Chapter 2

Variants of the Cori–Vauquelin–Schaeffer bijection

As is customary when studying on scaling limits of maps, this work strongly relies on powerful encodings of discrete maps by tree-like objects. We now present variants of the famous Cori–Vauquelin–Schaeffer (CVS) bijection [38,90] between plane quadrangulations and so-called *well-labeled trees*, and its generalizations by Chapuy–Marcus–Schaeffer [35] for higher genera and by Bouttier–Di Francesco–Guitter [28] for plane maps with faces of arbitrary degrees. We only give the constructions from the encoding objects to the considered maps and refer the reader to the aforementioned works for converse constructions and proofs.

2.1 Basic construction

Let $\mathbf{m}$ be a map, rooted or not, and f be a face of $\mathbf{m}$. Starting from a choice of a corner c_0 in f, we index the subsequent corners of f in counterclockwise order as $(c_i, i \in \mathbb{Z})$ (forming a periodic sequence). Let $\lambda\colon V(\mathbf{m}) \to \mathbb{Z}$ be a labeling of the vertices of $\mathbf{m}$ by integers. We extend the definition of λ to the corners of the map by setting $\lambda(c) = \lambda(v)$ if v is the vertex incident to the corner c. In what follows, we will either consider that λ is defined up to addition of a constant, or that the value of λ at some corner is fixed, for instance, that $\lambda(c_0) = 0$.

We say that $(\mathbf{m}, \lambda)$ is *well labeled* inside f if $\lambda(c_{i+1}) \geq \lambda(c_i) - 1$ for every $i \geq 0$. In particular, if $(\mathbf{m}, \lambda)$ is well labeled inside f and e is a half-edge of $\mathbf{m}$ such that both e and its reverse $\overline{e}$ are incident to f, then $|\lambda(e^+) - \lambda(e^-)| \leq 1$, where e^-, e^+ denote the origin and end of e. Note that this will be the case for every edge when $\mathbf{m}$ is a map with a single face.

Let $(\mathbf{m}, \lambda)$ be well labeled inside f. With the above notation, let us define $s(i) = \inf\{j > i : \lambda(c_j) = \lambda(c_i) - 1\} \in \mathbb{Z} \cup \{\infty\}$ and the *successor* of c_i as $s(c_i) = c_{s(i)}$, where c_∞ is by convention the unique corner incident to a vertex v_* that is added in the interior of f, and which naturally carries the label $\lambda(v_*) = \min\{\lambda(c_i), i \in \mathbb{Z}\} - 1$. Clearly, $s(c_i)$ is then well defined for all corners (distinct from c_∞), and only depends on the corner c_i and not on the particular choice of the index i. The *CVS construction inside the face* f consists in

- linking by an arc every corner c incident to f to its successor $s(c)$, in such a way that arcs do not cross, which is always possible due to the well-labeling condition,
- deleting all the edges of $\mathbf{m}$.

This construction results in an embedded graph[1] denoted by $\mathrm{CVS}(\mathbf{m}, \lambda; f)$, whose vertex set consists of v_* and the vertices of $\mathbf{m}$ incident to f, and whose edges are the arcs between the corners of f and their successors. By construction, the edges of $\mathrm{CVS}(\mathbf{m}, \lambda; f)$ are in bijection with the corners of $\mathbf{m}$ incident to f. If $\mathbf{m}$ is rooted inside f, say at the corner c_i, then $\mathrm{CVS}(\mathbf{m}, \lambda; f)$ naturally inherits a root at the corner preceding the arc linking c_i to $s(c_i)$. Note that the well-labeling condition, as well as the output $\mathrm{CVS}(\mathbf{m}, \lambda; f)$, are invariant under addition of a constant to λ, as they should.

We will also need an *interval variant* of this construction, where we fix a sequence, also referred to as an *interval*, of subsequent corners $I = \{c_0, c_1, \dots, c_r\}$ of a face f of $\mathbf{m}$, and only ask that $(\mathbf{m}, \lambda)$ is *well labeled* on I in the sense that $\lambda(c_{i+1}) \geq \lambda(c_i) - 1$ for $0 \leq i \leq r-1$. In this case, we set

$$\lambda_* = \min\{\lambda(c_i), 0 \leq i \leq r\} - 1 \quad \text{and} \quad \ell = \lambda(c_r) - \lambda_*.$$

Instead of a single extra corner c_∞, we introduce inside f a sequence of distinct consecutive corners $c_{r+1}, c_{r+2}, \dots, c_{r+\ell}$, incident to new vertices $v_{r+1}, v_{r+2}, \dots, v_{r+\ell}$ with labels $\lambda(c_r) - 1, \lambda(c_r) - 2, \dots, \lambda_*$. The successor mapping s is then defined for all corners except $c_{r+\ell}$. We let $\mathrm{CVS}(\mathbf{m}, \lambda; I)$ be the resulting (nonrooted) embedded graph whose edges are the arcs. In this embedded graph, the following are of particular interest:

(1) the *apex* $v_{r+\ell}$, which will usually be denoted with a subscript $*$;

(2) the *maximal geodesic*, which is the chain of arcs linking $c_0, s(c_0), s(s(c_0)), \dots, c_{r+\ell}$, and which will always be denoted with the letter γ and depicted in red (darker color) in the figures;

(3) the *shuttle*, which is the chain of arcs linking $c_r, c_{r+1}, \dots, c_{r+\ell}$, and which will always be denoted with the letter ξ and depicted in green (lighter color) in the figures.

Note that the two latter are paths from the first and last corners of I to the apex.

The construction generalizes to several intervals $I, J, \dots$ that pairwise share at most one extremity. In the case of a shared extremity, say $I = \{c_0, c_1, \dots, c_r\}$ and $J = \{c'_0, c'_1, \dots, c'_{r'}\}$ with $c_r = c'_0$, one first duplicates the common corner before applying the construction, in the sense that the copy c_r is used in the shuttle of I and c'_0 is used in the maximal geodesic of J; see Figure 2.1. In this construction, each interval yields a distinct apex, maximal geodesic, and shuttle, and the construction results in an embedded graph denoted by $\mathrm{CVS}(\mathbf{m}, \lambda; I, J, \dots)$. Plainly, the ordering of the intervals does not affect the construction.

[1] In general, this embedded graph is not a map of the surface into consideration. In all the constructions we will use in this work, it will, however, always turn out to be a map.

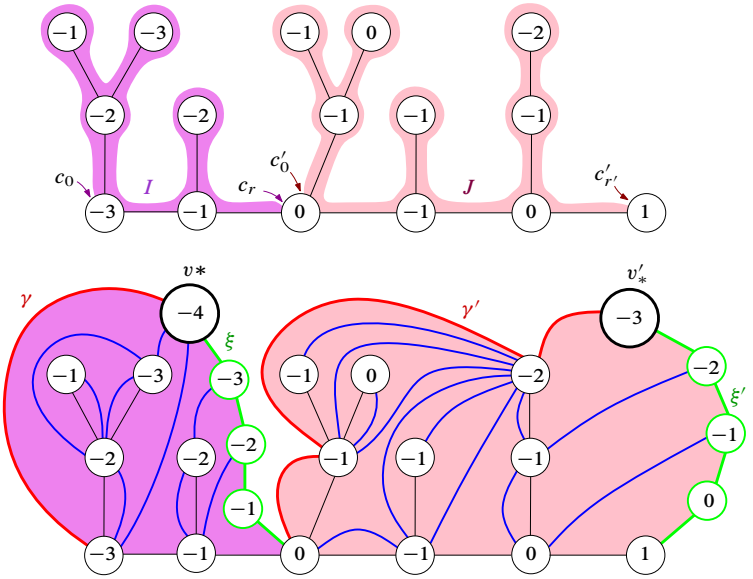

Figure 2.1. Performing the interval variant of the Cori–Vauquelin–Schaeffer bijection with two intervals sharing an extremity. The interval I consists of the corners in the purple (darker) area, starting with c_0 and ending with c_r, while J consists of the corners in the red (lighter) area, starting with c'_0 and ending with $c'_{r'}$. The interval I yields the apex v_*, maximal geodesic γ, and shuttle ξ, while J yields v'_*, γ', and ξ', respectively. As will be the case in all the figures, the maximal geodesics are in red (darker colored boundary) and the shuttles in green (lighter colored boundary).

We make the important observation that any chain $c, s(c), \dots, s^i(c)$ of consecutive successors induces a geodesic chain for the graph metric in the resulting embedded graph $\mathrm{CVS}(\mathbf{m}, \lambda; I, J, \dots)$, that is, a path of minimal length between its extremities. This is simply because, by construction, any arc of the resulting embedded graph links two vertices u and v such that $|\lambda(u) - \lambda(v)| = 1$, and because λ decreases by 1 at every step on a chain of consecutive successors. In particular, the maximal geodesics and shuttles of $\mathrm{CVS}(\mathbf{m}, \lambda; I, J, \dots)$ are geodesic chains.

2.2 The generalized Chapuy–Marcus–Schaeffer bijection

Encoding quadrangulations. As a first example, let us perform this construction on a particular class of maps. For $n \in \mathbb{Z}_{\geq 0}$ and $\boldsymbol{l} = (l^1, \dots, l^k) \in (\mathbb{Z}_{\geq 0})^k$, we let $\vec{\mathbf{M}}_{n,\boldsymbol{l}}^{[g]}$ be the set of labeled rooted maps $(\mathbf{m}, \lambda)$ satisfying the following properties:

- $\mathbf{m}$ is a map of genus g with $n + \sum_{i=1}^{k} l^i$ edges, one internal face f_* and k holes $h_1, \dots, h_k$, rooted at a corner of its internal face f_*,

- for all i, the hole h_i is of degree l^i; if it is an external face, then it has a simple boundary,[2]
- for any $i \neq j$, if h_i and h_j are faces, then they are not incident to any common edge,
- $(\mathbf{m}, \lambda)$ is well labeled inside f_*.

We similarly define the set $\mathbf{M}_{n,\boldsymbol{l}}^{[g]}$ of labeled nonrooted maps. Setting

$$\boldsymbol{l}0 = (l^1, \dots, l^k, 0),$$

the CVS construction applied to the internal face f_* provides a bijection between $\mathbf{M}_{n,\boldsymbol{l}}^{[g]}$ and $\mathbf{Q}_{n,\boldsymbol{l}0}^{[g]}$, through which the k first holes correspond, while the extra hole h_{k+1} of the quadrangulation is the extra vertex v_* of the construction. In case of rooted maps, it yields a one-to-two correspondence[3] between $\vec{\mathbf{M}}_{n,\boldsymbol{l}}^{[g]}$ and $\vec{\mathbf{Q}}_{n,\boldsymbol{l}0}^{[g]}$.

Decomposition into elementary pieces. Let us now perform the construction on the same set of maps $\mathbf{M}_{n,\boldsymbol{l}}^{[g]}$ but with well-chosen intervals. We will decompose a map of $\mathbf{M}_{n,\boldsymbol{l}}^{[g]}$ into a collection of labeled forests indexed by an underlying structure called the *scheme*. For the remainder of this section, we exclude the cases $(g, k) \in \{(0, 0), (0, 1)\}$ leading to encoding objects not entering the upcoming framework. We fix $(\mathbf{m}, \lambda) \in \mathbf{M}_{n,\boldsymbol{l}}^{[g]}$.

Let $\tilde{\mathbf{m}}$ be the nonrooted map obtained from $\mathbf{m}$ by iteratively removing all its vertices of degree 1 that are not holes. The resulting map $\tilde{\mathbf{m}}$ may be seen as a submap of $\mathbf{m}$: the map $\mathbf{m}$ is obtained from $\tilde{\mathbf{m}}$ by appending rooted labeled trees at its corners. We call *nodes* of $\mathbf{m}$ the following vertices:

- the external vertices of $\mathbf{m}$,
- the vertices of $\mathbf{m}$ having degree 3 or more in $\tilde{\mathbf{m}}$.

These nodes are linked in $\tilde{\mathbf{m}}$ by maximal chains of edges not containing any nodes other than their extremities. Replacing every such chain with a single edge yields a nonrooted map $\mathbf{s}$, called the *scheme* of $\mathbf{m}$. It has one internal face, still denoted by f_*, and k holes, still denoted by $h_1, \dots, h_k$; see Figure 2.2.

We denote by $\vec{E}(\mathbf{s})$ the set of half-edges incident to the internal face of $\mathbf{s}$; this set is partitioned into the set $\vec{I}(\mathbf{s})$ of half-edges whose reverses belong to $\vec{E}(\mathbf{s})$ as well, and the set $\vec{B}(\mathbf{s})$ of half-edges whose reverses do not belong to $\vec{E}(\mathbf{s})$. (We used the letter I for *internal* and B for *boundary*.) The set $\vec{B}(\mathbf{s})$ is further partitioned as

$$\vec{B}(\mathbf{s}) = \bigsqcup_{1 \leq r \leq k} \vec{B}_r(\mathbf{s}),$$

[2] A face has a *simple boundary* if it is incident to as many vertices as its degree.

[3] The factor 2 comes from the fact that the corners of f_* correspond to the edges of the resulting map, each edge corresponding to 2 half-edges. We refer the interested reader to [22, Section 3.1] for a presentation of the reverse mapping.

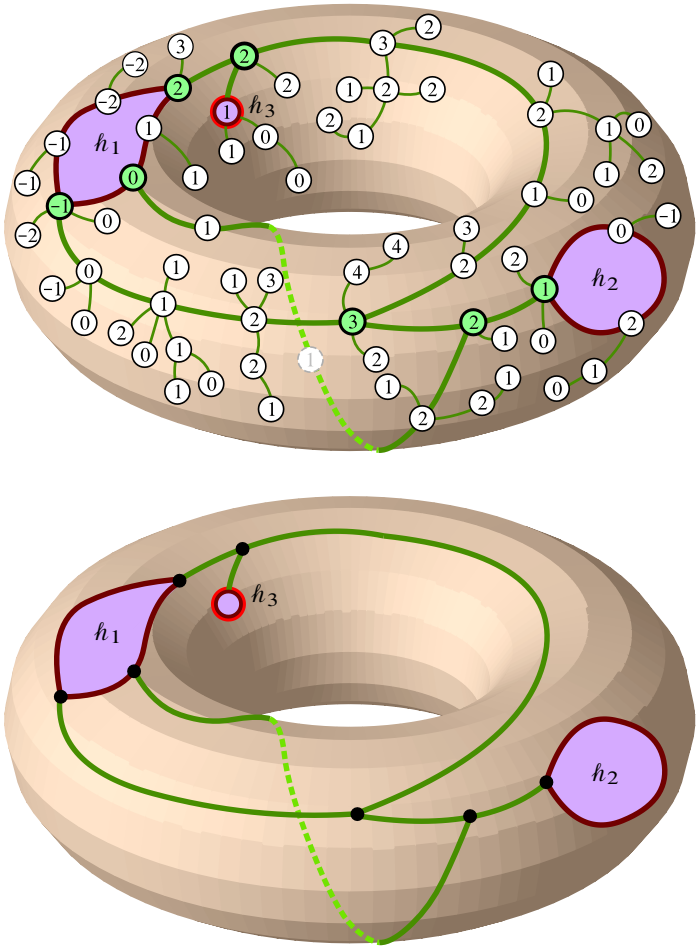

Figure 2.2. *Top:* A labeled map from $\mathbf{M}^{[1]}_{63,(6,3,0)}$. The outlined vertices are its nodes and the thicker edges correspond to the map $\tilde{\mathbf{m}}$. *Bottom:* The corresponding scheme.

where $\vec{B}_r(\mathbf{s})$ is either empty if $\hbar_r$ is a vertex, or the set of half-edges of $\vec{E}(\mathbf{s})$ whose reverses are incident to $\hbar_r$ if it is a face. We consider $e \in \vec{E}(\mathbf{s})$. It corresponds to a chain $e_1, \dots, e_j$ of half-edges in $\mathbf{m}$. Let us denote by c_e and c'_e the corners of $\tilde{\mathbf{m}}$ preceding e_1 and succeeding e_j in the contour order. In $\mathbf{m}$, there are several corners that make up c_e and c'_e. The corner interval I_e is the interval of corners of $\mathbf{m}$ from the *first* corner corresponding to c_e to the *first* corner corresponding to c'_e. Observe that, in $\mathbf{m}$, the tree grafted at c_e is thus covered by I_e, whereas the tree grafted at c'_e is not.

By construction, $\bigcup_{e\in\vec{E}(\mathbf{s})} I_e$ is equal to the set of corners of f_* and each extremity of these intervals is shared by exactly two such intervals. More precisely, the intervals I_e, $e \in \vec{E}(\mathbf{s})$, with their last corner removed give a partition of the corners

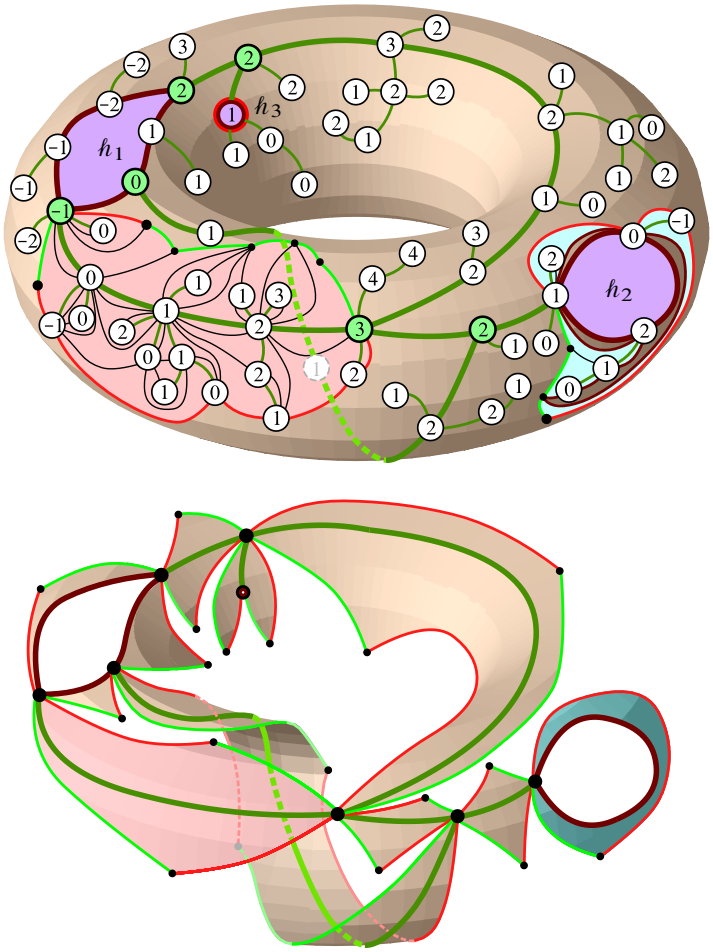

Figure 2.3. Performing the interval bijection on the labeled map from Figure 2.2. *Top.* Two elementary pieces are represented: one quadrilateral with geodesic sides in red (light color, bottom left), and one composite slice in blue (even lighter color, around h_2). *Bottom.* The interval bijection yields a decomposition into 4 composite slices and 7 quadrilateral with geodesic sides. Here, only the maximal geodesics and shuttles are depicted. We let the edges of the original scheme figure on this output map, but these are neither edges nor chains of edges of this output map (remember that the edges of the original map are never edges of the output map).

of f_*. Applying the interval CVS construction $\mathrm{CVS}(\mathbf{m}, \lambda; \{I_e, e \in \vec{E}(\mathbf{s})\})$ gives a natural decomposition of the quadrangulation $(\mathbf{q}, v_*) = \mathrm{CVS}(\mathbf{m}, \lambda; f_*)$ into submaps, whose study, starting in the next section, is the key to this work; see Figure 2.3. These submaps are called the *elementary pieces* of $(\mathbf{q}, v_*)$ and are of two types: the ones corresponding to half-edges of $\vec{B}(\mathbf{s})$ are called *(composite) slices* and the ones corresponding to half-edges of $\vec{I}(\mathbf{s})$ are called *quadrilaterals (with geodesic sides)*. They

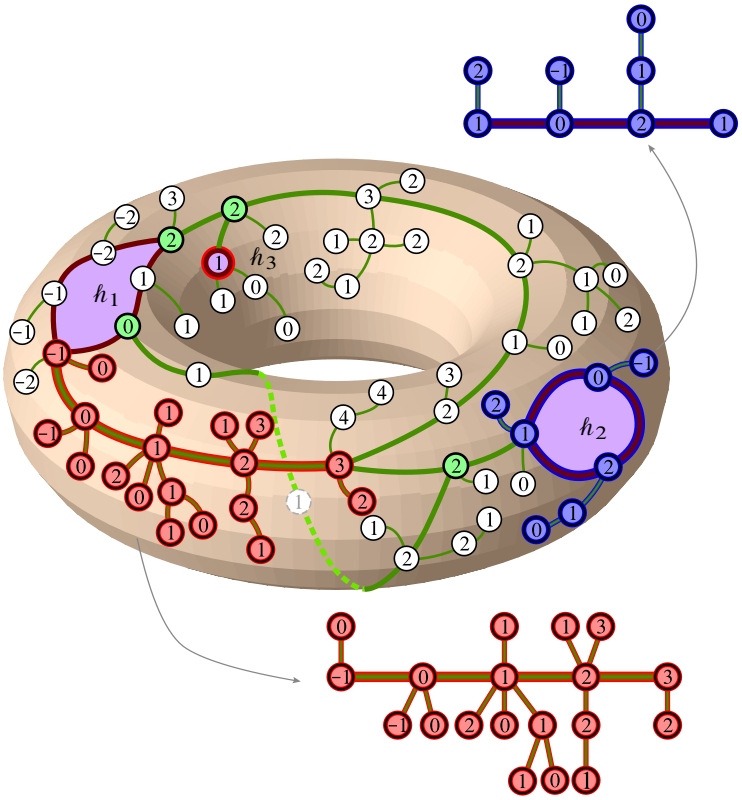

Figure 2.4. The parts of the labeled map from Figure 2.2 encoding the elementary pieces. The two parts (in red and blue) corresponding to the (red) quadrilateral with geodesic sides and the (blue) composite slice depicted in Figure 2.3 are extracted.

are not rooted and come with distinguished vertices on their boundaries that will be discussed later on.

The elementary piece corresponding to the half-edge $e \in \vec{E}(\mathbf{s})$ is encoded by the part of the labeled map $(\mathbf{m}, \lambda)$ corresponding to

- either the interval I_e if $e \in \vec{B}(\mathbf{s})$,
- or the union $I_e \cup I_{\overline{e}}$ if $e \in \vec{I}(\mathbf{s})$, where $\overline{e}$ denotes the reverse of e.

These encoding parts are depicted in Figure 2.4. Note that, when $e \in \vec{I}(\mathbf{s})$, the elementary pieces corresponding to e and to its reverse $\overline{e}$ are the same map; only the distinguished vertices on the boundary will differ (more precisely be given in a different order). We refer the reader to [22, Section 3.4.1] for more on this decomposition, keeping in mind that, in the latter reference, the maps are rooted and the root is encoded in the scheme, which essentially amounts in seeing the root of the map as an extra external vertex.

Finiteness of the number of schemes. We will elaborate more on elementary pieces in the next two sections and end this one with a simple combinatorial lemma. We say that a map with holes is a *scheme* if it has one internal face, all its external faces have a simple boundary and do not share a common incident edge, and all its internal vertices have degree 3 or more.

Lemma 2.1. *For fixed values of* $(g, k) \notin \{(0, 0), (0, 1)\}$*, there are finitely many genus* g *schemes with* k *holes and these have at most* $3(2g + k - 1)$ *edges and* $2(2g + k - 1)$ *vertices.*

Proof. As there is a finite number of maps with a given number of edges, the bound on the number of edges yields the finiteness of the considered set.

Let v, e, f be the number of vertices, edges and faces of a given scheme as in the statement, and let b be its number of external faces (so that $p = k - b$ are external vertices). By construction, we have $f = b + 1$ and the vertices are all of degree at least 3, except possibly up to p of them, which have degree at least 1. The sum of the degrees of the vertices being twice the number of edges, we obtain $2e \geq 3(v - p) + p$, and we see by the Euler characteristic formula $v - e + f = 2 - 2g$ that the considered scheme has at most $6g + 3k - p - 3 \leq 3(2g + k - 1)$ edges, and at most $2(2g + k - 1)$ vertices, using that $k = p + b$. ∎

2.3 Composite slices

We call *plane forest* a collection $\mathbf{f} = (\mathbf{t}^0, \dots, \mathbf{t}^{l-1}, \rho^l)$, for some $l \geq 1$, of rooted plane trees (the last one being reduced to the vertex-tree), which we view systematically as a map by taking an embedding of every $\mathbf{t}^i$ in the upper half-plane $\mathbb{R} \times \mathbb{R}_{\geq 0}$, with root ρ^i at the point $(i, 0)$, and in which ρ^i is linked to ρ^{i-1} by the line segment between $(i, 0)$ and $(i - 1, 0)$ for $1 \leq i \leq l$. The union of these line segments is called the *floor* of the forest. The resulting embedded graph, which we still denote by $\mathbf{f}$ (see, e.g., the left of Figure 2.5), is a nonrooted plane map coming with the two distinguished vertices $\rho = \rho^0$ and $\overline{\rho} = \rho^l$; it is in fact a plane tree, but we insist on calling it a forest.

We let a be the total number of edges of $\mathbf{t}^0, \dots, \mathbf{t}^{l-1}$, and $I = \{c_0, c_1, \dots, c_{2a+l}\}$ be the interval of corners of $\mathbf{f}$ that are incident to the upper half-plane (hence excluding the corners that are "below" the floor), starting from the root corner of $\mathbf{t}^0$ and ending with the only corner incident to ρ^l, arranged in the usual contour order. We now equip $\mathbf{f}$ with an integer-valued labeling function $\lambda\colon V(\mathbf{f}) \to \mathbb{Z}$, again defined up to addition of a constant, that we require to satisfy the well-labeling condition in the interval I. In this case, it means that

- $\lambda(u) - \lambda(v) \in \{-1, 0, 1\}$ when u and v are neighboring vertices of the same tree,
- for $1 \leq i \leq l$, we have $\lambda(\rho^i) \geq \lambda(\rho^{i-1}) - 1$.

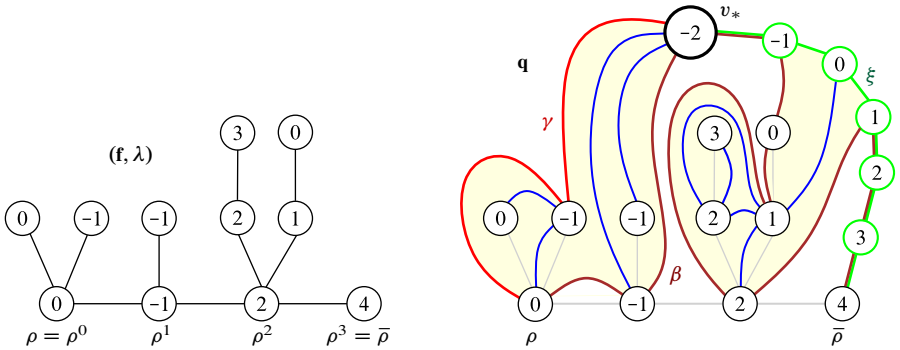

Figure 2.5. The interval Cori–Vauquelin–Schaeffer bijection giving composite slices. On this example, the forest has $a = 7$ edges and $l = 3$ trees (remember that the last vertex-tree does not count as a "real" tree). The boundary of **q** has three parts: the maximal geodesic γ (in red), the shuttle ξ (in green) and the base β (in burgundy). Its tilt is 4.

The map $\mathbf{sl} = \mathrm{CVS}(\mathbf{f}, \lambda; I)$ is then a nonrooted plane quadrangulation with one hole having a internal faces. Setting $\lambda_* = \min\{\lambda(v) : v \in V(\mathbf{f})\} - 1$, we see that the boundary of **sl** has length $2(\lambda(\rho^l) - \lambda_* + l)$. It contains three distinguished vertices: ρ, the apex v_* (the extra vertex with label λ_*), and $\overline{\rho}$, as well as three distinguished paths:

(1) the maximal geodesic γ, which has length $\lambda(\rho^0) - \lambda_*$;

(2) the shuttle ξ, which has length $\lambda(\rho^l) - \lambda_*$;

(3) the remaining boundary segment, called the *base* and denoted by β, consisting in the arcs connecting the root vertices of the trees. More precisely, if c_j denotes the last corner of the tree $\mathbf{t}^i$, then the part of the boundary of **sl** between ρ^i and ρ^{i+1} consists in the arc linking c_j to $s(c_j)$ and the successive arcs linking c_{j+1}, $s(c_{j+1})$, $s(s(c_{j+1}))$, …, $s(c_j)$. As a result, this base has length $\lambda(\rho^l) - \lambda(\rho^0) + 2l$. Moreover, any vertex of the base is at distance at most $\max_{1 \le i \le l} |\lambda(\rho^i) - \lambda(\rho^{i-1})| + 1$ from some element of the set $\{\rho^0, \dots, \rho^l\}$.

Note that, as is the case in Figure 2.5, the base may overlap with the other distinguished paths. Furthermore, as noted at the end of Section 2.1, the maximal geodesic and the shuttle are geodesic chains. On the contrary, the base is not a geodesic in general.

Definition 2.2. A map obtained by this construction will be called a *discrete composite slice*, or simply *slice* for short: its *area* is the integer a, its *width* is the integer l and its *tilt* is defined as the integer

$$\delta = \lambda(\overline{\rho}) - \lambda(\rho).$$

The terminology of composite slices, width and tilt are borrowed from [27]; however, the reader should mind that our exact definitions differ slightly from those in that reference.[4] Note also that, in the present work, we use the simplified terminology of *slice* in order to designate a composite slice. Beware that these have not to be confused with similar objects existing in the literature, in particular in our previous work [25], called *elementary slices* or also slices for short; they actually correspond to composite slices of width 0, objects that we do not consider here.

We record the following useful counting result.

Proposition 2.3. *The number of slices with area a, width l and tilt δ is equal to*

$$3^a \frac{l}{2a+l}\binom{2a+l}{a}\binom{2l+\delta-1}{l-1},$$

which can also be recast as

$$12^a 8^l 2^\delta Q_l(2a+l)P_l(\delta),$$

where $Q_\ell(u)$ is the probability that a simple random walk hits $-\ell$ for the first time at time u, and $P_\ell(j) = \mathbb{P}(G_1 + \cdots + G_\ell = j)$, where $G_1, G_2, \ldots$ are independent random variables with shifted Geometric$(1/2)$ *law, i.e., such that $\mathbb{P}(G_1 = j) = 2^{-j-2}$ for $j \geq -1$.*

Proof. The term $\frac{l}{2a+l}\binom{2a+l}{a}$ is the number of forests with l trees and a nonfloor edges, the term $\binom{2l+\delta-1}{l-1}$ counts the number of possible ways to well label the roots, and the term 3^a counts the number of ways to well label the other vertices, since it amounts to choosing a label difference in $\{-1, 0, 1\}$ along each edge.

The probabilistic form is a simple exercise using the encoding of forests and geometric walks by simple walks, yielding

$$\frac{l}{2a+l}\binom{2a+l}{a} = 2^{2a+l}Q_l(2a+l) \quad \text{and} \quad \binom{2l+\delta-1}{l-1} = 2^{2l+\delta}P_l(\delta).$$

See Section 4.4 and [21, Lemma 6]. ∎

An important feature of the construction is that the labels on $V(\mathbf{sl})$ inherited from those on $V(\mathbf{f})$ are exactly the relative distances to v_* in $\mathbf{sl}$:

$$d_{\mathbf{sl}}(v, v_*) = \lambda(v) - \lambda_*, \quad v \in V(\mathbf{sl}),$$

and that the following bound holds:

$$d_{\mathbf{sl}}(c_i, c_j) \leq \lambda(c_i) + \lambda(c_j) - 2\min_{i \leq r \leq j} \lambda(c_r) + 2, \quad i \leq j. \tag{2.1}$$

[4] In particular, in [27], the width is the length of the base, equal to $2l + \delta$ in our notation, and the tilt is the opposite $-\delta$ of what we call the tilt in this memoir.

If **sl** is a slice, using a slightly different convention from that of Section 1.4, we view it as the marked measured metric space in $\mathbb{M}^{(5,2)}$ given by

$$(V(\mathbf{sl}), d_{\mathbf{sl}}, \partial\mathbf{sl}, \mu_{\mathbf{sl}}, \nu_\beta) \quad \text{with } \partial\mathbf{sl} = (\beta, \rho, \gamma, \overline{\rho}, \xi), \tag{2.2}$$

where each boundary part is identified with the vertices it contains, where $\mu_{\mathbf{sl}}$ is the counting measure on the vertices of **sl** that *do not belong to the shuttle*, and where ν_β is the counting measure (with multiplicities) on $\beta \setminus \{\overline{\rho}\}$. The measures $\mu_{\mathbf{sl}}$ and ν_β are respectively called the *area measure* and the *base measure* of the slice. It might be surprising at this point to include ρ and $\overline{\rho}$ in the marking as these can be found from the other three marks; they are here to enter the framework of geodesic marks introduced in Section 3.2. The idea is that the data of (ρ, γ) suffice to recover the maximal geodesic as an *oriented* path, whereas the data of γ (as a set of vertices) do not give the orientation of the path.

2.4 Quadrilaterals with geodesic sides

Consider a *double forest*, that is, a pair $(\mathbf{f}, \overline{\mathbf{f}})$ of plane forests with the same number of trees. Let $h \geq 1$ denote this common number of trees and recall that this means that $\mathbf{f}$ and $\overline{\mathbf{f}}$ have h trees plus an additional vertex-tree. Similarly to the previous section, we represent it by letting

- the floors be both sent to the chain linking the points $(i, 0) \in \mathbb{R}^2$, where $0 \leq i \leq h$,
- the trees of $\mathbf{f}$ be contained in the upper half-plane $\mathbb{R} \times \mathbb{R}_{\geq 0}$, the i-th tree attached to $(i-1, 0)$ for $1 \leq i \leq h$,
- and the trees of $\overline{\mathbf{f}}$ be contained in the lower half-plane, the i-th tree attached to $(h-i+1, 0)$ for $1 \leq i \leq h$.

We obtain a nonrooted plane map, which we denote by $\mathbf{f} \cup \overline{\mathbf{f}}$, coming with the two distinguished vertices $\rho = (0, 0)$ and $\overline{\rho} = (h, 0)$. Here also, it is in fact a plane tree having two distinguished vertices.

We let $I = \{c_0, c_1, \ldots, c_{2a+h}\}$ be the interval of corners of $\mathbf{f} \cup \overline{\mathbf{f}}$, in facial order, that are incident to the upper half-plane, and $\overline{I} = \{\overline{c}_0, \overline{c}_1, \ldots, \overline{c}_{2\overline{a}+h}\}$ the interval of those incident to the lower half-plane, where a (resp. $\overline{a}$) is the number of edges in the trees of the upper (resp. lower) half-plane. As mentioned during Section 2.1, we use the slightly unusual convention that $c_{2a+h} \neq \overline{c}_0$ (and similarly $\overline{c}_{2\overline{a}+h} \neq c_0$): this means that the first corner incident to ρ is "split" in two corners, one in the upper half-plane and one in the lower half-plane.

Finally, assume that, in its unique face, the map $\mathbf{f} \cup \overline{\mathbf{f}}$ is well labeled by an integer function $\lambda\colon V(\mathbf{f} \cup \overline{\mathbf{f}}) \to \mathbb{Z}$ defined up to addition of a constant: this simply means that $\lambda(u) - \lambda(v) \in \{-1, 0, 1\}$ whenever u and v are neighboring vertices. See Figure 2.6 for an example. Note that, equivalently, a well-labeled double forest $((\mathbf{f}, \overline{\mathbf{f}}), \lambda)$ can be

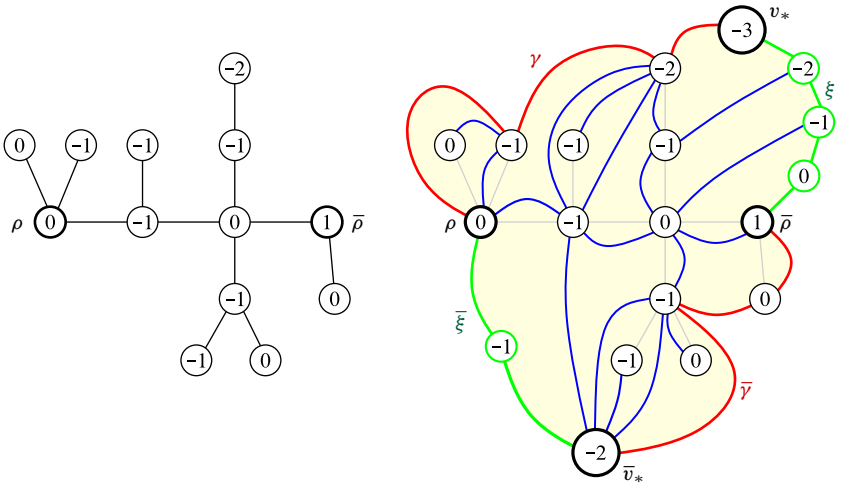

Figure 2.6. The interval Cori–Vauquelin–Schaeffer bijection giving quadrilaterals with geodesic sides. The quadrilateral with geodesic sides has half-areas 5 and 4, width 3, and tilt 1.

seen as a well-labeled *vertebrate*, that is, a well-labeled tree with two distinct distinguished vertices ρ, $\overline{\rho}$, where the interval I corresponds to the consecutive corners in the contour order from ρ to $\overline{\rho}$, which contains all the corners incident to ρ and stops at the first corner incident to $\overline{\rho}$, and $\overline{I}$ is defined similarly with the roles of ρ, $\overline{\rho}$ exchanged.

The map $\mathbf{qd} = \mathrm{CVS}(\mathbf{f} \cup \overline{\mathbf{f}}, \lambda; I, \overline{I})$ is then a nonrooted plane quadrangulation with one hole having $a + \overline{a} + h$ internal faces. Its boundary contains the four distinguished vertices ρ, the apex v_* associated with I, $\overline{\rho}$, and the apex $\overline{v}_*$ associated with $\overline{I}$, as well as the maximal geodesics γ, $\overline{\gamma}$ and shuttles ξ, $\overline{\xi}$, with obvious notation.

Definition 2.4. The *quadrilateral with geodesic sides*, or simply *quadrilateral* for short, associated with $((\mathbf{f}, \overline{\mathbf{f}}), \lambda)$ is by definition $\mathbf{qd} = \mathrm{CVS}(\mathbf{f} \cup \overline{\mathbf{f}}, \lambda; I, \overline{I})$. Its *width*, *half-areas*, and *tilt* are respectively the numbers h, a and $\overline{a}$, $\lambda(\overline{\rho}) - \lambda(\rho)$.

Observe that the parameters of a quadrilateral can be recovered from the map **qd** and the distinguished vertices ρ, $\overline{\rho}$. Furthermore, the quadrilateral associated with $((\overline{\mathbf{f}}, \mathbf{f}), \lambda)$ is obtained from the one associated with $((\mathbf{f}, \overline{\mathbf{f}}), \lambda)$ simply by switching the distinguished elements ρ with $\overline{\rho}$, γ with $\overline{\gamma}$, and ξ with $\overline{\xi}$. It has the same width, its half-areas are switched and its tilt is reversed.

Proposition 2.5. *The number of quadrilaterals with half-areas a, $\overline{a}$, width h and tilt δ is equal to*

$$12^{a+\overline{a}+h} Q_h(2a + h) Q_h(2\overline{a} + h) M_h(\delta),$$

where Q_ℓ has been defined in Proposition 2.3, and $M_\ell(j) = \mathbb{P}(U_1 + \cdots + U_\ell = j)$, where $U_1, U_2, \ldots$ are independent uniform random variables in $\{-1, 0, 1\}$.

Proof. As in the proof of Proposition 2.3, the number of forests with h floor edges and α nonfloor edges is $2^{2\alpha+h}Q_h(2\alpha+h)$, so the number of double forests with proper parameters is $4^{a+\overline{a}+h}Q_h(2a+h)Q_h(2\overline{a}+h)$. Then, the number of possible labelings of the floor vertices is the number of walks with h steps in $\{-1,0,1\}$ going from 0 to δ, which equals $3^h M_h(\delta)$. The final term $3^{a+\overline{a}}$ counts the possible labelings of the nonroot vertices in the double forest. ∎

If **qd** is a quadrilateral, we will view it as a marked measured metric space in $\mathbb{M}^{(6,1)}$ given by

$$(V(\mathbf{qd}), d_{\mathbf{qd}}, \partial\mathbf{qd}, \mu_{\mathbf{qd}}) \quad \text{with } \partial\mathbf{qd} = (\rho, \gamma, \xi, \overline{\rho}, \overline{\gamma}, \overline{\xi}), \tag{2.3}$$

where each boundary part is identified with the vertices it contains, and where $\mu_{\mathbf{qd}}$ is the counting measure on the vertices of **qd** that *do not belong to the shuttles*. We call this measure $\mu_{\mathbf{qd}}$ the *area measure* of the quadrilateral.

2.5 Scaling limits of elementary pieces

In this section, we state two important results that will be crucial in the proof of Theorem 1.1. These show that, under appropriate hypotheses, random discrete slices and quadrilaterals converge in distribution in the GHP topology toward "continuum analogs" of these objects.

We first fix three sequences $(a_n) \in (\mathbb{Z}_{\geq 0})^{\mathbb{N}}$, $(l_n) \in \mathbb{N}^{\mathbb{N}}$ and $(\delta_n) \in \mathbb{Z}^{\mathbb{N}}$ such that

$$\frac{a_n}{n} \underset{n\to\infty}{\longrightarrow} A > 0, \quad \frac{l_n}{\sqrt{2n}} \underset{n\to\infty}{\longrightarrow} L > 0 \quad \text{and} \quad \left(\frac{9}{8n}\right)^{1/4} \delta_n \underset{n\to\infty}{\longrightarrow} \Delta \in \mathbb{R}. \tag{2.4}$$

Recall that a slice is seen as an element of $\mathbb{M}^{(5,2)}$ given by (2.2) and that Ω_n is the scaling operator defined in (1.4).

Theorem 2.6. *Let* Sl_n *be uniformly distributed among composite slices with area* a_n*, width* l_n *and tilt* δ_n*. Then we have the convergence*

$$\Omega_n(\mathrm{Sl}_n) \underset{n\to\infty}{\overset{(d)}{\longrightarrow}} \mathsf{Sl}_{A,L,\Delta}$$

in distribution in the space $(\mathbb{M}^{(5,2)}, \mathrm{d}_{\mathrm{GHP}}^{(5,2)})$*. The limit is called a* (continuum composite) slice with area A, width L and tilt Δ.

This theorem will be proved in Chapter 4, where a detailed characterization of the limiting object will be given. For the time being, this theorem should be taken as a definition of the spaces $\mathsf{Sl}_{A,L,\Delta}$.

The following statement deals with the case of vanishing areas and widths, and will be useful in Section 3.7 below.

Corollary 2.7. *Let the sequences* $(a_n) \in (\mathbb{Z}_{\geq 0})^{\mathbb{N}}$ *and* $(l_n) \in (\mathbb{Z}_{\geq 0})^{\mathbb{N}}$ *satisfy* $l_n = o(\sqrt{n})$ *and* $a_n + l_n = \Theta((l_n)^2)$. *Let* Sl_n *be the vertex map whenever* $l_n = 0$, *or be uniformly distributed among slices with area* a_n, *width* l_n *and tilt* 0 *otherwise. Then we have the convergence toward the point space*

$$\Omega_n(\mathrm{Sl}_n) \xrightarrow[n\to\infty]{(d)} \{\varrho\}$$

in distribution in the space $(\mathbb{M}^{(5,2)}, \mathrm{d}_{\mathrm{GHP}}^{(5,2)})$.

Let us sketch in a few lines why this is indeed a consequence of Theorem 2.6. By the assumption that $a_n + l_n = \Theta((l_n)^2)$, the sequence $((a_n + l_n)/(l_n)^2)$, restricted to the values of n for which $l_n \neq 0$, is bounded away from 0 and ∞. This compact way of writing this property covers in fact the two following situations. If (l_n) is a bounded integer sequence, it simply means that (a_n) is a bounded integer sequence, in which case the statement becomes trivial. If (l_n) is unbounded, then it means that $(a_n/(l_n)^2)$ is bounded away from 0 and ∞. In this case, Theorem 2.6 easily implies that the diameters of $\Omega_{a_n}(\mathrm{Sl}_n)$ form a tight family of random variables. Since $a_n = o(n)$, the conclusion follows. Note that, up to extracting subsequences, we may always assume that we are in one of the two situations discussed above.

We will derive Theorem 2.6 from the known convergence of the uniform infinite half-planar quadrangulation toward the Brownian half-plane. The former naturally contains a family of slices and the latter contains a continuous "flow" of continuum slices. These consist in free versions of the objects considered here so that we will need to finish with a conditioning argument.

We now turn to quadrilaterals, which are seen as elements of $\mathbb{M}^{(6,1)}$ given by formula (2.3). We consider four sequences (a_n), $(\overline{a}_n) \in (\mathbb{Z}_{\geq 0})^{\mathbb{N}}$, $(h_n) \in \mathbb{N}^{\mathbb{N}}$ and $(\delta_n) \in \mathbb{Z}^{\mathbb{N}}$ such that, as $n \to \infty$,

$$\frac{a_n}{n} \to A > 0, \quad \frac{\overline{a}_n}{n} \to \overline{A} > 0, \quad \frac{h_n}{\sqrt{2n}} \to H > 0, \quad \left(\frac{9}{8n}\right)^{1/4} \delta_n \to \Delta \in \mathbb{R}. \tag{2.5}$$

Theorem 2.8. *Let* Qd_n *be a random variable uniformly distributed among quadrilaterals with half-areas* a_n *and* $\overline{a}_n$, *width* h_n *and tilt* δ_n. *Then we have the convergence*

$$\Omega_n(\mathrm{Qd}_n) \xrightarrow[n\to\infty]{(d)} \mathsf{Qd}_{A,\overline{A},H,\Delta}$$

in distribution in the space $(\mathbb{M}^{(6,1)}, \mathrm{d}_{\mathrm{GHP}}^{(6,1)})$. *The limit is called a* continuum quadrilateral with half-areas A and $\overline{A}$, width H and tilt Δ.

As for slices, the proof of this result is postponed to Chapter 5, where a detailed characterization of the limiting object will be given. The idea of the proof will be similar to that of Theorem 2.6, using the uniform infinite planar quadrangulation and Brownian plane as reference spaces instead of the half-planar versions mentioned above.