we arbitrarily choose among the two possibilities). We denote by \mathbb{p} and \mathbb{p}' the left-most geodesics from h_2 and h_4 to v . It is easy to see that \mathbb{p} and \mathbb{p}' merge as soon as they meet. Note that \mathbb{p} and \mathbb{p}' may even be equal if $h_2^+ = h_4^+$.

Let us first suppose that \mathbb{p} and \mathbb{p}' only meet when reaching v . The paths $\text{rev}(h_3) \bullet \mathbb{p} \bullet \text{rev}(\mathbb{p}') \bullet \text{rev}(h_4)$ and $h_1 \bullet h_2 \bullet \mathbb{p} \bullet \text{rev}(\mathbb{p}')$ are both simple loops, in the sense that they do not go twice through the same vertex¹; as a result, they separate the map in two parts. We call left part the loop $\text{rev}(h_3) \bullet \mathbb{p} \bullet \text{rev}(\mathbb{p}') \bullet \text{rev}(h_4)$ together with all the elements of the map located to the left of this loop. In the case where $h_3 = \text{rev}(h_4)$, the path $\text{rev}(h_3) \bullet \text{rev}(h_4)$ is a “flat” simple loop and the left part is defined as the map consisting in a single edge joining two vertices. In this part, we denote by \mathbb{l} the loop. The right part is the loop $h_1 \bullet h_2 \bullet \mathbb{p} \bullet \text{rev}(\mathbb{p}')$ together with all the elements of the map located to its right. We use the same convention as above if $h_1 = \text{rev}(h_2)$. In the right part, the loop is denoted by \mathbb{r} . Note that f belongs to neither parts.

The map q is the map obtained by gluing back the two parts while matching \mathbb{l}_k with \mathbb{r}_k , for all $1 \leq k \leq [\mathbb{l}] = 2d + 2$. The edge e is the one corresponding to \mathbb{l}_{d+2} and e' is the one corresponding to \mathbb{l}_1 . Remark that when $h_2^+ = h_4^+$, e and e' have the same extremities. When, moreover, two half-edges incident to f are reverse one from another, namely $h_1 = \text{rev}(h_2)$ or $h_3 = \text{rev}(h_4)$, e and e' correspond to the same edge. In this case, we distinguish e and e' in the obvious way: when, for example, $h_1 = \text{rev}(h_2)$, the operation merely consists in suppressing the face f by gluing h_3 and h_4 together, the edge e and e' being distinguished in such a way that e is on the side where h_4 was and e' is on the side where h_3 was. See Fig. 11.

Remark that choosing the other possibility for the half-edges h_1, h_2, h_3, h_4 merely changes the orientation of the loops considered and exchanges the left part with the right part. The final output, $(q; e, e')$, remains the same. To convince oneself that this is true, one can think of the following similar situation: we cut a sphere in two halves, turn clockwise of some fixed angle one of the two

¹ Note that this remains true if $h_2^+ = h_4^+$. In this case, as we supposed that \mathbb{p} and \mathbb{p}' do not meet before reaching v , we must have $v = h_2^+$.




Fig. 11. Removing a confluent face. On the first line, $d = 2$. The special cases where $h_2^+ = h_4^+$ and where two half-edges incident to f are reverse one from another are illustrated on the second line. Beware that, when $h_2^+ = h_4^+$ but neither $h_1 = \text{rev}(h_2)$ nor $h_3 = \text{rev}(h_4)$, the two parts separated by f are rotated before being glued back together.




Fig. 12. Removing a confluent face when the loops are pinched. Here, $d = 3$.

hemispheres and glue the two parts together. The result is the same no matter what hemisphere we chose to rotate.

Let us suppose now that p and p' meet before reaching v . A path will be called a *pinched loop* if it is of the form $a \bullet b \bullet \text{rev}(b) \bullet c$, where b is a self-avoiding path and $a \bullet c$ is a simple loop that intersects b only at its origin. The paths $\text{rev}(h_3) \bullet p \bullet \text{rev}(p') \bullet \text{rev}(h_4)$ and $h_1 \bullet h_2 \bullet p \bullet \text{rev}(p')$ are now both pinched loops. We use the same conventions as above to define a left part and a right part, to glue them together, and to define e and e' . The only difference is that we cut along the pinched part of the loop in the one of the two parts where it is possible, similarly as in Section 4. See Fig. 12.

In both cases, we set $m := c$ (c as confluent) and $\Phi_{n+1 \downarrow, p}(q'; f, v) := (q; e, e', m)$.

5.1.2. Simple face

Let us now turn to the case where f is simple with respect to v . The operation in this case is a little easier as it involves only one geodesic. We still let d be the distance between f and v , and we let h_1, h_2, h_3, h_4 be the four half-edges incident to f read in the counterclockwise order, h_4^+ being the closest to v . We denote by p the left-most geodesic from h_4 to v .

We cut along p , starting from the corner delimited by h_4 and $\text{rev}(h_1)$ and stopping at v . We obtain a map with a face of degree $2d + 4$. In this map, every half-edge of p has two images, exactly one being incident to the face of degree $2d + 4$. For $1 \leq k \leq d$, let p_k^r be the image of p_k that is incident to the face of degree $2d + 4$, and p_k^l the other one. We set

$$\mathbb{l} := \text{rev}(h_1) \bullet p^l \bullet \text{rev}(p_d^r) \quad \text{and} \quad \mathbb{r} := h_2 \bullet h_3 \bullet h_4 \bullet p_{1 \rightarrow d-1}^r.$$

The map q is then defined by gluing back \mathbb{l}_k to \mathbb{r}_k for $1 \leq k \leq d + 2$. The edge e is the one corresponding to \mathbb{l}_{d+2} and e' the one corresponding to \mathbb{l}_1 . See Fig. 13.




Fig. 13. Removing a simple face. On this picture, $d = 3$ and $m := s_1$.




Fig. 14. Setting the value of the mark m . Beware that, on the first line, the external face is not necessarily in the infinite component of the plane whereas, on the second line, it is. The half-edges that will make up e and e' are highlighted.

The difficulty in this case is to set the value of the mark m . To this end, we need to consider the left-most geodesic p' from h_1 to v and see how it merges with p (see Fig. 14). Let $k \in \{0, \dots, d+1\}$ be the smallest integer such that $p_k'^+ \in \{p_0^-, \dots, p_d^-, p_d^+\}$ (with the convention $p_0^- := h_4$ and $p_0^+ := h_1$). After cutting along p ,

- ◊ if $p_k'^+ \in \{p_1^-, \dots, p_d^-, p_d^+\}$, we set $m := s_1$;
- ◊ if $p_k'^+ \in \{h_4^-, p_1^-, \dots, p_{d-1}^-\}$, we set $m := s_2$;
- ◊ if $p_k'^+ = p_d^-$ and the external face of q' is to the right of the cycle $p_{1 \rightarrow d} \bullet \text{rev}(p_{1 \rightarrow d-1}) \bullet h_1$, then $m := s_1$;
- ◊ if $p_k'^+ = p_d^-$ and the external face of q' is to the left of $p_{1 \rightarrow d} \bullet \text{rev}(p_{1 \rightarrow d-1}) \bullet h_1$, then $m := s_2$.

The reason why we choose this rule may seem surprising at this point but should become clear in the next sections. As a first insight, it can be noticed at this stage that h_1 and h_2 will make up e' and that p_{d-1}^r and p_d^r will make up e . Here again, we set $\Phi_{n+1, p}(q'; f, v) := (q; e, e', m)$.

5.2. Adding a face

We now present the mapping $\Phi_{n \uparrow, p}$ from $\mathcal{Q}_{n, p}^{\text{em}}$ to $\mathcal{Q}_{n+1, p}^{\text{fv}}$. Let $(q; e, e', m) \in \mathcal{Q}_{n, p}^{\text{em}}$. Notice that the four distances between the extremities of e and the extremities of e' are either $d, d+1, d+1, d$ or $d, d+1, d+1, d+2$, for some $d \geq 0$. In the first case, we will say that e and e' are *parallel*. See Fig. 15.

5.2.1. Confluent case

Let us start with the construction in the case $m = c$. In a first time, we moreover suppose that e and e' are parallel and we let d be the distance between them. We arbitrarily choose a half-edge \bar{e} corresponding to e and we consider the right-most geodesics p and p' from \bar{e} and from $\text{rev}(\bar{e})$ toward e' .




Fig. 15. Parallel edges on the left, nonparallel edges on the right.




Fig. 16. Choosing the path \mathbb{P} when e and e' are nonparallel (top) or parallel (bottom) and when $m = s_1$ (left) or $m = s_2$ (right).

As there are no cycles of odd length, it is easy to see that these paths do not intersect and, in particular, their endpoints are distinct: $\mathbb{P}_d^+ \neq \mathbb{P}'_d^+$. Let \vec{e}' denote the half-edge corresponding to e' , directed from \mathbb{P}_d^+ to \mathbb{P}_d^+ . The loop

$$\vec{e}' \bullet \text{rev}(\mathbb{P}) \bullet \text{rev}(\vec{e}) \bullet \mathbb{P}'$$

separates the map in a left and a right part as before, both of them having a face of degree $2d + 2$. (If e and e' correspond to the same edge, one part will consist in a single edge and the other one will have a face of degree 2 where e and e' were.) We denote by \mathbb{L} and \mathbb{R} the images of the loop in the left and in the right part. We define q' by gluing the parts together while matching \mathbb{L}_k with \mathbb{R}_{k+1} for $2 \leq k \leq 2d + 1$. This creates an extra face f incident to $\mathbb{R}_1, \mathbb{R}_2, \text{rev}(\mathbb{L}_1), \text{rev}(\mathbb{L}_{2d+2})$. We set $v := \mathbb{R}_{d+2}^+$ in the map q' and $\Phi_{n \uparrow, p}(q; e, e', m) := (q'; f, v)$. See Fig. 11.

Now, if e and e' are nonparallel, we set \vec{e} the half-edge corresponding to e directed toward e' , as well as \vec{e}' the half-edge corresponding to e' directed away from e . We also let $d - 1$ be the distance between e and e' . We consider the right-most geodesic \mathbb{P} from \vec{e} toward \vec{e}'^+ , and the right-most geodesic \mathbb{P}' from \vec{e} toward \vec{e}'^- . The path

$$\vec{e}' \bullet \text{rev}(\mathbb{P}) \bullet \text{rev}(\vec{e}) \bullet \vec{e} \bullet \mathbb{P}'$$

is a pinched loop of length $2d + 2$. (Note that, as soon as \mathbb{P} and \mathbb{P}' split, they cannot meet again.) As a result, it separates the map in a left and a right part as before and we use the same convention to define q' and f . The vertex v is defined as the endpoint of the pinched part in the part where the loop is pinched. We set $\Phi_{n \uparrow, p}(q; e, e', m) := (q'; f, v)$.

5.2.2. Simple case

We now suppose that $m \in \{s_1, s_2\}$. We will always proceed as follows: we will find a path \mathbb{P} linking e to e' , cut along it and create a new face by sliding the two sides of the cut as in Fig. 13. Let us first describe how to choose the path \mathbb{P} along which we will cut. See Fig. 16.

We first suppose that e and e' are nonparallel. Let \vec{e} be the half-edge corresponding to e directed toward e' . If $m = s_1$, then \mathbb{P} is the right-most geodesic from \vec{e} to e' . If $m = s_2$, then \mathbb{P} is the right-most geodesic from \vec{e} to the extremity of e' the farther away from e .

Now, if e and e' are parallel, we consider the two right-most geodesics from both half-edges corresponding to e toward e' . As there are no cycles of odd length, these two geodesics do not meet. As a

result, it is possible to concatenate a half-edge corresponding to e , one of these geodesics, a half-edge corresponding to e' and the reverse of the other geodesic so that the result is a simple loop. This may be done in two different ways, depending on which geodesic we choose to visit first, and choosing one way or the other merely changes the orientation of the loop. If $m = s_1$ (resp. $m = s_2$), we denote by \mathbb{p} the geodesic such that if we visit it first, then the external face of the map lies to the right (resp. to the left) of this loop.

In all cases, we denote by \vec{e} the half-edge corresponding to e directed toward the beginning of \mathbb{p} and \vec{e}' the half-edge corresponding to e' directed away from the endpoint of \mathbb{p} . (We use an obvious extension in the case where \mathbb{p} is the empty path.) We also set $d := [\mathbb{p}]$. By cutting along

$$\text{rev}(\vec{e} \bullet \mathbb{p} \bullet \vec{e}'),$$

we obtain a map with a face of degree $2d + 4$. We denote by \mathbb{r} the image of the previous path that is incident to this face, and \mathbb{l} the image that is not incident to this face. Beware that in the case where e and e' are nonparallel and $m = s_2$, one of these paths is not self-avoiding if $\mathbb{p}_d^- = \vec{e}'^+$. In this case, the part of the map separated by the loop $\vec{e}' \bullet \mathbb{p}_d$ that does not contain e remains attached to \vec{e}'^+ . We then define q' by gluing back \mathbb{l}_k with \mathbb{r}_{k+2} for $2 \leq k \leq d$, as well as \mathbb{l}_{d+1} with $\text{rev}(\mathbb{l}_{d+2})$. This creates an extra face f incident to $\mathbb{r}_1, \mathbb{r}_2, \mathbb{r}_3$ and $\text{rev}(\mathbb{l}_1)$. We set $v := \mathbb{l}_{d+1}^+$ in the map q' and $(q'; f, v) := \Phi_{n \uparrow, p}(q; e, e', m)$. See Fig. 13.

5.3. These mappings are reverse one from another

Theorem 4. *The mappings $\Phi_{n \uparrow, p} : \mathcal{Q}_{n, p}^{\text{em}} \rightarrow \mathcal{Q}_{n+1, p}^{\text{fv}}$ and $\Phi_{n+1 \downarrow, p} : \mathcal{Q}_{n+1, p}^{\text{fv}} \rightarrow \mathcal{Q}_{n, p}^{\text{em}}$ are one-to-one and reverse one from another.*

Proof. We will proceed in a similar manner as in the proof of Theorem 3. Our constructions from Sections 5.1 and 5.2 are clearly reverse one from another, the only thing we have to check is that the paths along which we cut correspond. In other words, we need to verify that the path denoted by \mathbb{l} in both sections is the same if we apply our constructions to a map and to its image through one of our mappings, and that the same goes for the path \mathbb{r} .

Confluent case. Let $(q'; f, v) \in \mathcal{Q}_{n+1, p}^{\text{fv}}$ be such that f is confluent with respect to v and let $(q; e, e', m) := \Phi_{n+1 \downarrow, p}(q'; f, v)$. We use here the notation of Section 5.1. Let us first suppose that the loop considered in Section 5.1.1 is not pinched (case of Fig. 11). In order to conclude that $\Phi_{n \uparrow, p}(q; e, e', m) = (q'; f, v)$, it will be sufficient to show that $\text{rev}(\mathbb{r}_{2 \rightarrow d+1})$ and $\mathbb{r}_{d+3 \rightarrow 2d+2}$ are the two right-most geodesics from the two half-edges corresponding to e toward e' . Note that, in particular, this will entail that e and e' are parallel. This fact is again shown by the winding argument: if, for example, the right-most geodesic from $\text{rev}(\mathbb{r}_{d+2})$ to e' leaves $\text{rev}(\mathbb{r}_{2 \rightarrow d+1})$ and enters the left part or the right part, then it has to leave it or it would contradict either the fact that \mathbb{p} and \mathbb{p}' are left-most geodesics from f to v or the fact that h_1^+ and h_3^+ are at distance $d + 1$ from v . Repeating the argument, it will indefinitely wind around e' .

If the loop of Section 5.1.1 is pinched (case of Fig. 12), we conclude by exactly the same argument. Note that in this case, e and e' are nonparallel.

Simple case. Let us suppose now that $(q'; f, v) \in \mathcal{Q}_{n+1, p}^{\text{fv}}$ is such that f is simple with respect to v and let again $(q; e, e', m) := \Phi_{n+1 \downarrow, p}(q'; f, v)$. We use the notation of Section 5.1.2 (see Fig. 17). In q , let us set $\vec{e} := \text{rev}(\mathbb{l}_{d+2})$ and $\vec{e}' := \text{rev}(\mathbb{l}_1)$. It is not hard (although it requires some care) to show that $\text{rev}(\mathbb{l}_{2 \rightarrow d+1})$ is the right-most geodesic from \vec{e} to \vec{e}'^- . This entails in particular that $d_q(\vec{e}^+, \vec{e}'^-) = d$. It is also easy to see that $d_q(\vec{e}^-, \vec{e}'^+) = d + 1$.

The remaining pair of distances

$$(d_q(\vec{e}^+, \vec{e}'^+), d_q(\vec{e}^-, \vec{e}'^+))$$

may be either $(d + 1, d + 2)$, $(d - 1, d)$ or $(d + 1, d)$, depending on how \mathbb{p}' merges with \mathbb{p} . The first two cases correspond to the top line of Fig. 14. In these cases, e and e' are nonparallel and \vec{e} is directed




Fig. 17. Reminder of the construction $\Phi_{n+1\downarrow,p}$ in the case of a simple face and some notation.

toward e' . As a result, the path $\text{rev}(\mathbb{I}_{2 \rightarrow d+1})$ corresponds to the path of Section 5.2.2. The last case corresponds to the bottom line of Fig. 14; we obtain that e and e' are parallel, and distinguishing whether $m = s_1$ or $m = s_2$, we see that $\text{rev}(\mathbb{I}_{2 \rightarrow d+1})$ still corresponds to the path of Section 5.2.2. This is sufficient to conclude that $\Phi_{n\uparrow,p}(q; e, e', m) = (q'; f, v)$.

The fact that $\Phi_{n+1\downarrow,p} \circ \Phi_{n\uparrow,p}$ is the identity over $\mathcal{Q}_{n,p}^{\text{sem}}$ also follows from the same kind of arguments and is left to the reader. \square

5.4. Interpretation through the Bouttier–Di Francesco–Guitter bijection

The Bouttier–Di Francesco–Guitter bijection is a classical bijection allowing to encode maps with simpler objects. In our case of quadrangulations, it specializes into a bijection between quadrangulations having a distinguished vertex and so-called well-labeled forests. We will not describe this bijection in this work and refer to [3] for a complete description. See also [2,4] for an exposition in the particular case of quadrangulations.

Definition 3. A well-labeled forest is a pair (f, l) where $f = (t_1, \dots, t_p)$ is a forest and l is an integer-valued function on the vertices of f satisfying the following:

- ◊ $|l(u) - l(v)| \leq 1$ whenever u and v are vertices of the same tree linked by an edge,
- ◊ $l(\rho_{i+1}) \geq l(\rho_i) - 1$ for all $1 \leq i \leq p$, where ρ_i denotes the root vertex of t_i , and $\rho_{p+1} := \rho_1$,
- ◊ $l(\rho_1) = 0$.

Our quadrangulations in $\mathcal{Q}_{n+1,p}^{\text{fv}}$ come with a distinguished vertex, so it seems natural to try to understand what happens to the coding well-labeled forest through $\Phi_{n+1\downarrow,p}$. However, after applying our bijection, the forest is no longer a forest, it becomes some map with two faces with rather complicated rules on its labels. We may also mention that the maps in $\mathcal{Q}_{n,p}^{\text{sem}}$ do not come a priori with a natural distinguished vertex.

A quite remarkable fact that follows from our theorems is that a third of the quadrangulations in $\mathcal{Q}_{n+1,p}^{\text{fv}}$ are such that the distinguished face is confluent with respect to the distinguished vertex, one third are organized as on the left of Fig. 14 and one third are organized as on the right of Fig. 14. This fact is actually very easy to show by using the Bouttier–Di Francesco–Guitter bijection.

In fact, quadrangulations in $\mathcal{Q}_{n+1,p}^{\text{fv}}$ correspond to well-labeled forests having a distinguished edge. The first third correspond to the case where the edge links two vertices having the same label. In the other cases, the edge links a vertex v labeled ℓ to a vertex v' labeled $\ell + 1$. Removing the edge breaks one of the trees t in two parts. We consider, on the one hand, the part of t not containing its root and, on the other hand, the part of t containing the root, together with all the other trees. We denote by S_v the one of these two sets containing v and $S_{v'}$ the one containing v' . Finally, we set $m_v := \min_{S_v} l - \ell$ and $m_{v'} := \min_{S_{v'}} l - (\ell + 1)$. Then, the top-left part of Fig. 14 corresponds to the case where $m_{v'} < m_v$, the top-right part corresponds to the case where $m_v < m_{v'}$. The bottom

line corresponds to the case where $m_v = m_{v'}$; on the left, the root of t belongs to S_v , on the right, it belongs to $S_{v'}$. By symmetry, we recover the distribution into three thirds.

References

- [1] Jérémie Bettinelli, Scaling limits for random quadrangulations of positive genus, *Electron. J. Probab.* 15 (52) (2010) 1594–1644.
- [2] Jérémie Bettinelli, Scaling limit of random planar quadrangulations with a boundary, *Ann. Inst. H. Poincaré Probab. Statist.* (2013), in press, <http://imstat.org/aihp/accepted.html>, arXiv:1111.7227.
- [3] Jérémie Bouttier, Philippe Di Francesco, Emmanuel Guitter, Planar maps as labeled mobiles, *Electron. J. Combin.* 11 (1) (2004), Research Paper 69, 27 pp. (electronic).
- [4] Jérémie Bouttier, Emmanuel Guitter, Distance statistics in quadrangulations with a boundary, or with a self-avoiding loop, *J. Phys. A* 42 (46) (2009), 465208, 44 pp.
- [5] Robert Cori, Un code pour les graphes planaires et ses applications, *Astérisque* 27 (1975), with an English abstract.
- [6] J. Pitman, *Combinatorial Stochastic Processes*, Lecture Notes in Math., vol. 1875, Springer-Verlag, Berlin, July 2006, Lectures from the 32nd Summer School on Probability Theory held in Saint-Flour, July 7–24, 2002, with a foreword by Jean Picard.
- [7] Jean-Luc Rémy, Un procédé itératif de dénombrement d'arbres binaires et son application à leur génération aléatoire, *RAIRO Inform. Théor.* 19 (2) (1985) 179–195.
- [8] Gilles Schaeffer, *Conjugaison d'arbres et cartes combinatoires aléatoires*, PhD thesis, Université de Bordeaux 1, 1998.
- [9] William T. Tutte, A census of slicings, *Canad. J. Math.* 14 (1962) 708–722.