

# VEERING TRIANGULATIONS AND CANNON-THURSTON MAPS

FRANÇOIS GUÉRITAUD

ABSTRACT. Any hyperbolic surface bundle over the circle gives rise to a continuous surjection from the circle to the sphere, by work of Cannon–Thurston and Bowditch. We prove that the order in which this surjection fills out the sphere is dictated by a natural triangulation of the surface bundle (introduced by Agol) when all singularities of the invariant foliations are at punctures of the fiber.

## 1. INTRODUCTION

Surface bundles over the circle are historically an important source of examples in hyperbolic 3-manifold theory. Thurston proved that, barring natural topological obstructions, they always carry complete hyperbolic metrics, which was a first step towards Perelman’s geometrization of 3-manifolds. Cannon and Thurston [6] found surprising sphere-filling curves naturally associated to hyperbolic surface bundles.

In [1], Agol singled out a special class of hyperbolic surface bundles: the ones for which singularities of the invariant foliations occur only at punctures of the fiber. He proved that such surface bundles come with a natural (topological) ideal triangulation.

The purpose of this paper is to exhibit, for this family of surface bundles, a correspondence between Agol’s triangulation and the corresponding Cannon–Thurston map. The correspondence takes the form of a pair of tessellations of the plane  $\mathbb{C}$ : (1) the link of a vertex  $\Omega$  of the universal cover of Agol’s triangulation; (2) a plane tiling recording the order in which the Cannon–Thurston map fills out the sphere  $\mathbb{C} \cup \{\infty\}$ , switching colors at each passage through the parabolic fixed point  $\infty$ . Object (1) is clearly a triangulation of the plane, though it may not be realized by non-overlapping Euclidean triangles in  $\mathbb{C}$  (Agol’s triangulation is only topological, not geodesic). Object (2) is clearly a partition of  $\mathbb{C}$ , although it will take work to determine that it is actually a tessellation into topological disks (typically with fractal-looking boundaries). In the end (Theorem 1.3 below), the two tessellations turn out to fully determine each other at the combinatorial level, and in particular share the same vertex set. This connection between tessellations was previously known for punctured torus bundles by results of Cannon–Dicks [4] and Dicks–Sakuma [5], whose work was a crucial inspiration. These results were announced in [9].

---

*Date:* June 2015.

Partially supported by the Agence Nationale de la Recherche under grants ETTT (ANR-09-BLAN-0116-01), DiscGroup (ANR-11-BS01-013), Labex CEMPI (ANR-11-LABX-0007-01).

**1.1. Hyperbolic mapping tori and invariant foliations.** Let  $S$  be an oriented surface with at least one puncture, and  $\varphi : S \rightarrow S$  an orientation-preserving homeomorphism. Define the *mapping torus*  $M_\varphi := S \times [0, 1] / \sim_\varphi$ , where  $\sim_\varphi$  identifies  $(x, 1)$  with  $(\varphi(x), 0)$ . The topological type of the 3-manifold  $M_\varphi$  depends only on the isotopy type of  $\varphi$ .

Suppose  $S$  has a half-translation structure, *i.e.* a singular Euclidean metric with a finite number of conical singularities of cone angle  $k\pi$  ( $k \geq 3$ ), and total cone angle  $k'\pi$  ( $k' \geq 1$ ) around each puncture. Every straight line segment in  $S$  then belongs to a unique (singular) foliation by parallel straight lines. The surface  $S$  with cone points removed admits an isometric atlas over  $\mathbb{R}^2$  whose chart maps are all of the form  $(x, y) \mapsto (\alpha, \beta) \pm (x, y)$  for some reals  $\alpha, \beta$ . The group  $\mathrm{SL}_2(\mathbb{R})$  acts on the space of such atlases by composition with the charts, hence  $\mathrm{PSL}_2(\mathbb{R})$  acts on the space of (isometry classes of) half-translation surfaces endowed with a privileged pair of perpendicular foliations by straight lines. A landmark result of Thurston's [7, 15] is

**Fact A.** *Suppose the isotopy class  $[\varphi]$  preserves no finite system of simple closed curves on  $S$  ( $\varphi$  is called pseudo-Anosov). Then there exists a half-translation structure on  $S$  such that  $\varphi : S \rightarrow S$  is realized by a diagonal element  $(\alpha \ \alpha^{-1})$ , where  $\alpha > 1$ ; the vertical and horizontal foliations of  $S$  for this structure, called  $\lambda^+$  and  $\lambda^-$ , are preserved by  $\varphi$  and come equipped with transverse measures that are preserved up to a factor  $\alpha$  (resp.  $\alpha^{-1}$ ).*

*Moreover, the mapping torus  $M_\varphi$  admits a (unique) complete hyperbolic metric:  $M_\varphi \simeq \Gamma \backslash \mathbb{H}^3$  for some discrete group  $\Gamma < \mathrm{PSL}_2\mathbb{C} = \mathrm{Isom}_0(\mathbb{H}^3)$ .*

In Thurston's proof, which gave the first abundant source of examples of hyperbolic 3-manifolds, the foliations  $\lambda^\pm$  are in fact an important tool to construct the hyperbolic metric on  $M_\varphi$ . The half-translation structure on  $S$  in the result above is itself unique up to the action of diagonal elements of  $\mathrm{PSL}_2(\mathbb{R})$ .

**1.2. Combinatorics of veering triangulations.** We now describe a construction of Agol's triangulation, an alternative to [1] which may be of separate interest. An ideal tetrahedron is a space diffeomorphic to a compact tetrahedron minus its 4 vertices. An ideal triangulation of a 3-manifold  $M$  is a realization of  $M$  as a union of finitely many ideal tetrahedra, glued homeomorphically face-to-face.

**Definition.** *A taut structure on an ideal triangulation of an oriented 3-manifold  $M$  (into  $n$  tetrahedra) is a map from the set of all  $6n$  dihedral angles of the tetrahedra into  $\{0, \pi\}$  such that each tetrahedron has one pair of opposite edges labeled  $\pi$  and all other edges labeled 0; and each degree- $k$  edge of  $M$  is adjacent to precisely two angles labelled  $\pi$  and  $(k - 2)$  edges labelled 0.*

A taut structure can be viewed as a crude attempt at endowing the tetrahedra of  $M$  with geometric shapes in order to realize the hyperbolic metric (very crude indeed since all tetrahedra look flat!). The term is inspired by the fact that, up to a degree 2 cover, such a structure gives rise to a taut foliation.

In a rhombus of  $\mathbb{R}^2$  symmetric across both coordinate axes, we call the two edges with positive slope *rising*, the other two edges *falling*. The diagonals, which are segments of the coordinate axes, are called vertical and horizontal.

**Definition.** *A taut structure on an ideal triangulation of an oriented 3-manifold  $M$  is called veering if its edges can be 2-colored, in red and blue, so that every*

tetrahedron can be sent by an orientation-preserving map to the one pictured in Figure 1: a thickened rhombus in  $\mathbb{R}^2 \times \mathbb{R}$  with  $\pi$  on the diagonals and 0 on other edges; with the vertical diagonal in front, the horizontal diagonal in the back, rising edges red, and falling edges blue. (The diagonals might be any color.)




FIGURE 1. *Left*: a flattened tetrahedron with 4 colored edges. *Right*: the triangular link at any of the 4 vertices. Angles 0 and  $\pi$  are indicated by a graphical, train-track-like convention. The tip and base, drawn in grey, receive colors (blue/red) from the adjacent triangles. The triangle is called *hinge* if and only if the tip and base have different colors.

In [11] and [8], veering triangulations are shown to admit *positive angle structures*: this is a less crude (linearized) version of the problem of finding the complete hyperbolic metric on  $M$  endowed with a geodesic triangulation. Interestingly however, Hodgson, Issa and Segerman found veering triangulations that are in fact not realized geodesically but have instead some tetrahedra turning “inside out” [10].

In [1], Agol described a *canonical, veering* triangulation of a general hyperbolic mapping torus  $M = M_\varphi$ , provided all singularities of the foliations  $\lambda^+, \lambda^-$  occur at punctures of the fiber  $S$ . Our first main result is an alternative construction of Agol’s triangulation (details in Section 2).

**Theorem 1.1.** *Suppose all singularities of the invariant foliations  $\lambda^\pm$  of the pseudo-Anosov monodromy  $\varphi : S \rightarrow S$  are at punctures of  $S$ . Any maximal immersed rectangle  $R$  in  $S$  with edges along leaf segments of  $\lambda^\pm$  contains one singularity in each of its four sides. Connecting these four ideal points, and thickening in the direction transverse to  $S$ , yields a tetrahedron  $\Delta_R \subset S \times \mathbb{R}$ . The tetrahedra  $\Delta_R$  glue up to yield a veering triangulation of  $S \times \mathbb{R}$ , compatible with the equivalence relation  $(x, t + 1) \sim_\varphi (\varphi(x), 0)$ , hence descending to  $M_\varphi$ .*

The veering structure in the above theorem is given as follows:  $\Delta_R$  has its  $\pi$ -angles at the edges connecting points belonging to opposite sides of  $R$  (the edge connecting horizontal sides being in front); an edge of positive slope is red; an edge of negative slope is blue.

Remark: unlike Agol’s original definition, this construction does not rely on any auxiliary choices (*e.g.* of train tracks). As a result, it should generalize to the Cannon-Thurston maps of degenerate surface groups built by Mahan Mj [14], when the ending laminations define foliations with no saddle-point connection. The focus of this paper being combinatorics, we choose however to restrict to surface bundles.

In order to state the main result, we now point out some features inherent to any veering triangulation  $\mathcal{T}$ : see Figure 2, and [8] for detailed proofs. The link of a vertex  $\Omega$  of the universal cover  $\tilde{\mathcal{T}}$  is a tessellation  $\Delta$  of the plane. The vertices and edges of  $\Delta$  receive colors (red/blue) from the edges of  $\mathcal{T}$ .




FIGURE 2. Three adjacent ladders in the vertex link  $\Delta$  of a veering triangulation. Ladderpoles (vertical) are slightly thicker. The middle ladder is ascending (tips above base rungs). The colors of the rungs are determined by the combinatorics of this chunk of  $\Delta$ , except the 3 dotted rungs. Triangles with no dotted edge are shaded when they are hinge.

Among the triangles of  $\Delta$ , we may distinguish two types: a triangle coming from truncating a tetrahedron (thickened rhombus) whose diagonals are of opposite colors is called *hinge*; other triangles are *non hinge*. See Figure 1 and its caption.

An edge of  $\Delta$  connecting two vertices of the same color is called a *ladderpole edge* (and is always of the opposite color). The other edges are called *rungs*. It turns out that each vertex belongs to exactly two ladderpole edges, and ladderpole edges arrange into infinitely many, disjointly embedded simplicial lines of alternating colors, called *ladderpoles*. Every rung connects two vertices from two consecutive ladderpoles. The region between two consecutive ladderpoles is called a *ladder*. In each ladder, all triangles have their  $\pi$ -angle, or *tip*, on the same side of their base rung (say *above* the base rung if we arrange the ladder vertically with suitable orientation); in the next ladder the tips are on the other side (*below* the base rungs). Ladders of the former type are called ascending, of the latter type descending; see Figure 2.

Note that vertices of  $\Delta$  have well-defined coordinates in the plane  $\mathbb{C}$  (up to similarity), given by any developing map of the hyperbolic metric on  $M_\varphi$  that takes  $\Omega$  to  $\infty \in \mathbb{P}^1\mathbb{C} = \partial_\infty\mathbb{H}^3$ . However, we will usually draw the ladderpoles as vertical lines (with regular meanders to respect the train-track convention for angles 0 and  $\pi$  as in Figure 2), emphasizing the combinatorics at the expense of the geometry.

**1.3. The Cannon-Thurston map, and the Main Result.** Let  $\Sigma$  (a disk) be the universal cover of the fiber  $S$  of the hyperbolic surface bundle  $M = M_\varphi$ , and  $\mathbb{S}$  (a circle) the natural boundary of  $\Sigma$ . The inclusion  $S \rightarrow M$  lifts to a map  $\iota : \Sigma \rightarrow$

$\mathbb{H}^3$  between the universal covers. Bowditch [3], generalizing work of Cannon and Thurston [6], proved the following surprising fact.

**Fact B.** *The map  $\iota$  extends continuously to a boundary map  $\bar{\tau} : \mathbb{S} \rightarrow \mathbb{P}^1\mathbb{C}$  which is a surjection from the circle to the sphere. The endpoints of any leaf of  $\lambda^\pm$  (lifted to  $\Sigma$ ) have the same image under  $\bar{\tau}$ , and this in fact generates all the identifications occurring under  $\bar{\tau}$ .*

In this note, we prove a correspondence between the combinatorics of the Cannon-Thurston map  $\bar{\tau}$  (the “order in which  $\bar{\tau}$  fills out the sphere”) and the triangulation  $\Delta$  of the plane. To state the correspondence, we will first need to prove facts about  $\bar{\tau}$  (details in Section 4).

Recall the chosen parabolic fixed point  $\infty$  of the Kleinian group  $\Gamma$ . The surjection  $\bar{\tau} : \mathbb{S} \rightarrow \mathbb{C} \cup \{\infty\}$  goes infinitely many times through the point  $\infty$ , by Fact B (indeed there are infinitely many leaves terminating at a given parabolic boundary point of  $\Sigma$ ). We may imagine that  $\bar{\tau}$  *changes color* (red/blue) each time it goes through  $\infty$ : the resulting coloring of the plane  $\mathbb{C}$  becomes an interesting object to look at.

**Theorem 1.2.** *There exists a  $\mathbb{Z}$ -family of Jordan curves  $J_i$  of  $\mathbb{C} \cup \{\infty\}$ , bounding domains  $D_i$ , with the following properties:*

- For all  $i$  the curve  $J_i$  goes through  $\infty$ ;
- $\cdots \supset D_{-1} \supset D_0 \supset D_1 \supset D_2 \supset \cdots$ ;
- $\bigcap_{i \in \mathbb{Z}} D_i = \emptyset$  and  $\bigcup_{i \in \mathbb{Z}} D_i = \mathbb{C}$ ;
- $J_i \cap J_{i'} = \{\infty\}$  if and only if  $|i' - i| > 1$ ;
- For every  $i \in \mathbb{Z}$ , the closure of  $D_i \setminus D_{i+1}$  in  $\mathbb{C}$  is the union of a family  $\{\delta_s^i\}_{s \in \mathbb{Z}}$  of closed disks, all disjoint except that each  $\delta_s^i$  shares one boundary point with  $\delta_{s+1}^i$ ;
- Between the  $i$ -th and  $(i+1)$ -st color switches, the map  $\bar{\tau}$  fills out the  $\delta_s^i$  one by one; the order of filling switches with the parity of  $i$ .

By this theorem, the trajectory of the plane-filling curve  $\bar{\tau}$  is reminiscent of that of a plowing ox (or boustrophedon, the name of an ancient writing style): we consequently call the closure of  $D_i \setminus D_{i+1}$  a *furrow*; see Figure 3. The disks  $\{\delta_s^i\}_{(i,s) \in \mathbb{Z}^2}$  making up all the furrows define a tessellation of  $\mathbb{C}$  in which every vertex has order 4, and is adjacent to two consecutive disks of the  $i$ -th furrow, one disk of the  $(i-1)$ -st furrow and one disk of the  $(i+1)$ -st furrow (for some  $i$ ). Each disk  $\delta$  of the  $i$ -th furrow has:

- 2 vertices (the *gates* of  $\delta$ ) shared with other disks of the  $i$ -th furrow;
- some nonnegative number of vertices (called *spikes* of  $\delta$ ) shared with disks of the  $(i+2)$ -nd or  $(i-2)$ -nd furrow.

Edges of the tessellation always separate disks from consecutive furrows; they come in two types (Figure 3):

- *in-furrow* edges, connecting consecutive gates of the same furrow (or equivalently, two consecutive spikes of some disk of the adjacent furrow);
- *cross-furrow* edges, connecting two gates of adjacent furrows (or equivalently, the first or last spike on one side of some disk to the adjacent gate).

With this terminology, we can state our main result, which consists of (1) a full dictionary between the various features of the two tessellations, and (2) a recipe book to reconstruct (combinatorially) one tessellation from the other.




FIGURE 3. Six Jordan curves  $J_i$  through  $\infty$  bound 5 consecutive furrows, alternatingly red and blue. The domain  $\overline{D_i}$  is the area to the right of  $J_i$ . The disk  $\delta$  belongs to the furrow  $\overline{D_2} \setminus D_3$ . Arrows materialize in which order furrows are filled out by  $\tau$ .

**Theorem 1.3. (1) Dictionary.** *The Cannon-Thurston tessellation and the link  $\Delta$  of the Agol triangulation have the same vertex set, and there are natural bijections between the following objects:*

Cannon-Thurston tessellation		Triangulation $\Delta$	
	Vertices	Vertices	
	Furrows	Ladderpoles	
	Disks	Ladderpole edges	
	Spikes	Rungs	
In-furrow	} edges	Non-hinge	} triangles
Cross-furrow		Hinge	

**(2) Recipe book.** *Given the link  $\Delta$ , we can obtain topologically the 1-skeleton of the Cannon-Thurston tessellation by drawing, for each triangle, an arc from its tip to the tip of the next triangle across the base rung (this may create double edges along the ladderpoles: keep them). See Figure 4, left.*

*Conversely, given the Cannon-Thurston tessellation, we can obtain the 1-skeleton of the triangulation  $\Delta$  by adding edges connecting each gate of a blue (resp. red) cell to all the vertices clockwise (resp. counterclockwise) until the other gate, and deleting redundant edges. See Figure 4, right.*

In (2), the redundant edges to be deleted in the last step can be either present in the initial Cannon-Thurston tessellation (namely, the boundary of a cell  $\delta_s^i$  with no spikes consists of two mutually isotopic edges along a ladderpole; Figure 3 shows four such “small” cells) — or they can be created during the process of adding extra edges.

**Remark 1.4.** *As mentioned above, the vertices of the two tessellations of Theorem 1.3 are well-defined complex algebraic numbers. Due to the orientation issue raised*




FIGURE 4. *Left*: two ladders (ascending and descending) of the Agol triangulation, with, superimposed in green, the tip-to-tip edges to be drawn to obtain the topological 1-skeleton of the Cannon-Thurston tessellation. Hinge triangles are shaded.

*Right*: two disks (from a red and a blue furrow) of the Cannon-Thurston tessellation, with, superimposed in green, the edges to be inserted to obtain the Agol cusp triangulation (before deletion of redundant edges).

in [10], however, edges of  $\Delta$  must be thought of combinatorially, not as straight segments. On the other hand, edges of the Cannon-Thurston tessellation are precise, fractal-looking plane curves: some examples are beautifully rendered in [4, 5].

**Notation.** Throughout the paper, we will denote by  $\Sigma$  the universal cover of the fiber  $S$  of the surface bundle  $M_\varphi$ , and by  $\overline{\Sigma}$  and  $\overline{S}$  the metric completions of  $\Sigma$  and  $S$  respectively, for the locally Euclidean

$$(1.1) \quad \begin{array}{ccc} \Sigma & \leftrightarrow & \overline{\Sigma} \\ \downarrow & & \downarrow \\ S & \leftrightarrow & \overline{S}. \end{array}$$

However, note that the rightmost vertical map  $\overline{\Sigma} \rightarrow \overline{S}$  is *not* a universal covering: it has infinite branching above all the points representing punctures of  $S$ .

**1.4. Plan of the paper.** In Section 2 we prove Theorem 1.1. In Section 3 we study geodesics in a half-translation surface to produce a combinatorial description of the source circle  $\mathbb{S}$  of the Cannon-Thurston map  $\tau$ . In Section 4 we use this understanding to prove Theorem 1.2 on the combinatorics of  $\tau$ . Finally, in Section 5 we prove Theorem 1.3, making each line of the “dictionary” correspond to a certain type of rectangle in the foliated surface  $(\Sigma, \lambda^\pm)$ . In Section 6 we give some extra illustrations.

I am grateful to Makoto Sakuma for pointing out a mistake in an earlier version of Section 3, and to the anonymous referee for remarks that helped improve the exposition. This work was completed during an extended visit to the Wolfgang-Pauli Institute, Vienna.

## 2. THE CANONICAL VEERING TRIANGULATION

We now prove Theorem 1.1. Let  $S, \lambda^\pm, \varphi$  be as in the theorem, and  $\bar{S}, \Sigma, \bar{\Sigma}$  as in (1.1). Let  $g_0$  be a singular Euclidean metric on  $\bar{\Sigma}$  that makes the measured foliations  $\lambda^+$  and  $\lambda^-$  vertical and horizontal, and  $g_t := (e^t e^{-t})g_0$ .

By a singularity-free rectangle in  $\bar{\Sigma}_t := (\bar{\Sigma}, g_t)$ , we mean an embedded rectangle whose sides are leaf segments of  $\lambda^\pm$  and which contains no singularity except possibly in its boundary. Note that a singularity-free rectangle contains *at most* one singularity in each edge: indeed no leaf of  $\lambda^+$  or  $\lambda^-$  can connect two singularities, otherwise there would be arbitrarily short such leaves (by applying  $\varphi^{\pm 1}$  many times), contradicting the fact that the singular set is finite in the quotient  $\bar{S}$ .

A singularity-free rectangle in  $\bar{\Sigma}_t$  receives a height and a width (depending on  $t$ ) from the metric  $g_t$ . We speak of a singularity-free *square* if the height and width are equal.

The following construction is an analogue of the Delaunay triangulation relative to the singular set, with circles replaced by squares.<sup>1</sup> In the context of this paper, we will just refer to it as the Delaunay cellulation.

**Proposition 2.1.** *Connecting the singularities found in the boundary of every maximal singularity-free square of  $\bar{\Sigma}_t$  produces, in the quotient  $\bar{S}_t$ , finitely many triangles and (exceptionally) quadrilaterals, which define a cellulation of  $\bar{S}_t$ .*

*Proof.* First, there exist maximal singularity-free squares in  $\bar{\Sigma}_t$ : to find one, start with a singularity  $p$ ; construct a small square  $Q$  containing  $p$  in the interior of one edge; then scale  $Q$  up with respect to  $p$  until  $Q$  bumps into another singularity  $p'$ . If  $p'$  belongs to the interior of the edge opposite to  $p$  then  $Q$  is maximal. Otherwise,  $Q$  has a corner  $q$  whose two adjacent closed edges contain  $\{p, p'\}$  in their union; scale  $Q$  further up with respect to  $q$  until  $Q$  bumps into a third singularity  $p''$ , necessarily in one of the two remaining edges. Then  $Q$  is maximal. (For exceptional values of  $t$ , a fourth vertex  $p'''$  may simultaneously appear in the last edge.)

We call the convex hull of the singularities contained in the boundary of a maximal singularity-free square a *Delaunay cell*. Delaunay cells can be either edges, or triangles, or quadrilaterals (the latter do not occur for generic  $t$ ). We claim that:

(2.1) *Two distinct Delaunay cells in  $\bar{\Sigma}_t$  can only intersect (if at all) along an edge or a vertex.*

To see this, consider two maximal singularity-free squares  $Q = ABCD$  and  $Q' = A'B'C'D'$  in  $\bar{\Sigma}_t$ . If  $Q$  and  $Q'$  have disjoint interiors, the conclusion is immediate. Otherwise, denote by  $D_Q, D_{Q'}$  the Delaunay cells in  $Q$  and  $Q'$  respectively. Up to permuting  $Q, Q'$ , rotating and relabelling the vertices, we are in one of the following 5 situations (see Figure 5):

- (1) The open segment  $(BC)$  intersects  $(A'B')$  at a point  $E$ , and  $(CD)$  intersects  $(A'D')$  at a point  $F$ . Since no singularities lie in the interior of either square, all singularities in  $\partial Q$  belong to the broken line  $[EBADF]$ , so  $D_Q$  is contained in the pentagon  $EBADF$ . Similarly  $D_{Q'} \subset EB'C'D'F$ . These pentagons share just one edge  $EF$ , hence the result.

<sup>1</sup>An analogous construction works with circles (or squares) replaced by any convex, centrally symmetric plane shape  $\mathcal{C}$ , provided no straight segment between singularities is parallel to a segment of  $\partial\mathcal{C}$ . See [12] for related ideas in dimension  $\geq 2$ , as well as [2] and references therein.






FIGURE 5. Five possible relative positions of maximal singularity-free squares  $Q = ABCD$  and  $Q' = A'B'C'D'$ . The subsets of  $\partial Q$  and  $\partial Q'$  that may contain singularities are marked by thicker lines, respectively green and purple. Convex regions that contain the Delaunay polygons  $D_Q$  and  $D_{Q'}$  are shaded.

- (2) The open segment  $(BC)$  intersects  $(A'B')$  at a point  $E$  and  $(C'D')$  at a point  $F$ , with  $B, E, F, C$  lined up in that order. Then all singularities in  $\partial Q$  belong to the broken line  $[EBADCF]$ , so  $D_Q$  is contained in its convex hull  $ABCD$ ; similarly  $D_{Q'} \subset EB'C'F$ . These rectangles share just one edge  $EF$ , hence the result.
- (3) The open segment  $(BC)$  intersects  $(A'B')$  at a point  $E$ , the points  $D, D', C, C'$  are lined up in that order, and  $[DC]$  contains a singularity  $F$ . Since the leaves of  $\lambda^\pm$  through  $F$  contain no other singularity than  $F$ , the singularities of  $\partial Q$  (resp.  $\partial Q'$ ) lie in the broken line  $[FDABE]$  (resp.  $[FC'B'E]$ ). The convex hulls of these broken lines are polygons sharing just one edge  $EF$ .
- (4) The open segment  $(BC)$  intersects  $(A'B')$  at a point  $E$ , the points  $D, D', C, C'$  are lined up in that order, and  $[DC]$  contains *no* singularity. The singularities of  $\partial Q$  (resp.  $\partial Q'$ ) lie in the broken line  $[DABE]$  (resp.  $CC'B'E$ ). The convex hulls of these broken lines are polygons sharing just one vertex  $E$ .
- (5) The points  $A, A', B, B'$  are lined up in that order, and so are  $D, D', C, C'$ . Let  $E$  be a point on  $[A'B']$ , equal to the singularity if there is one on this segment. Let  $F$  be a point on  $[DC']$ , equal to the singularity if there is one on this segment. The singularities in  $\partial Q$  (resp.  $\partial Q'$ ) are in the broken line  $[EADF]$  (resp.  $[EB'C'F]$ ). The convex hulls of these broken lines are polygons sharing just one edge  $EF$ . This proves (2.1).

Since the Delaunay polygons have disjoint interiors in  $\bar{\Sigma}_t$ , in particular (2.1) implies that the projections of these interiors in the quotient surface  $S_t$  are *embedded* (not just immersed).

Next, we claim that every side  $[AB]$  of a Delaunay polygon is a side of exactly one other Delaunay polygon, adjacent to the first one. This is clear if one considers the 1-parameter family of all squares in the Euclidean plane containing a pair of points  $\{A, B\}$  in their boundary (Figure 6):  $[AB]$  subdivides each square of the family into two regions, which vary monotonically (for the inclusion) in opposite directions with the parameter  $\tau$ . Thus for each side of  $[AB]$  there is an extremal value of  $\tau$  at which the region on that side bumps for the first time into one or (exceptionally) two singularities. The convex hull of the union of these singularities and  $\{A, B\}$  is the Delaunay polygon on that side of  $[AB]$ .

The Delaunay polygons therefore define a cell decomposition of some region of  $\bar{S}$ . This region is open and closed, because there are only finitely many Delaunay polygons in  $\bar{S}$  (the diameter of  $\bar{S}$  gives an upper bound on the possible sizes of




FIGURE 6. The extremal squares circumscribed to a segment  $[AB]$ , and two intermediate squares (colored). Dotted: the Delaunay triangles containing  $[AB]$ .

singularity-free squares, so a compactness argument applies). Therefore, the Delaunay polygons give a cell decomposition (generically a triangulation, but possibly nonsimplicial) of  $\overline{S}$  itself.  $\square$

**Remark 2.2.** *The ideal Delaunay decomposition of  $S_t = (S, g_t)$  (obtained by removing the singularities of  $\overline{S}$ ) varies with  $t$ . The changes occur at the values of  $t$  such that the decomposition contains a quadrilateral inscribed in a square  $Q$  (or several such quadrilaterals). At such times  $t$ , the triangulation undergoes a diagonal exchange: before  $t$  the triangulation contains an edge connecting the singularities on the vertical sides of  $Q$ ; after  $t$  the triangulation contains an edge connecting the singularities on the horizontal sides of  $Q$ .*

Interpreting such diagonal exchanges as (flattened) tetrahedra as in Figure 1, we obtain a so-called *layered* taut ideal triangulation of  $S \times \mathbb{R}$  which naturally descends to the quotient  $S$ -bundle  $M_\varphi$ . This taut triangulation is veering by construction if we just color edges in red or blue according to the sign of their slope in the translation surface  $S$ . Theorem 1.1 is proved.

An important result of [1] is that there is at most one layered veering triangulation of  $M_\varphi$ , so it has to be the one constructed in Theorem 1.1.

### 3. NORMAL FORMS OF POINTS OF $\partial_\infty \Sigma$

In this section,  $S$  is a half-translation surface with at least one puncture, no singularities (except at punctures), and capable of carrying a complete hyperbolic metric  $h$ , for which the punctures become cusps. The universal cover  $\mathbb{H}^2$  of  $(S, h)$  has a boundary at infinity which is a topological circle  $\mathbb{S}$ . This circle  $\mathbb{S}$  is topologically independent of the choice of  $h$ , in the sense that if  $h'$  is another hyperbolic metric on  $S$ , then the identity map from  $(S, h)$  to  $(S, h')$  lifts to a self-homeomorphism of  $\mathbb{H}^2$  which extends to a self-homeomorphism of  $\mathbb{S}$ .

Alternatively we can obtain the circle  $\mathbb{S}$  in the following way. The fundamental group  $\pi_1(S)$  is free, and its Gromov boundary  $\partial_\infty(\pi_1(S))$  is a Cantor set. Moreover,  $\partial_\infty(\pi_1(S))$  carries a natural cyclic order induced by the orientation of  $S$  (or its universal cover, the disk  $\Sigma$ ). The circle  $\mathbb{S}$  is naturally identified with the quotient of the Cantor set  $\partial_\infty(\pi_1(S))$  under the equivalence condition  $\sim$  that *collapses any two points that are not separated by a third (distinct) point* for the cyclic order. Indeed, since any two such points are the attracting and repelling fixed points

$\xi_\alpha^+, \xi_\alpha^- \in \partial_\infty(\pi_1(S))$  of a peripheral element  $\alpha$  of  $\pi_1(S)$ , a holonomy representation  $\rho$  of the hyperbolic metric  $h$  on  $S$  takes  $\alpha$  to a parabolic isometry  $\rho(\alpha)$  of  $\mathbb{H}^2$  fixing a unique ideal point  $\eta_\alpha \in \partial_\infty \mathbb{H}^2$ . The map  $\{\xi_\alpha^+, \xi_\alpha^-\} \mapsto \eta_\alpha$  extends naturally to a  $\rho$ -equivariant homeomorphism

$$\psi : \partial_\infty(\pi_1(S))/\sim \xrightarrow{\sim} \partial_\infty \mathbb{H}^2 = \mathbb{S}.$$

In the proposition below, we describe the points of the circle  $\mathbb{S}$  as “normalized” paths (possibly semi-infinite) in the singular Euclidean surface  $\overline{\Sigma}$ ; see [13] for a similar construction. The singular points of  $\overline{\Sigma}$  are infinite branching points of the covering map towards the completion  $\overline{S}$  of  $S$ : let  $\Omega \in \overline{\Sigma}$  denote such a singular point, fixed throughout the paper.

**Proposition 3.1.** *There exists a natural parameterization  $\Psi$  of the circle  $\mathbb{S}$  by the collection of all piecewise straight paths  $\gamma$  in  $\overline{\Sigma}$  starting at  $\Omega$  and turning only at singularities (leaving angles  $\geq \pi$  as they turn), possibly terminating at a singularity.*

*Given two distinct, nontrivial such paths  $\gamma$  and  $\gamma'$ , the 3 ideal points  $\Psi(\gamma)$ ,  $\Psi(\gamma')$ ,  $\Psi(\underline{\Omega})$  in  $\partial_\infty \mathbb{H}^2$  (where  $\underline{\Omega}$  is the “trivial” path that stays at  $\Omega$ ) span an ideal triangle of  $\mathbb{H}^2$ . This triangle is clockwise oriented if and only if the trajectory of  $\gamma'$  in  $\overline{\Sigma}$  departs from that of  $\gamma$  to the right, possibly after a finite common prefix. This must happen at a singularity which the paths may pass on either side, or stop at: see Figure 7.*




FIGURE 7. Illustration of Proposition 3.1. *Top*: a sequence of paths passing through a singularity and progressing clockwise in the space  $\mathbb{S}$  of paths. *Bottom*: at a singularity, a path  $\gamma'$  (green) departs from another path  $\gamma$  (purple) to the right after a nonempty common prefix (black), in all of 5 possible ways.

*Proof of Proposition 3.1.* A local minimizing argument shows that the geodesics of  $\overline{\Sigma}$  are exactly the piecewise straight curves  $\gamma$  that turn only at singularities, with all turning angles being  $\geq \pi$ . Since  $\overline{\Sigma}$  is a CAT(0) space, geodesics exist between any two points  $p, q$  of  $\overline{\Sigma}$  and are unique.

The map  $\psi$  already identifies  $\mathbb{S}$  with the quotient space  $\partial_\infty(\pi_1(S))/\sim$ . Let us identify the latter set with the space of geodesics of  $\overline{\Sigma}$  starting at  $\Omega$ , via a map  $\Phi$

to be defined. (Note that we *cannot* identify  $\partial_\infty(\pi_1(S))$  or  $\partial_\infty(\pi_1(S))/\sim$  with a ‘‘Gromov boundary of  $\overline{\Sigma}$ ’’, as  $\overline{\Sigma}$  is not Gromov-hyperbolic.)

We can find a connected, polygonal fundamental domain  $P$  of  $\overline{\Sigma}$  with all vertices at singularities, for example by taking an appropriate union of cells of one of the Delaunay cellulations of Section 2. Up to fixing a connecting path between the chosen singularity  $\Omega$  and a lift of the basepoint of  $S$ , we can also view any  $\xi \in \partial_\infty(\pi_1(S))$  as an infinite sequence  $(P_i)_{i \geq 0}$  of distinct copies of  $P$  in  $\overline{\Sigma}$  such that  $\partial P_0$  contains  $\Omega$  and each  $P_i$  shares an edge with  $P_{i+1}$ . Conversely, any such sequence  $(P_i)_{i \geq 0}$  defines a point  $\xi \in \partial_\infty(\pi_1(S))$ .

If  $\xi \in \partial_\infty(\pi_1(S))$  is a fixed point of a peripheral  $\alpha \in \pi_1(S)$ , fixing a singularity  $\Omega' \neq \Omega$  of  $\overline{\Sigma}$ , then there exists a (smallest)  $i_0$  such that all the  $P_i$  for  $i \geq i_0$  share  $\Omega'$  as a vertex. The other fixed point  $\xi'$  of  $\alpha$  is obtained by replacing this suffix  $(P_i)_{i \geq i_0}$ , formed of a sequence of polygons turning around  $\Omega'$  in one direction, by the sequence turning in the other direction. The *pair*  $\{\xi, \xi'\} \in \partial_\infty(\pi_1(S))/\sim$  of peripheral fixed points can thus naturally be identified with the unique geodesic path  $\gamma$  in  $\overline{\Sigma}$ , issued from  $\Omega$ , visiting the polygons  $P_0, \dots, P_{i_0}$ , and terminating at  $\Omega'$ . (Exceptionally,  $\gamma$  may be reduced to one edge of  $P_0$ .) Thus, set  $\Phi(\gamma) = \{\xi, \xi'\} \in \partial_\infty(\pi_1(S))/\sim$ .

By extension, if  $\{\xi, \xi'\} \in \partial_\infty(\pi_1(S))/\sim$  is the pair of fixed points of the peripheral fixing  $\Omega$  itself, then  $\{\xi, \xi'\}$  identifies with the trivial path  $\underline{\Omega}$ : that is,  $\Phi(\underline{\Omega}) = \{\xi, \xi'\} \in \partial_\infty(\pi_1(S))/\sim$ .

Finally, if  $\xi \in \partial_\infty(\pi_1(S))$  is *not* a fixed point of any a peripheral element of  $\pi_1(S)$ , then any singularity of  $\overline{\Sigma}$  belongs to at most finitely many of the polygons  $P_i$ . As a consequence, a local compactness argument applies to show that any sequence of geodesics  $\gamma_i$  of  $\overline{\Sigma}$ , connecting  $\Omega$  to a point of  $P_i$ , converges to an infinite geodesic ray of  $\overline{\Sigma}$ . This ray  $\gamma$  is unique by standard CAT(0) estimates, since the  $P_i$  have bounded diameter but escape far from  $\Omega$ . We thus set  $\Phi(\gamma) = \{\xi\} \in \partial_\infty(\pi_1(S))/\sim$ .

In short, the space of finite (resp. infinite) geodesics  $\gamma \subset \overline{\Sigma}$  issued from  $\Omega$  identifies, via  $\Phi$ , with the unordered pairs (resp. singletons) in the domain  $\partial_\infty(\pi_1(S))/\sim$  of the map  $\psi$ . The map  $\Psi := \psi \circ \Phi$  is a bijection from this set of geodesics in  $\overline{\Sigma}$  to the circle  $\mathbb{S} \simeq \partial_\infty \mathbb{H}^2$ .

Lastly, to check the statement on clockwise orientations, pick points  $p \neq p'$  in  $\mathbb{S} \setminus \{\Psi(\underline{\Omega})\}$ , let  $(\gamma, \gamma') = (\Psi^{-1}(p), \Psi^{-1}(p'))$  be the corresponding geodesic paths in  $\overline{\Sigma}$ , and let  $\mathcal{L}, \mathcal{L}'$  be the oriented geodesics in  $\mathbb{H}^2$  from  $\Psi(\underline{\Omega})$  to  $p, p'$ .

Recall the holonomy representation  $\rho : \pi_1(S) \rightarrow \text{Isom}_0(\mathbb{H}^2)$  of the hyperbolic metric  $h$  on  $S$ . The polygon  $P \subset \overline{\Sigma}$  determines an ideal hyperbolic polygon  $Q$  of  $\mathbb{H}^2 \simeq \Sigma$ , whose  $\rho(\pi_1(S))$ -orbit tiles  $\mathbb{H}^2$ .

The result then follows from the following facts:

- $p'$  lies to the right of  $p$  as seen from  $\Psi(\underline{\Omega})$  if and only if  $\mathcal{L}'$  lies to the right of  $\mathcal{L}$  in  $\mathbb{H}^2$ ;
- $\mathcal{L}$  (resp.  $\mathcal{L}'$ ) visits the interior of a tile  $\rho(\gamma) \cdot Q$  of  $\mathbb{H}^2$  if and only if  $\gamma$  (resp.  $\gamma'$ ) visits the interior of  $\gamma \cdot P$ ;
- there exists an equivariant, orientation-preserving homeomorphism from  $\mathbb{H}^2$  to the regular part  $\Sigma$  of  $\overline{\Sigma}$ , taking  $Q$  to  $P$ .

□

In the remainder of this paper, the word *path* will usually refer to a geodesic path in  $\overline{\Sigma}$ , issued from the singular point  $\Omega$ . The topology on the space of paths is induced by its identification with  $\mathbb{S}$  (dropping  $\Psi$  from the notation). It may also be

described thus: *two nontrivial paths are close if, when followed from  $\Omega$ , they stay close in  $\overline{\Sigma}$  for a great amount of length; or if they coincide up to a singularity at which they both turn by a great amount.* (Terminating at a singularity is interpreted as infinite turning, a “cusp excursion of infinite depth” in the hyperbolic metric  $h$ .)

#### 4. THE COMBINATORICS OF THE CANNON-THURSTON MAP $\bar{\tau}$

In this section we prove Theorem 1.2 about the combinatorics of the Cannon-Thurston map  $\bar{\tau}$ . Our tools are Proposition 3.1 for the description of the domain of  $\bar{\tau}$  (paths) with its cyclic order, and Fact B for the fibers of  $\bar{\tau}$ .

As in the previous section, let  $\Omega$  be a singularity of  $\overline{\Sigma}$  (the metric completion of the universal cover of the flat punctured surface  $S$ ), and let the circle  $\mathbb{S}$  be the space of  $\overline{\Sigma}$ -geodesic paths  $\gamma$  issued from  $\Omega$ . By Fact B, the points  $\gamma$  of  $\mathbb{S}$  that have the same image as the trivial path  $\underline{\Omega}$  under the map  $\bar{\tau}$  (the “ $\bar{\tau}$ -fiber of  $\underline{\Omega}$ ”) are  $\underline{\Omega}$  itself, and all the rays shooting out from  $\Omega$  along a leaf of  $\lambda^+$  or  $\lambda^-$ . Since  $\Omega$  is a branching point of infinite order, these rays (which hit no other singularity than  $\Omega$ ) form a  $\mathbb{Z}$ -family, counting clockwise: for any integer  $i$ , the  $i$ -th ray  $\ell_i$  is vertical for  $i$  even and horizontal for  $i$  odd.

**4.1. Fibers of  $\bar{\tau}$  and colors.** We apply a Möbius map to normalize so that  $\bar{\tau}(\underline{\Omega}) = \infty \in \mathbb{P}^1\mathbb{C}$ .

Every point  $\gamma$  of  $\mathbb{S}$  that is *not* in the fiber of  $\underline{\Omega}$  falls inbetween  $\ell_i$  and  $\ell_{i+1}$  for a unique integer  $i$ , which is even (resp. odd) if and only if the initial segment of (the geodesic representative of)  $\gamma$  has positive (resp. negative) slope. We say that the directions between  $\ell_i$  and  $\ell_{i+1}$  form a *quadrant* at  $\Omega$ . There are countably many quadrants.

Therefore, the rule that  $\bar{\tau}$  changes color at each passage through  $\infty$  means that the color of  $\bar{\tau}(\gamma)$  is determined by the sign of the slope of the initial segment of  $\gamma$ , or equivalently, by the parity of the quadrant that contains this initial segment.

To understand the interfaces between the two colors in the image of  $\bar{\tau}$ , we must therefore understand when a  $\bar{\tau}$ -fiber (described by Fact B) contains two paths  $\gamma, \gamma'$  whose initial segments belong to *different* quadrants at  $\Omega$ . In general, let  $Q(\gamma)$  denote the quadrant containing the initial segment of  $\gamma$ : for example, a sequence  $(\gamma_n)$  in  $\mathbb{S}$  converges to the trivial path  $\underline{\Omega}$  exactly when  $Q(\gamma_n)$  escapes any finite collection of quadrants.

**Definition 4.1.** For any  $i \in \mathbb{Z}$  and  $t \in \overline{\mathbb{R}^+} := [0, +\infty]$ , let  $\ell_i$  be the  $i$ -th leaf of  $\lambda^\pm$  issued from  $\Omega$  for the clockwise order.

The right (resp. left)  $t$ -hook along  $\ell_i$ , written  $\Gamma_i^+(t)$  (resp.  $\Gamma_i^-(t)$ ), is the geodesic straightening of the path obtained by following the leaf  $\ell_i$  from  $\Omega$ , for a length  $t$ , then making a right (resp. left) turn to follow a leaf of the other foliation — all the way to infinity, or to another singularity.

In particular,

$$(4.1) \quad \Gamma_i^\pm(0) = \ell_{i\pm 1} \quad \text{and} \quad \Gamma_i^\pm(+\infty) = \ell_i \quad \text{and} \quad \bar{\tau}(\Gamma_i^+(t)) = \bar{\tau}(\Gamma_i^-(t))$$

where the last identity follows from Fact B. Also,  $Q(\Gamma_i^+(t))$  and  $Q(\Gamma_i^-(t))$  are consecutive quadrants for any  $0 < t < +\infty$ .

**Definition 4.2.** A singularity  $p$  of  $\overline{\Sigma}$  is called a *ruling singularity* if the geodesic from  $\Omega$  to  $p$  consists of a single straight Euclidean segment  $[\Omega, p]$ , equal to the diagonal of a singularity-free rectangle. We also call  $[\Omega, p]$  a *ruling segment*.

**Proposition 4.3.** *Let  $\mathcal{F}$  be a fiber of  $\bar{\tau}$ , not containing  $\underline{\Omega}$ . One of the following holds:*

- $Q(\mathcal{F})$  is a single quadrant;
- $|Q(\mathcal{F})| = 2$  and there exists a unique  $(i, t) \in \mathbb{Z} \times \mathbb{R}_+^*$  such that  $\Gamma_i^\pm(t) \in \mathcal{F}$ ;
- $|Q(\mathcal{F})| = 3$  and there exists a unique  $(i, t, t') \in \mathbb{Z} \times \mathbb{R}_+^* \times \mathbb{R}_+^*$  such that  $\{\Gamma_i^\pm(t), \Gamma_{i+1}^\pm(t')\} \subset \mathcal{F}$ ; this happens exactly when  $\mathcal{F}$  contains a ruling segment  $[\Omega, p]$ , inscribed in a singularity-free rectangle of sidelengths  $t, t'$ .

*Proof.* By Fact B, a fiber  $\mathcal{F}$  may have cardinality 1, 2, or  $\infty$ . If  $|\mathcal{F}| = 1$  we are in the first case.




FIGURE 8. An ideal triangle in  $\bar{\Sigma}$ , with one vertex  $\Omega$  and the opposite edge equal to a horizontal leaf. We also draw the geodesic from  $\Omega$  to  $x$  (a nonsingular point on that leaf). In Figures 8–9–10–11, the shaded area is a portion of  $\bar{\Sigma}$  containing no singularities.

Suppose now that  $\mathcal{F}$  contains exactly two elements  $\gamma, \gamma'$ . By Fact B, the paths  $\gamma$  and  $\gamma'$  are then geodesic representatives of two paths that coincide up to a point  $x$  at which they shoot off on opposite rays of a leaf of  $\lambda^+$  or  $\lambda^-$  containing no singularity. We may assume this leaf is horizontal. The ideal triangle spanned by the endpoints of  $\gamma, \gamma'$  and by their common origin  $\Omega$  has one fully horizontal edge, and two edges whose directions (followed towards  $\Omega$ ) depart monotonically from horizontal until they merge at a singular point  $p$ , and then continue on together (with arbitrary changes of direction) until the point  $\Omega$ . See Figure 8. The total




FIGURE 9. Illustration of Proposition 4.3 when  $|\mathcal{F}| = 2$ .

change of direction of  $\gamma$  between  $p$  and its endpoint at infinity of  $\bar{\Sigma}$ , plus the total change of direction of  $\gamma'$  between  $p$  and its endpoint at infinity of  $\bar{\Sigma}$ , is less than  $\pi$ .

If the initial quadrants  $Q(\gamma)$  and  $Q(\gamma')$  at the point  $\Omega$  are distinct, then they are consecutive and the only possibility is that  $p$  coincides with  $\Omega$  and the segment  $[p, x]$  can be taken vertical along a leaf  $l_i$ . We then have the situation of Figure 9: the number  $t$  is just the vertical distance  $px$ , and  $\mathcal{F} = \{\Gamma_i^+(t), \Gamma_i^-(t)\} = \{\gamma, \gamma'\}$ ; the second case of Proposition 4.3 holds.

It remains to treat the case  $|\mathcal{F}| = \infty$ . By Fact B,  $\mathcal{F}$  then contains exactly:

- one path  $\gamma$  terminating at a singularity  $p$ , and
- all the paths  $\hat{\gamma}$  obtained from  $\gamma$  by tacking on a leaf  $l_s$  of  $\lambda^+$  or  $\lambda^-$  issued from  $p$  (here  $s$  ranges over  $\mathbb{Z}$ , the order being clockwise as seen from  $p$ ).




FIGURE 10. Possibilities for the path  $\hat{\gamma}$  include  $\gamma', \gamma'', \hat{\gamma}, \check{\gamma}$ .

If  $|Q(\mathcal{F})| > 1$ , then in particular one of these paths,  $\gamma'$ , satisfies  $Q(\gamma') \neq Q(\gamma)$ . It follows that  $\{\gamma, \gamma'\} = \{\Gamma_i^+(t), \Gamma_i^-(t)\}$  for some  $(i, t) \in \mathbb{Z} \times \mathbb{R}_+^*$ : the argument is similar to the case  $|\mathcal{F}| = 2$  above, except that the horizontal leaf through  $x$  terminates at  $p$  (Figure 10, ignoring for the moment the paths labelled  $\hat{\gamma}, \check{\gamma}, \gamma''$ ).


It remains to find the quadrants  $Q(\hat{\gamma})$  for the other elements  $\hat{\gamma}$  of the fiber  $\mathcal{F}$ , obtained by tacking on to  $\gamma$  a leaf  $l_s$  of  $\lambda^\pm$  issued from  $p$ . Only four paths  $\hat{\gamma} \in \mathcal{F} \setminus \{\gamma\}$  have geodesic representatives that *do not go* through  $p$ : these are the paths  $\gamma', \gamma'', \hat{\gamma}, \check{\gamma}$  shown in Figure 10. These possibilities correspond to tacking on to  $\gamma$  any one of four consecutive leaves  $l_{s_0-1}, l_{s_0}, l_{s_0+1}, l_{s_0+2}$ :

- either a boundary leaf (say  $l_{s_0+1}$  or  $l_{s_0}$ ) of the quadrant at  $p$  containing the last segment of  $\gamma$  (this yields geodesic representatives  $\gamma'$  and  $\gamma''$ );
- or one of the next closest leaves  $l_{s_0+2}, l_{s_0-1}$  (this yields  $\hat{\gamma}$  and  $\check{\gamma}$ ).

All other possible  $\hat{\gamma}$  go through  $p$ , and in particular satisfy  $Q(\hat{\gamma}) = Q(\gamma)$ . Closer inspection shows that actually  $Q(\hat{\gamma})$  and  $Q(\check{\gamma})$  are equal to  $Q(\gamma)$ , too (see Figure 10). The remaining possibility,  $\hat{\gamma} = \gamma''$ , can give two outcomes:

- If there is no ruling segment  $[\Omega, p]$  (Figure 10), then  $Q(\gamma'') = Q(\gamma)$  and we are still in the second case of Proposition 4.3.
- If  $[\Omega, p]$  is a ruling segment, then  $Q(\gamma'') \neq Q(\gamma)$ . This situation is portrayed in Figure 11, and corresponds exactly to the third case of Proposition 4.3.

This concludes the proof of Proposition 4.3.  $\square$

FIGURE 11. When  $[\Omega, p]$  is a ruling segment.

**4.2. Color interfaces of the Cannon-Thurston map.** We will use Proposition 4.3 to prove Theorem 1.2 concerning the combinatorics of the Cannon-Thurston map  $\bar{\tau}$ . Here is a first step.

**Proposition 4.4.** Recall  $\overline{\mathbb{R}^+} := [0, +\infty]$ . For all  $i \in \mathbb{Z}$ , the map

$$J_i : \begin{array}{ccc} \overline{\mathbb{R}^+} & \longrightarrow & \mathbb{P}^1\mathbb{C} \\ t & \longmapsto & \bar{\tau}(\Gamma_i^+(t)) \end{array}$$

is continuous, injective on  $[0, +\infty)$ , and  $J_i(\overline{\mathbb{R}^+})$  is a Jordan curve.

(In this proposition, the image of  $J_i$  is the curve labelled “ $J_i$ ” in Figure 3. The  $t$ -hook  $\Gamma_i^+(t)$  in the definition of  $J_i$  could be replaced by  $\Gamma_i^-(t)$ , due to (4.1) in Definition 4.1.)

*Proof.* We first prove continuity. To fix ideas, suppose the leaf  $\ell_i$  shoots off from  $\Omega$  vertically, upwards. We write  $\Gamma(t)$  for  $\Gamma_i^+(t)$ .

If  $t \in (0, +\infty)$  and  $\Gamma(t)$  does not terminate on a singularity, then the geodesic representative of  $\Gamma(t)$  makes infinitely many turns at singularities  $p_1, p_2, \dots$ . Continuity of  $J_i$  at  $t$  then follows from continuity of  $\bar{\tau}$ : indeed, for any integer  $N$ , if  $t'$  is close enough to  $t$  then the geodesic representative of  $\Gamma(t')$  will coincide with that of  $\Gamma(t)$  at least up to  $p_N$ .

If  $t \in (0, +\infty)$  and  $\Gamma(t)$  terminates on a singularity  $p$ , consider  $t'$  very close to  $t$ . Suppose first that  $t' < t$ . Let  $\Gamma(t^-)$  be the path obtained from  $\Gamma(t)$  by tacking on a leaf of  $\lambda^\mp$  making an angle  $\pi$  with  $\Gamma(t)$  below  $p$ . Since  $\bar{\tau}(\Gamma(t)) = \bar{\tau}(\Gamma(t^-))$  by Fact B, it is enough to prove that  $\Gamma(t')$  approaches  $\Gamma(t^-)$  in the space of paths as  $t'$  approaches  $t$  from below. This is the case, because the geodesic straightening of  $\Gamma(t^-)$  again goes through infinitely many singularities  $p_1, p_2, \dots$  (not including  $p$ ), and agrees with that of  $\Gamma(t')$  up to any given  $p_N$  provided  $t - t' > 0$  is small enough.



In the case  $t' > t$ , define similarly  $\Gamma(t^+)$ , the path obtained from  $\Gamma(t)$  by tacking on a leaf of  $\lambda^\mp$  making an angle  $\pi$  with  $\Gamma(t)$  above  $p$ . This time,  $p$  is the final turn of  $\Gamma(t^+)$ . Since  $\bar{\tau}(\Gamma(t)) = \bar{\tau}(\Gamma(t^+))$  by Fact B, it is enough to prove that  $\Gamma(t')$  approaches  $\Gamma(t^+)$  as  $t'$  approaches  $t$  from above. This holds true because the next turn of  $\Gamma(t')$  after  $p$  lies arbitrarily far out in the horizontal direction if  $t' - t > 0$  is small enough.

If  $t = 0$ , the argument is similar to the case  $t' > t$  just treated. If  $t = +\infty$ , just observe that for very large  $t'$ , the geodesic straightening of  $\Gamma(t')$  starts with an arbitrarily long, nearly vertical segment, so again  $\Gamma(t')$  approaches  $\Gamma(+\infty) = \ell_i$  in the space of paths as  $t' \rightarrow +\infty$ .

Next, the identity  $J_i(0) = J_i(+\infty) = \bar{\tau}(\underline{\Omega})$  is clear by Fact B. Finally, we check the injectivity properties of  $J_i$ . Note that, after straightening,  $\Gamma(t)$  and the other  $t$ -hook  $\Gamma'(t)$  along  $\ell_i$  (with the same initial segment) start into different quadrants:  $Q(\Gamma(t)) \neq Q(\Gamma'(t))$ . Therefore, Proposition 4.3 applies, particularly the uniqueness of  $t$ .  $\square$

**4.3. Subdivision of quadrants.** In a given quadrant at  $\Omega$ , bounded by rays  $\ell_i$  and  $\ell_{i+1}$  and coordinatized by  $(\mathbb{R}^+)^2$ , the ruling singularities  $\{p_s^i\}_{s \in \mathbb{Z}}$  form a naturally ordered  $\mathbb{Z}$ -sequence, with vertical and horizontal coordinates varying monotonically in opposite directions. To see this, one may for example consider for every  $t > 0$  the initial length- $t$  segment of the leaf  $\ell_i$ , and push this segment in the direction of  $\ell_{i+1}$  until it bumps into a singularity. The area swept out is a rectangle, and the union of all these rectangles for all  $t > 0$  forms a “staircase”, the  $p_s^i$  being the turning points at the back of each stair. See Figure 12. The indexing convention is such that the path  $[\Omega, p_s^i]$  converges to  $\ell_{i+1}$  (resp.  $\ell_i$ ) as  $s \rightarrow +\infty$  (resp.  $s \rightarrow -\infty$ ). This defines  $s$  only up to an additive shift depending on  $i$ , but we make no attempt to harmonize these shifts (i.e. any choice will do).




FIGURE 12. A staircase in the  $i$ -th quadrant, with ruling segments  $[\Omega, p_s^i]$ . We mark the directions for initial segments of paths in the subinterval  $\mathcal{I}_s^i$  of  $\mathbb{S}$ , as well as  $\mathcal{I}^i = \bigcup_{s \in \mathbb{Z}} \mathcal{I}_s^i$ .

**Definition 4.5.** Inside the circle  $\mathbb{S}$  of geodesic paths of  $\bar{\Sigma}$  issued from  $\Omega$ , we let:

- $\mathcal{I}^i$  be the closed interval of paths between the rays  $\ell_i$  and  $\ell_{i+1}$ ;

- $\mathcal{I}_s^i \subset \mathcal{I}^i$  be the closed interval of paths between the ruling segments  $[\Omega, p_s^i]$  and  $[\Omega, p_{s+1}^i]$ .

**Remark 4.6.** An important feature is that the order on the  $\mathcal{I}_s^i$  induced by the cyclic order on  $\mathbb{S}$  is the lexicographic order on pairs  $(i, s)$ . In particular, fixing  $i \in \mathbb{Z}$ ,

- $\Gamma_i^-(t) \in \mathcal{I}_s^{i-1}$  for an index  $s$  that is a nondecreasing function of  $t$ ;
- $\Gamma_i^+(t) \in \mathcal{I}_s^i$  for an index  $s$  that is a nonincreasing function of  $t$ .

Recall the Jordan maps  $J_i = \bar{\tau} \circ \Gamma_i^\pm : \overline{\mathbb{R}^+} \rightarrow \mathbb{P}^1\mathbb{C}$  from Proposition 4.4. The following proposition classifies the intersection points of the Jordan curves  $J_i(\overline{\mathbb{R}^+})$ .

**Proposition 4.7.** For  $(i, t)$  and  $(i', t')$  distinct elements of  $\mathbb{Z} \times \overline{\mathbb{R}^+}$ , the relationship  $J_i(t) = J_{i'}(t')$  holds if and only if:

- $\{t, t'\} \subset \{0, +\infty\}$ , or
- $|i - i'| = 1$  and the quadrant between  $\ell_i$  and  $\ell_{i'}$  contains a ruling singularity  $p$  at coordinates  $(t, t')$ .

*Proof.* For the “if” direction, the first case is again the characterization of the  $\bar{\tau}$ -fiber of  $\underline{\Omega}$  by Fact B. The second case follows similarly from the characterization of the  $\bar{\tau}$ -fiber of  $p$  (more precisely, of the path represented by the ruling segment  $[\Omega, p]$ ).

For the “only if” direction, suppose first that  $J_i(t) = J_{i'}(t')$  is the point  $\bar{\tau}(\underline{\Omega})$ : by Fact B, the corresponding hooks are degenerated to full leaves of  $\lambda^\pm$ , i.e.  $t$  and  $t'$  belong to  $\{0, +\infty\}$ . If  $J_i(t) = J_{i'}(t')$  is some other point  $P$ , then we can apply the second and third cases of Proposition 4.3: the fiber  $\mathcal{F} = \bar{\tau}^{-1}(P)$  contains at most two (opposite)  $t$ -hooks along any leaf  $\ell_i$ , and this happens either for one value of  $i$  (in which case  $P$  belongs only to the Jordan curve  $J_i(\overline{\mathbb{R}^+})$ ) or for two consecutive values of  $i$  (in which case  $P$  belongs to the two curves). The latter case arises precisely when  $\mathcal{F}$  contains a ruling segment, and  $(t, t')$  are then its coordinates.  $\square$

**4.4. Proof of Theorem 1.2.** Propositions 4.4 and 4.7 show that two Jordan curves  $J_i(\overline{\mathbb{R}^+}) \subset \mathbb{P}^1\mathbb{C}$  intersect only at  $\infty$  and at a *discrete* subset of  $\mathbb{C}$ , because ruling singularities do not accumulate. More precisely,  $J_i(\overline{\mathbb{R}^+}) \cup J_{i+1}(\overline{\mathbb{R}^+}) \setminus \{\infty\}$  is the boundary of the union of an infinite string of disks  $(\delta_s^i)_{s \in \mathbb{Z}}$ , each disk sharing just one boundary point with the next.

To finish proving Theorem 1.2, it remains to check that the Jordan curves  $J_i(\overline{\mathbb{R}^+})$  never *cross* each other (i.e. they bound nested disks  $D_i$ ), and that the Cannon-Thurston map  $\bar{\tau}$  fills the string of disks  $D_i \setminus D_{i+1}$  in linear order.

First, define  $D'_i$  as the image under  $\bar{\tau}$  of all paths lying weakly clockwise from the ray  $\ell_i$  (and counterclockwise from the trivial path  $\underline{\Omega}$ ):

$$D'_i := \bar{\tau} \left( \overline{\bigcup_{j \geq i} \mathcal{I}^j} \right).$$

**Proposition 4.8.** The set  $D'_i$  is the closure of one complementary component of the Jordan curve  $J_i(\overline{\mathbb{R}^+})$  in the sphere  $\mathbb{P}^1\mathbb{C}$ .

*Proof.* First,  $D'_i$  is closed, as it is the image of a compact interval under a continuous map. Therefore  $D'_i \setminus J_i(\overline{\mathbb{R}^+})$  is closed in  $\mathbb{P}^1\mathbb{C} \setminus J_i(\overline{\mathbb{R}^+})$ ; let us prove that it is also open. Let  $\gamma$  be a nontrivial path lying clockwise from  $\ell_i$ , and not equal to any leaf  $\ell_j$ .

Suppose  $D'$  does not contain any neighborhood of  $\bar{\tau}(\gamma)$ . By surjectivity of the Cannon-Thurston map  $\bar{\tau}$ , this means some paths lying *counterclockwise* from  $\ell_i$  are mapped by  $\bar{\tau}$  arbitrarily close to  $\bar{\tau}(\gamma)$ . Taking limits it implies, by continuity of  $\bar{\tau}$ , that the  $\bar{\tau}$ -fiber  $\mathcal{F}$  of  $\gamma$  contains paths lying weakly counterclockwise from  $\ell_i$ . Such a limit cannot be  $\ell_i$  itself (or  $\underline{\Omega}$ ), since these belong to the fiber  $\bar{\tau}^{-1}(\infty) \neq \mathcal{F}$ . So  $\mathcal{F}$  contains a path lying *strictly* counterclockwise from  $\ell_i$ .

By Proposition 4.3, the fiber  $\mathcal{F}$  contains paths belonging to (the interiors of)  $k$  consecutive quadrants for some  $k \in \{1, 2, 3\}$ . The discussion above implies  $k \geq 2$ . The case  $k = 2$  means, by Proposition 4.3, that some element of  $\mathcal{F}$  is (the geodesic straightening of) a  $t$ -hook along  $\ell_i$ , hence  $\bar{\tau}(\gamma)$  belongs to the Jordan curve  $J_i(\overline{\mathbb{R}^+})$ . By Proposition 4.7, the case  $k = 3$  means that  $\bar{\tau}(\gamma)$  belongs to the intersection of two Jordan curves:  $J_i(\overline{\mathbb{R}^+}) \cap J_{i+1}(\overline{\mathbb{R}^+})$  or  $J_i(\overline{\mathbb{R}^+}) \cap J_{i-1}(\overline{\mathbb{R}^+})$ . In any case,  $\bar{\tau}(\gamma) \in J_i(\overline{\mathbb{R}^+})$ . This proves openness of  $D'_i \setminus J_i(\overline{\mathbb{R}^+})$  in  $\mathbb{P}^1\mathbb{C} \setminus J_i(\overline{\mathbb{R}^+})$ .

To see that  $D'$  is equal to only one side of  $J_i(\overline{\mathbb{R}^+})$ , just remark that there are  $\bar{\tau}$ -fibers all of whose elements start off counterclockwise from  $\ell_i$  (a  $\bar{\tau}$ -fiber, other than that of  $\Omega$ , occupies at most 3 consecutive quadrants by Proposition 4.3).  $\square$

The proposition above implies that  $D'_i$  equals  $D_i$ , the disk bounded by the  $i$ -th Jordan curve  $J_i(\overline{\mathbb{R}^+})$ . The inclusion  $D_i \supset D_{i+1}$  is immediate from the definition of  $D'_i$ . Therefore, by Proposition 4.8, the closure of  $D_i \setminus D_{i+1}$  is the union of a  $\mathbb{Z}$ -sequence of disks  $(\delta_s^i)_{s \in \mathbb{Z}}$ , each having only one (boundary) point  $\bar{\tau}([\Omega, p_s^i])$  in common with the previous one.

**Proposition 4.9.** *For any  $(i, s) \in \mathbb{Z}^2$ , we have  $\bar{\tau}(\mathcal{I}_s^i) = \delta_s^i$ .*

*Proof.* As we just mentioned, the gate  $\delta_s^i \cap \delta_{s-1}^i$  is the image under  $\bar{\tau}$  of the  $s$ -th ruling segment  $[\Omega, p_s^i]$  in the  $i$ -th quadrant.

Writing  $p_s^i = (x_s, y_s)$  the natural coordinates in the quadrant ( $y_s$  is along  $\ell_i$ , and  $x_s$  along  $\ell_{i+1}$ ), the boundary of  $\delta_s^i$  is a Jordan curve that can be broken up into two arcs:

$$\begin{aligned} \bar{\tau}(\Gamma_i^+([y_{s+1}, y_s])) &= J_i([y_{s+1}, y_s]), \\ \text{and } \bar{\tau}(\Gamma_{i+1}^-([x_s, x_{s+1}])) &= J_{i+1}([x_s, x_{s+1}]). \end{aligned}$$

The strategy is now similar to the proof of Proposition 4.8, replacing  $\bigcup_{j \geq i} \overline{\mathcal{I}^j}$  with  $\mathcal{I}_s^i$ . The set  $\bar{\tau}(\mathcal{I}_s^i) \setminus \partial\delta_s^i$  is closed in  $\mathbb{P}^1\mathbb{C} \setminus \partial\delta_s^i$  because  $\mathcal{I}_s^i$  is compact; let us prove that it is also open. Let  $\gamma \in \mathcal{I}_s^i$  be a path. By surjectivity and continuity of  $\bar{\tau}$ , it is enough to prove that if the  $\bar{\tau}$ -fiber  $\mathcal{F}$  of  $\gamma$  does not contain (the geodesic straightening of) a  $t$ -hook (i.e.  $\bar{\tau}(\gamma) \notin \partial\delta_s^i$ ), then  $\mathcal{F}$  is contained in  $\mathcal{I}_s^i$ . Conversely, prove that if  $\mathcal{F}$  intersects two distinct subintervals  $\mathcal{I}_s^i$  and  $\mathcal{I}_{s'}^i$ , then  $\mathcal{F}$  intersects two distinct intervals  $\mathcal{I}^i$  and  $\mathcal{I}^{i''}$ .

If  $\mathcal{F}$  is a singleton, then there is nothing to prove. If  $\mathcal{F}$  consists of two elements, then these are paths  $\gamma$  and  $\gamma'$  asymptotic to the two ends of a foliation leaf  $\ell$  (horizontal, say), as in Figure 8 above: the two geodesics coincide up to a singularity  $p$  (which they both pass on the same side unless  $p = \Omega$ ); then diverge from each other, forming together with  $\ell$  the boundary of a triangle containing no singularity. If  $\gamma' \notin \mathcal{I}_s^i$ , then the initial segments of  $\gamma$  and  $\gamma'$  are distinct, meaning that  $p = \Omega$ . There are several possible situations (Figure 13):

- If the height  $y$  of the horizontal line  $\ell$  is larger than  $y_s$ , then the paths  $\gamma$  and  $\gamma'$  both lie in  $\mathcal{I}_s^i$ . Indeed, the singularity  $p_s$  prevents the geodesic straightening  $\gamma'$  from having an initial segment with slope any larger than  $\frac{y_s}{x_s}$ .




FIGURE 13. Case  $|\mathcal{F}| = 2$ . The shaded regions contain no singularity.

(This will be the actual value if  $y - y_s$  is small enough, but otherwise another singularity can force the slope to be even lower, as in the left panel of Figure 13).

- If  $y \in (y_{s+1}, y_s)$ , then the absence of singularity in the rectangle  $[0, x_{s+1}] \times [0, y_s]$  causes the initial segment of  $\gamma'$  to lie in a another quadrant:  $\mathcal{F}$  intersects a second interval  $\mathcal{I}^{i\pm 1}$  in addition to  $\mathcal{I}^i$ . See the right panel of Figure 13. (In particular,  $\gamma$  and  $\gamma'$  are  $t$ -hooks, with  $t = y$ , and  $\bar{v}(\gamma) \in \partial\delta_s^i$ .)
- The case  $y < y_{s+1}$  is ruled out, as neither  $\gamma$  nor  $\gamma'$  would belong to the subinterval  $\mathcal{I}_s^i$ . This finishes the case  $|\mathcal{F}| = 2$ .

Finally we discuss the case that the fiber  $\mathcal{F}$  is infinite. The elements of  $\mathcal{F}$  are a certain path  $\bar{\gamma}$  terminating at a (not necessarily ruling) singularity  $p$ , and all the paths  $(\gamma_n)_{n \in \mathbb{Z}}$  obtained from  $\bar{\gamma}$  by tacking on a vertical or horizontal leaf issued from  $p$ . Each of these paths belongs to some  $\mathcal{I}_{\mathbf{s}(\gamma_n)}^{i(\gamma_n)}$ , for some integers  $i(\gamma_n)$  and  $\mathbf{s}(\gamma_n)$  determined by the initial segment of (the geodesic straightening of)  $\gamma_n$ . If  $\bar{\gamma} \in \mathcal{F}$  is a ruling segment, we have shown (see Figure 11 and Proposition 4.3) that  $\{i(\gamma_n)\}_{n \in \mathbb{Z}}$  consists of three consecutive integers and the proof is finished. If not, then  $(\mathbf{i}, \mathbf{s})$  is a well-defined function on  $\mathcal{F} = \{\bar{\gamma}\} \cup \{\gamma_n\}_{n \in \mathbb{Z}}$ , and we can assume  $(\mathbf{i}, \mathbf{s})(\bar{\gamma}) \neq (\mathbf{i}, \mathbf{s})(\gamma_n)$  for some  $n \in \mathbb{Z}$ . Since  $\bar{\gamma}$  and  $\gamma_n$  differ only by the final leaf from  $p$ , the situation is analogous to the case  $|\mathcal{F}| = 2$  (Figure 13). The only difference is that the leaf  $\ell$  terminates at a point  $p$ , but the argument is exactly the same (applied to  $\bar{\gamma}, \gamma_n$  instead of  $\gamma, \gamma'$ ): if the initial segments of the two paths fall in distinct subintervals ( $\mathcal{I}_{\mathbf{s}(\gamma_n)}^{i(\gamma_n)} \neq \mathcal{I}_{\mathbf{s}(\bar{\gamma})}^{i(\bar{\gamma})}$ ), then they fall in completely different quadrants:  $\mathbf{i}(\gamma_n) \neq \mathbf{i}(\bar{\gamma})$ .  $\square$

Theorem 1.2 is now proved: the boustrophedonic order of filling of the discs  $\delta_s^i$  derives from the opposite directions of monotonicity in Remark 4.6. In particular, we can speak of the objects defined in Section 1.3: the  $i$ -th furrow  $\bigcup_{s \in \mathbb{Z}} \delta_s^i$ ; cross-furrow and in-furrow edges in  $\partial\delta_s^i$ ; gates  $\delta_s^i \cap \delta_{s \pm 1}^i$ , and spikes.

## 5. PROOF OF THEOREM 1.3

The proofs of the previous section have made it clear that the maximal singularity-free rectangles play an important role in the combinatorics of the Cannon-Thurston map  $\bar{v}$ . Since these rectangles also govern the Agol triangulation by Theorem 1.1,

a result such as Theorem 1.3 should now come as no surprise. In this section we establish the detailed correspondence.

**5.1. The dictionary or Theorem 1.3.(1).** More precisely, we now discuss and illustrate the various entries in the right column (Agol triangulation features) of the dictionary table of Theorem 1.3. To each such feature, we associate first a type of rectangle in the singular surface  $\overline{\Sigma}$ , and then a feature of the Cannon-Thurston tessellation (left column of the dictionary table). These correspondences will be clearly bijective, proving Theorem 1.3.(1).

All figures in this section will obey the following convention:

- Left panel: the Cannon-Thurston tessellation.
- Middle panel: the singular Euclidean surface  $\overline{\Sigma}$ .
- Right panel: the Agol triangulation (more precisely, its vertex link).

Each of the 3 panels may consist of several diagrams showing different possibilities (orientation or color reversals, even/odd quadrants etc): these diagrams are in natural correspondence between the 3 panels.

Note that the left and right panel both live in the (same) complex plane, while the middle panel lives in the singular surface  $\overline{\Sigma}$ . Each diagram in the middle panel is considered up to a  $180^\circ$  rotation, since  $S$  is only a *half*-translation surface. Up to this ambiguity, the figures are designed to systematically capture all possible cases. Note also the following correspondence:

In $\overline{\Sigma}$ (middle panel)	In the link of $\Omega$ (other panels)
Clockwise rotation	Rightwards motion
Vertical direction	Top end
Horizontal direction	Bottom end.

**Vertices (Figure 14).** Ruling segments  $[\Omega, p]$  in  $\overline{\Sigma}$  correspond to vertices of the Agol triangulation by Theorem 1.1, and to vertices of the Cannon-Thurston tessellation by Proposition 4.7.




FIGURE 14. In the middle panel, the green rectangle contains no singularity in its interior and  $\Omega$  is the dark vertex; this convention is upheld in all pictures of Section 5.1.

**Ladderpole edges (Figure 15).** Recall that ladderpole edges in the Agol triangulation are by definition edges connecting two vertices of the same color. These two vertices necessarily come from two consecutive ruling edges of a given quadrant. The paths (of  $\mathcal{I}_s^i$ ) inbetween these two ruling segments are mapped by  $\bar{t}$  to exactly one 2-cell  $\delta_s^i$  of the Cannon-Thurston tessellation, by Proposition 4.9.

Furrows  $\bigcup_{s \in \mathbb{Z}} \delta_s^i$  correspond to full quadrants  $\bigcup_{s \in \mathbb{Z}} \mathcal{I}_s^i = \mathcal{I}^i$  in  $\overline{\Sigma}$ , which correspond to full sequences of ruling singularities  $(p_s^i)_{s \in \mathbb{Z}}$ , and in turn to sequences of ladderpole edges, i.e. ladderpoles, in the Agol triangulation.




FIGURE 15. In the middle panel, we have also indicated ruling singularities in the two quadrants adjacent to the one containing the singularity-free rectangle. These singularities correspond to spikes in the left panel. In the left panel, each arc indicated by an arrow is isotopic to the ladderpole edge seen in the corresponding diagram of the right panel.

**Nonhinge triangles (Figure 16).** A triangle in the Agol triangulation corresponds by definition to a maximal singularity-free rectangle in  $\bar{\Sigma}$  containing exactly one singularity in every edge, one of these singularities being  $\Omega$ . The triangle is non-hinge exactly when the two diagonals of the quadrilateral spanned by the four singularities have slopes of the same sign. Up to symmetry, we may then assume that  $\Omega = (0, 0)$  is the bottom vertex, and that the other three vertices  $T, R, L$  (Top, Right, Left) satisfy:  $x_L < x_T < 0 < x_R$  and  $0 < y_R < y_L < y_T$ , where  $(x_p, y_p)$  are the coordinates of a point  $p$ .

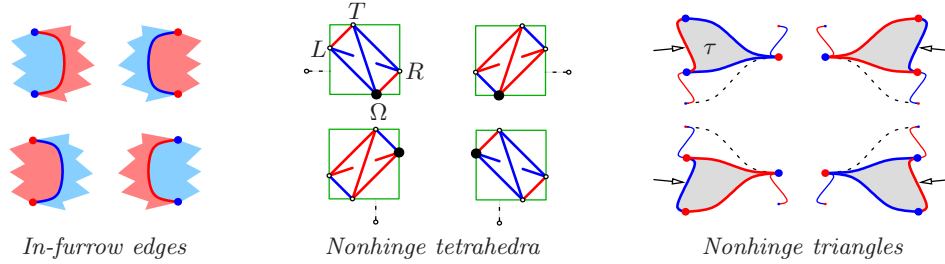


FIGURE 16. In the right panel, the nonhinge triangle under consideration is shaded; we have also depicted the *next* triangle across its base rung. The color of the base rung of *that* second triangle (dashed) is determined by the position, in the middle panel, of the next ruling singularity after  $L$  (also dashed), relative to the height  $y_R$  of  $R$ . Finally, each edge indicated by an arrow in the right panel is isotopic to the edge seen in the corresponding diagram of the left panel.

Note that  $R$  is a ruling singularity of the upper right quadrant at  $\Omega$ , while  $L$  and  $T$  are two consecutive ruling singularities of the upper left quadrant. The ruling staircases in these two quadrants climb from opposite sides to the upwards ray at  $\Omega$ . The inequalities above mean precisely that no ruling singularity in the quadrant of  $R$  has a vertical coordinate in  $[y_L, y_T]$ . By Proposition 4.7, this means that the Jordan curve  $J_i(\overline{\mathbb{R}^+})$ , associated to  $t$ -hooks along the vertical leaf issued from  $\Omega$ , does not intersect  $J_{i+1}(\overline{\mathbb{R}^+})$  between  $J_i(y_L)$  and  $J_i(y_T)$ ; the corresponding

cell  $\tau([\Omega L, \Omega T])$  of the Cannon-Thurston tessellation thus has no spike on this side, but instead has an in-furrow edge (vertical edge)  $J_i([y_L, y_T])$ .

**Remark 5.1.** *This edge “LT” is topologically the same as an edge of the initial nonhinge triangle  $\tau$  of the Agol triangulation: namely the edge connecting the tip of  $\tau$  to the tip of the next triangle across the base rung of  $\tau$ . This edge is indicated by an arrow in Figure 16.*

**Hinge triangles (Figure 17).** Similarly, a hinge triangle in the Agol triangulation

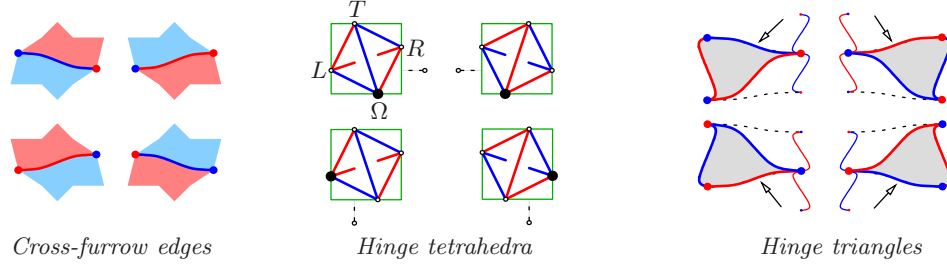


FIGURE 17. In the right panel, the hinge triangle under consideration is shaded; we have also depicted the *next* triangle across its base rung  $\cdot$ . The color of the base rung of *that* second triangle (dashed) is determined by the position, in the middle panel, of the next ruling singularity after  $R$  (also dashed), relative to the height  $y_L$  of  $L$ . Finally, each edge indicated by an arrow in the right panel is isotopic to the edge seen in the corresponding diagram of the left panel.

corresponds in  $\bar{\Sigma}$  (up to symmetries) to a maximal singularity-free rectangle whose edges contain singularities  $(0, 0) = \Omega, L, T, R$  such that  $x_L < x_T < 0 < x_R$  and  $0 < y_L < y_R < y_T$ . Equivalently,  $L, T$  are consecutive ruling singularities in the upper left quadrant and  $R$  is the *highest* of the ruling singularities in the upper right quadrant whose vertical coordinate lies in  $[y_L, y_T]$ . By Proposition 4.7, this means the separator  $J_i(\mathbb{R}^+)$ , followed downwards from the gate  $J_i(y_T)$  of the 2-cell  $\tau([\Omega L, \Omega T])$ , has its first spike at  $J_i(y_R)$ . That is to say, the arc  $J_i([y_R, y_T])$  is a cross-furrow edge.

**Remark 5.2.** *This edge “RT” is topologically the same as an edge of the initial nonhinge triangle  $\tau$  of the Agol triangulation: namely the edge connecting the tip of  $\tau$  to the tip of the next triangle across the base rung of  $\tau$ . This edge is indicated by an arrow in Figure 17.*

**Rungs (Figure 18).** By definition, a rung in the Agol triangulation is an edge connecting two vertices of distinct colors. If we call  $L, R$  the corresponding ruling singularities in  $\bar{\Sigma}$  (belonging to adjacent quadrants), then we can assume up to symmetry that  $y_L < y_R$ . The existence of a singularity-free rectangle circumscribed to  $\{\Omega, R, L\}$  means that if  $L'$  denotes the next higher ruling singularity after  $L$  (in the quadrant of  $L$ ), then  $y_R \in [y_L, y_{L'}]$ . This in turn means the arc  $J_i([y_L, y_{L'}])$ , which connects the two gates of one cell of the Cannon-Thurston tessellation, has a spike at  $J_i(y_R)$ .

All correspondences above are clearly bijective. This proves the first, “dictionary” part of Theorem 1.3.




FIGURE 18. In the right panel, each vertex indicated by an arrow is the vertex of the spike shown in the corresponding diagram of the left panel. Also, in the right panel, we have drawn next to each rung a ladderpole edge (thinner) which is the one corresponding to the 2-cell in the left panel to which the spike belongs.

**5.2. The recipe book or Theorem 1.3.(2).** Remarks 5.1 and 5.2 above, and the bijectivity of the dictionary, imply that the edges of the Cannon-Thurston tessellation are obtained from the Agol triangulation by drawing an edge between the tip of each triangle and the tip of the next triangle across its base rung. This gives the recipe in the first direction of Theorem 1.3.(2).

Call  $\mathcal{L}_i$  the path of tip-to-tip edges thus formed inside the  $i$ -th ladder. The fact that the second recipe in Theorem 1.3.(2) returns the original Agol triangulation can be seen purely at the level of the tessellation itself, based on the fact that  $\mathcal{L}_i$  visits all vertices of the  $i$ -th ladder exactly once, and that between any two consecutive vertices on one ladderpole,  $\mathcal{L}_i$  visits a (possibly empty) sequence of consecutive vertices on the other ladderpole. Theorem 1.3 is proved.

## 6. ILLUSTRATIONS

We finish with some extra illustrations to make the combinatorics of the filling curve  $\bar{\tau}$  and the Cannon-Thurston tessellation more concrete. We will make no effort to draw its fractal-looking edges realistically: the information is still purely combinatorial.

**6.1. Global illustration.** In Figure 19 we show a relatively large sample of the Cannon-Thurston tessellation (left) and of the Agol triangulation (right), for the same underlying combinatorics. The picture shows in interaction many of the local features described and illustrated in Section 5.1.

In the left panel, only the solid colored areas contain the information on the Cannon-Thurston tessellation. However, we overlaid them with the 1-skeleton of the Agol triangulation to help visualize the correspondences. Since all Cannon-Thurston edges are isotopic to Agol edges, each of the former receives a natural color (blue or red) even though it always just separates a blue area from a red one. Also, each ladderpole edge in the Agol triangulation is isotopic to 0, 1, or 2 in-furrow edges of the Cannon-Thurston tessellation: in the resulting overlay (left panel) we draw all 1,2 or 3 edges separately but with the same color. This was also the convention in the left panel of Figure 15.

In short, this synthesis is useful to illustrate Theorem 1.3 and the interplay between the various lines of the “dictionary”. However, it has two drawbacks:

- We cannot draw the corresponding features in the flat surface  $\bar{\Sigma}$ , as the figure would become too crowded.



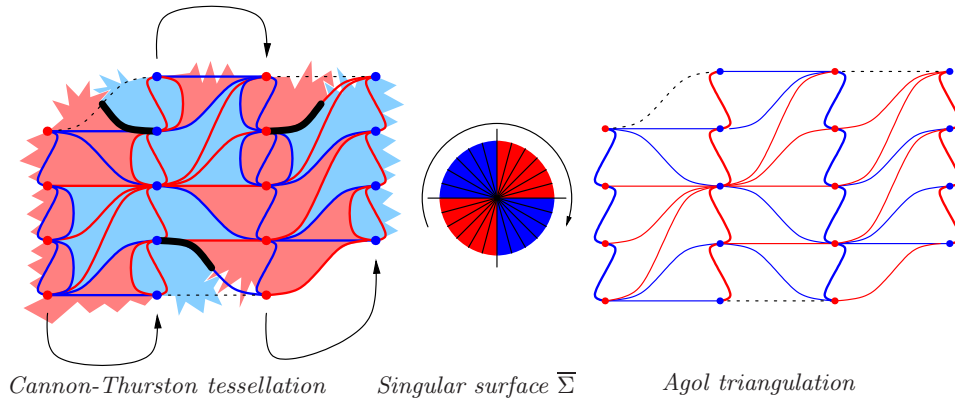


FIGURE 19. Corresponding chunks of the two plane tessellations (the combinatorics in  $\bar{\Sigma}$  are not shown). *Right panel*: the 3 dotted edges have undetermined color. *Left panel*: correspondingly, each of the 3 unfinished (thick black) edges has undetermined endpoint (at either end of the nearby dotted edge).

- In order to show exactly the same information in the two diagrams, we must leave some details ambiguous near the outer boundaries.

For example, forcing a *color* upon the top-left dotted edge in the Agol-triangulation panel would force an *endpoint* (and a color) upon the top-left, unfinished edge of the Cannon-Thurston panel (interrupted here on a dotted arc). Conversely, choosing specific numbers of spikes for the outer cells of the Cannon-Thurston tessellation (presently truncated in a ragged style) would force some partial information upon the adjacent ladders (not pictured) in the Agol triangulation.

The first drawback is unavoidable if we draw too large a portion of the tessellations, but the second is unavoidable as long as we draw only a bounded portion.

**6.2. Semi-local illustration.** Thus, we try to strike a middle ground in Figure 20. This figure shows two full cells of the Cannon-Thurston tessellation (left), along with the corresponding data in  $\bar{\Sigma}$  (middle) that allows to count their spikes, and the corresponding chunk of the Agol triangulation (right). In the middle panel, the whole green area is singularity-free.

The numbers of spikes chosen are 0, 1 in the top cell (on its left and right sides respectively), and 2, 3 in the bottom cell. Up to varying these numbers, one can build the Cannon-Thurston tessellation entirely out of 2-cells like the blue one shown, and red ones obtained by a horizontal reflection and an exchange of colors.

The portrayed cell  $\delta_s$  in the left panel is the image, under the Cannon-Thurston map  $\bar{\tau}$ , of a subinterval  $\mathcal{I}_s^i$  in the *upper left* quadrant of the middle panel. In the right panel, the ladderpole edge corresponding to  $\mathcal{I}_s^i$  is the thick red one, connecting blue vertices. The ruling singularities that we show in the middle panel are:

- the two endpoints of  $\mathcal{I}_s^i$ , mapping to the gates of  $\delta_s$ , and connected to  $\Omega$  in blue (upper left quadrant);
- the ruling singularities in neighboring quadrants that correspond to spikes of  $\delta_s$ , connected to  $\Omega$  in red (their coordinates, along the direction common to the two consecutive quadrants, fall between the endpoints of  $\mathcal{I}_s^i$ );

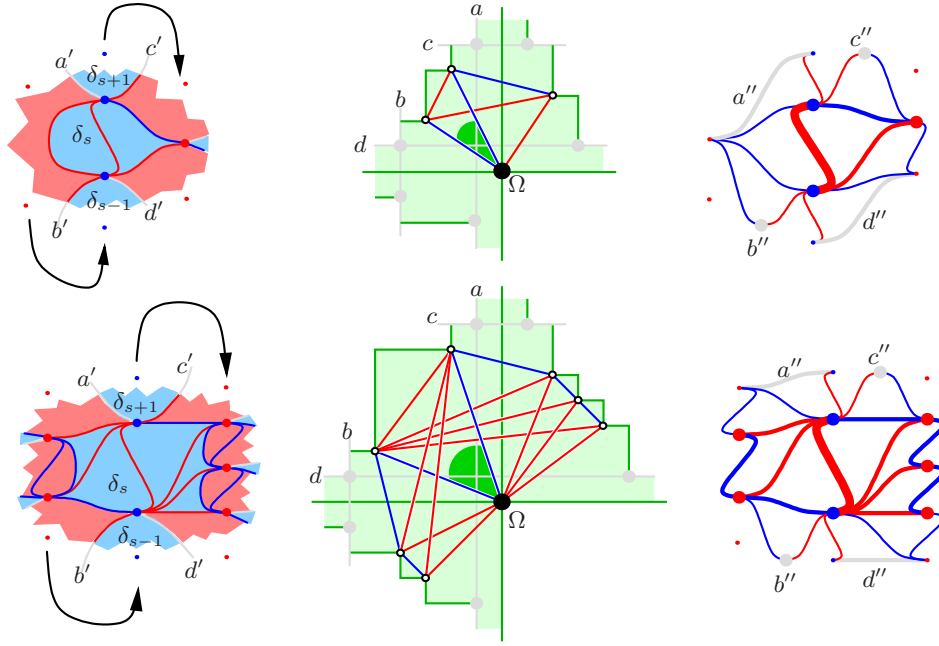


FIGURE 20. The combinatorics associated to full cells (left panel) of the Cannon-Thurston tessellation: the first cell has 0 spikes on one side and 1 on the other; the second cell has 2 and 3. The corresponding subinterval  $\mathcal{I}_s^i$  is shown as a green sector in the middle panel.

- two more ruling singularities (gray) in each of the three quadrants, just before and just after the ones already given. The corresponding points in the left panel are six small, isolated dots.

Indeterminacies about how neighbors of  $\delta_s$  glue up together give rise to corresponding indeterminacies in the other diagrams, of which we keep careful track. Namely, each panel has some features left in gray: saddle point connections  $a, b, c, d$  between gray singularities (middle panel); unfinished edges  $a', b', c', d'$  (left panel); edges  $a'', d''$  and vertices  $b'', c''$  (right panel).

Each of the unfinished edges  $a', b', c', d'$  may end at one of two possible vertices, blue or red. The two possible endpoints for  $a'$  are the two endpoints of the edge  $a''$ ; same for  $d'$  and  $d''$ . The two possible endpoints for  $b'$  are the two possible positions of the vertex  $b''$  (on one of two adjacent ladderpoles, of opposite colors); same for  $c'$  and  $c''$ . The saddle point connections  $a, b, c, d$  are in fact never horizontal or vertical, but we drew them nearly so, to insist that their slope could be of any sign, deciding the endpoints of  $a', b', c', d'$  and colors of  $a'', b'', c'', d''$ . More precisely:

If the saddle point connection  $a$  has negative slope, then the unfinished edge  $a'$  ends at the blue vertex (meaning the cell  $\delta_{s+1}$  above  $\delta_s$  has no spike on the left), and the edge  $a''$  is blue. Conversely, if  $a$  has positive slope, then  $a'$  ends at the red vertex ( $\delta_{s+1}$  has a spike on the left), and  $a''$  is red. The same holds of  $d, d', d''$ , replacing spikes on the left in  $\delta_{s+1}$  with spikes on the right in  $\delta_{s-1}$ .

If the saddle point connection  $b$  has negative slope, then  $b'$  ends at the red vertex ( $\delta_{s-1}$  has a spike on the left) and  $b''$  is red. Conversely, if  $b$  has positive slope, then  $b'$  ends at the blue vertex ( $\delta_{s-1}$  has no spike on the left) and  $b''$  is blue. The same holds of  $c, c', c''$ , replacing spikes on the left in  $\delta_{s-1}$  with spikes on the right in  $\delta_{s+1}$ . To summarize,

$\text{slope}(a) < 0$	$a'$ ends on blue	$a''$ blue
$\text{slope}(a) > 0$	$a'$ ends on red	$a''$ red
$\text{slope}(b) < 0$	$b'$ ends on red	$b''$ red
$\text{slope}(b) > 0$	$b'$ ends on blue	$b''$ blue
$\text{slope}(c) < 0$	$c'$ ends on red	$c''$ red
$\text{slope}(c) > 0$	$c'$ ends on blue	$c''$ blue
$\text{slope}(d) < 0$	$d'$ ends on blue	$d''$ blue
$\text{slope}(d) > 0$	$d'$ ends on red	$d''$ red.

Note that the edges  $b'$  and  $c'$  are necessarily red, while  $a'$  and  $d'$  could be any color (the color opposite to that of their unknown endpoint). On the other hand, the vertex  $b''$  (resp.  $c''$ ), if blue, is an endpoint of the edge  $d''$  (resp.  $a''$ ).

**6.3. Further subdivisions of the Cannon-Thurston map.** In Figure 21, we show a schematic diagram of the order in which a cell  $\delta_s^i = \tau(\mathcal{I}_s^i)$  of the Cannon-Thurston tessellation is itself filled out. This involves two consecutive ruling singularities  $p_s, p_{s+1}$  of the top right quadrant, as well as a maximal sequence of  $k \geq 2$  singularities such that any two consecutive of them, together with  $p_s$  and  $p_{s+1}$ , span a maximal singularity-free rectangle with  $[p_s, p_{s+1}]$  as the bottom left edge. We show the cases  $k = 2, 3, 6$ .

Note that for each path  $\gamma$  that terminates at a singularity, such as the paths numbered 0, 3, 8, 11 in the case  $k = 2$  (top), the curve  $\bar{\tau}$  actually goes infinitely many times through  $\tau(\gamma)$  in a neighborhood of  $\gamma$  (we just draw one pass for simplicity). Indeed, near  $\tau(\gamma)$ , the trajectory of  $\bar{\tau}$  is actually the image of the full boustrophedon under a Möbius map, with  $\tau(\gamma)$  playing the role of the point at infinity. This also explains why so-called spikes do look “spiky”, in the sense that they are pinched between two tangent circles, the Möbius image of a pair of furrow-parallel lines.

#### REFERENCES

- [1] I. Agol, *Ideal Triangulations of Pseudo-Anosov Mapping Tori*, arXiv:1008.1606, in *Contemp. Math.* 560, Amer. Math. Soc., Providence, RI, (2011), 1–17
- [2] N. Bonichon, C. Gavaille, N. Hanusse, L. Perković, *Tight stretch factors for  $L^1$ - and  $L^\infty$ -Delaunay triangulations*, *Computational Geometry* **48**-3 (2015), 237–250
- [3] B.H. Bowditch, *The Cannon-Thurston map for punctured-surface groups*, *Mathematische Zeitschrift* **255**-1 (2007), 35–76
- [4] J.W. Cannon, W. Dicks, *On hyperbolic once-punctured-torus bundles II: fractal tessellations of the plane*, *Geom. Dedicata* 123 (2006), 11–63
- [5] W. Dicks and M. Sakuma, *On hyperbolic once-punctured-torus bundles III: comparing two tessellations of the complex plane*, *Topology Appl.* **157** (2010), 1873–1899
- [6] J.W. Cannon, W.P. Thurston, *Group Invariant Peano Curves*, *Geometry and Topology* 11 (2007), 1315–1356
- [7] A. Fathi, F. Laudenbach, V. Poénaru, *Travaux de Thurston sur les surfaces*, *Astérisque* 66–67, Soc. Math. France (1991, 2nd ed.), 286 p.
- [8] D. Futer, F. Guéritaud, *Explicit angle structures for veering triangulations*, *Algebraic & Geometric Topology* **13** (2013), 205–235

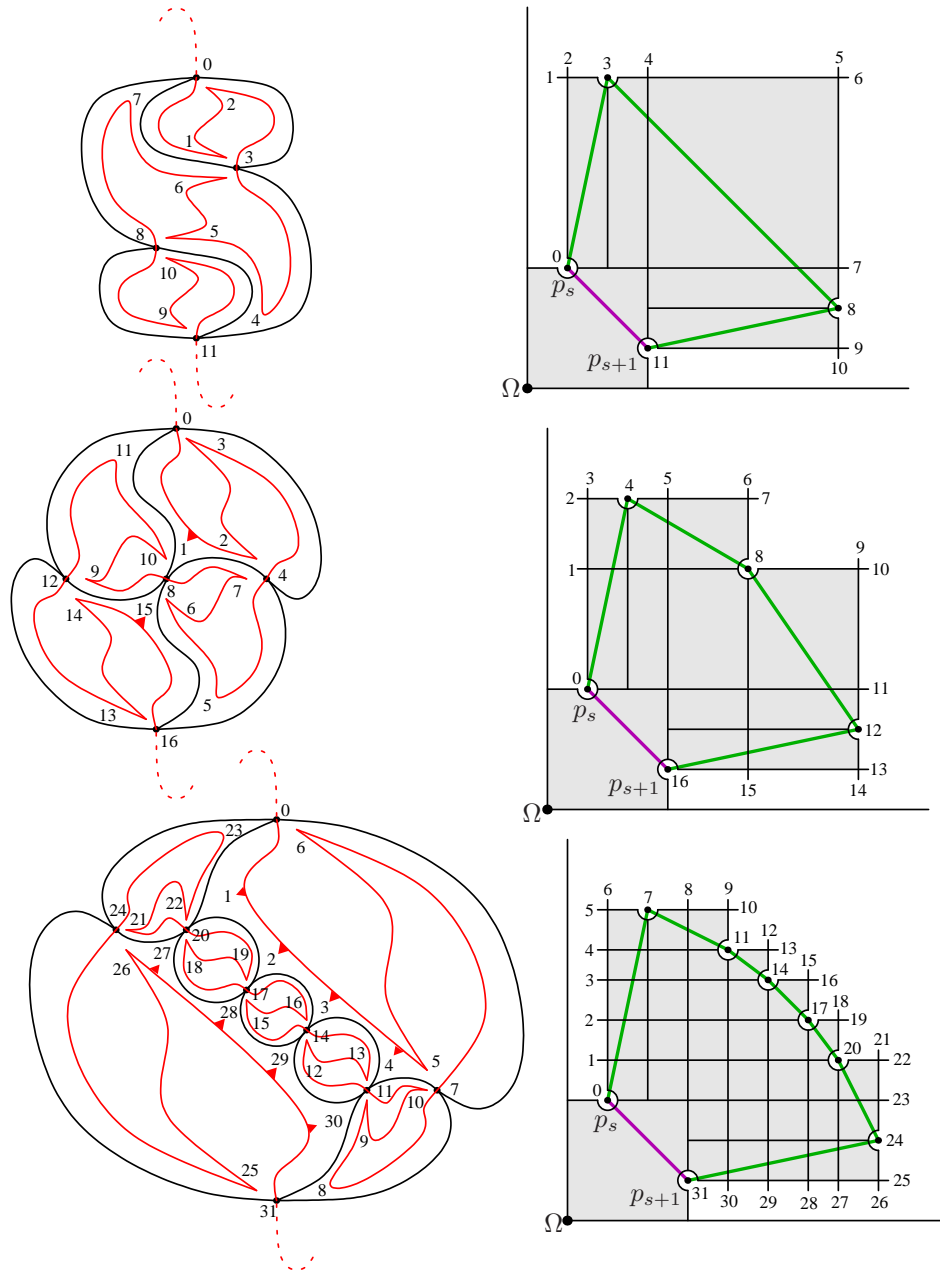


FIGURE 21. Left: the filling curve  $\tau$  (red), filling a disk  $\delta_s^i$ . Right: the paths issued from  $\Omega$  and crossing the purple segment  $[p_s, p_{s+1}]$  can be partitioned according to which green segment they cross next;  $\tau$  maps this partition to a partition of  $\delta_s^i$  into smaller disks. A form of recursion then applies (we show only the first step).

[9] F. Guéritaud, *Veering triangulations and the Cannon-Thurston map*, pp 1419–1420 in: W. Jaco, F. Lutz, D. Gómez-Pérez, J. Sullivan, *Triangulations*, Oberwolfach Reports **9**, Issue 2 (2012), 1405–1486.

- [10] C.D. Hodgson, A. Issa, H. Segerman, *Non-geometric veering triangulations*, Experimental Mathematics **25**, no. 1 (2016), 17–45
- [11] C.D. Hodgson, J.H. Rubinstein, H. Segerman, S. Tillmann, *Veering triangulations admit strict angle structures*, Geometry & Topology **15** (2011), 2073–2089
- [12] V.A. Klyachin, *On a Generalization of the Delaunay Condition*, Vestn. Tomsk. Gos. Univ. Mat. Mekh. 2008 no. 1(2), 48–50 (Russian)
- [13] C.T. McMullen, *Local connectivity, Kleinian groups and geodesics on the blowup of the torus*, Invent. Math. **146** (2001), 35–91
- [14] Mahan Mj, *Cannon-Thurston Maps for Surface Groups*, Annals of Mathematics **179**-1 (2014), 1–80
- [15] J.-P. Otal, *Le théorème d'hyperbolisation pour les variétés fibrées de dimension 3*, SMF – Astérisque 235 (1996), 159 pages

CNRS AND UNIVERSITÉ LILLE 1, LABORATOIRE PAUL PAINLEVÉ, 59655 VILLENEUVE D'ASCQ  
CEDEX, FRANCE

WOLFGANG-PAULI INSTITUTE, UNIVERSITY OF VIENNA, CNRS-UMI 2842, AUSTRIA

*E-mail address:* francois.gueritaud@math.univ-lille1.fr