



Nicolas Curien · Wendelin Werner

The Markovian hyperbolic triangulation

Received July 14, 2011 and in revised form April 3, 2012

Abstract. We construct and study the unique random tiling of the hyperbolic plane into ideal hyperbolic triangles (with the three corners located on the boundary) that is invariant (in law) with respect to Möbius transformations, and possesses a natural spatial Markov property that can be roughly described as the conditional independence of the two parts of the triangulation on the two sides of the edge of one of its triangles.

1. Introduction

The study of the scaling limit of critical two-dimensional discrete models from statistical physics has given rise to various random objects in the continuum that combine conformal invariance with a “spatial Markov property” that is inherited from the locality of the interactions in the discrete models (one can think of course about Schramm’s SLE processes [11]).

In the present paper we shall exhibit and study a special Möbius-invariant random triangulation of the Poincaré disk \mathbb{D} endowed with its hyperbolic complex structure, which has a certain spatial Markov property. Let us first very briefly explain what type of triangulations we have in mind. A (hyperbolic) triangle T will be determined by its three corners, always taken on $\partial\mathbb{D}$, and T will be the “inside” of the three hyperbolic lines joining these three points (recall that these hyperbolic lines are circular chords when viewed in the Euclidean setting). We say that \mathbf{T} is a *complete hyperbolic triangulation* of \mathbb{D} if it is a disjoint collection of such triangles, and if the complement of the union of all these triangles has zero hyperbolic measure.

We say that a random triangulation \mathbf{T} is *Möbius-invariant* if its law is invariant under all conformal transformations from \mathbb{D} onto itself. In other words, for any Möbius transformation ϕ of the unit disk onto itself, the law of $\phi(\mathbf{T})$ is the same as that of \mathbf{T} . Our main statement can be described as follows:

N. Curien: DMA-École Normale Supérieure, 45 rue d’Ulm, 75230 Paris Cedex 05, France;
e-mail: nicolas.curien@ens.fr

W. Werner: Laboratoire de Mathématiques Université Paris-Sud, Bât. 425, 91405 Orsay Cedex, France; e-mail: wendelin.werner@math.u-psud.fr

Mathematics Subject Classification (2010): Primary 60D05; Secondary 30F45, 60D05

There exists a unique random complete Möbius-invariant triangulation of \mathbb{D} that has a spatial Markov property that can be loosely speaking described as follows: Given a triangle $T = (abc)$ in this triangulation \mathbf{T} (in fact the rigorous statement is that T is the triangle that contains the origin in \mathbf{T}), the restrictions of the triangulation \mathbf{T} to the three connected components of the complement of T in \mathbb{D} are conditionally independent, and moreover, the part that is beyond (bc) is independent of the position of a .

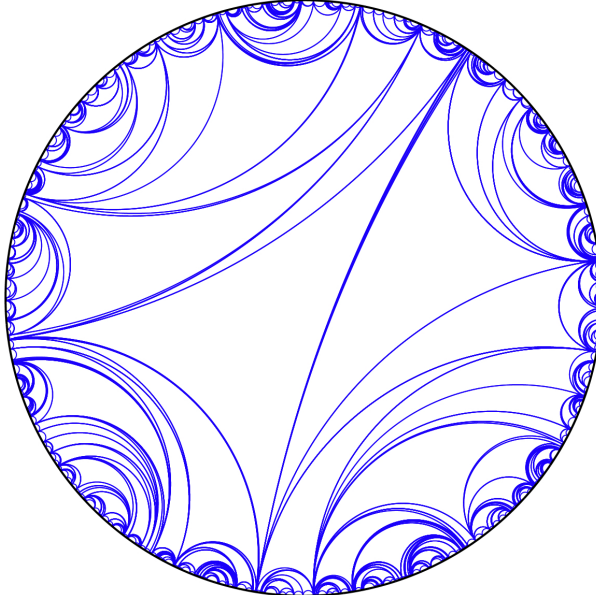


Fig. 1. Sample of our triangulation in the disk.

This will be stated more rigorously in Theorem 2. Uniqueness means of course here uniqueness of the law of the triangulation. Heuristically, the spatial Markov property means that there is conditional independence of both sides of an edge in the triangulation, so that the role of the edges in our triangulations is reminiscent of that of an interface in a nearest-neighbor interaction model from statistical physics.

Discrete models, such as triangulations of convex polygons, have been thoroughly studied in combinatorics, physics or geometry. Some triangulations of the disk can be viewed as continuous counterparts to these discrete models, and various random triangulations of the disk have been defined and studied, in particular in recent years (see for instance [1, 8] and the references therein). However, the particular random triangulation that we construct and study in the present paper is different (we would like to stress that it is *not* the same as the uniform triangulation defined by Aldous [1] that can be viewed as the scaling limit of the uniform triangulation of an n -gon; we shall for instance see that our triangulation is much thinner), and despite its rather striking properties it does not seem (to our knowledge) to have been studied before.

In order to help getting a feeling of what is going on, let us provide a brief heuristic discussion. Assume that a random triangulation \mathbf{T} is complete, Möbius-invariant and Markovian. We can first sample the triangle $T(0)$ of \mathbf{T} that contains the origin—and it will be easy (see Section 2) to identify its law from the conformal invariance and completeness of \mathbf{T} . We write a , b and c for the three apexes of $T(0)$ ordered anti-clockwise. Then, we can start exploring the three pieces of the complement of $T(0)$ independently, because of the spatial Markov property. A first naive guess is that the edge (bc) will also be one of the edges of another triangle of \mathbf{T} , which “neighbors” $T(0)$. We can wonder what the conditional law of its third corner a' will be. Conformal invariance (and the fact that a' is conditionally independent of a given (bc)) imposes that this (conditional) law is invariant under all Möbius transformations of the disk that fix b and c . But all (non-zero) measures supported on the arc of $\partial\mathbb{D}$ between b and c that are invariant under all these transformations necessarily have an infinite mass (in fact, they are multiples of the image of the measure $\xi(dx) := dx/x$ on \mathbb{R}_+ under any fixed Möbius transformation that maps the upper half-plane onto the unit disk, and 0 and ∞ to b and c respectively), and more precisely this infinite mass lies in the neighborhood of b and of c . So, this attempt to construct a neighboring triangle in a Möbius-invariant way fails, but it suggests to those of us who are acquainted with Lévy processes which way to go: When exploring the triangles “outwards” starting from the (bc) boundary of $T(0)$, one will use a Poisson point process, with intensity given by ξ , that will be used at the “time” of choosing each new corner. In particular (because ξ is an infinite measure), almost surely, two different triangles T and T' in the triangulation will never be adjacent (there will always be infinitely many other very thin—in the Euclidean sense—triangles that separate T from T'). In fact, it will turn out that there are not that many triangles either: For any large n , the number of triangles of (Euclidean) width in $[2^{-n-1}, 2^{-n})$ that are separating T and T' is of constant order (it is random but its mean is roughly constant). This explains why one could at first glance think that big triangles can happen to be adjacent to each other by looking at the simulation depicted in Figure 1.

The paper is organized as follows. In Section 2, we collect and derive some rather general or elementary facts, we write down definitions and state our main result, Theorem 2. In Section 3, inspired by the previous heuristic, we show that if \mathbf{T} is a Möbius-invariant complete and Markovian triangulation, then it necessarily corresponds to some Poisson point process that we describe. This argument will prove that the law of \mathbf{T} is *unique*, if it exists. In Section 4, we define explicitly a random triangulation (again, using a Poisson point process), and we check that it indeed has all the required properties, so that *existence* of the Möbius-invariant complete Markovian triangulation follows.

In Section 5, we discuss and state results dealing with Möbius-invariant Markovian tilings of the disk into other hyperbolic polygons than triangles (analogous statements and constructions exist for instance for tilings into conformal squares). In the final section, we make a few comments, and list a couple of open questions.

We are going to assume that the reader is acquainted with basic properties of Möbius transformations on the one hand, and basic knowledge about Poisson point processes, pure jump processes and subordinators (as can be found in [3, 4, 10]) on the other hand. As we are aware that this is not quite a usual mix of backgrounds, we will recall some of the basic features that we will use.

2. Simple preliminaries

2.1. Hyperbolic triangles

We will mostly use the unit disk $\mathbb{D} = \{z \in \mathbb{C} : |z| < 1\}$ to represent the hyperbolic plane. At some point, it will also be convenient to work in the upper half-plane $\mathbb{H} = \{z \in \mathbb{C} : \text{Im}(z) > 0\}$. Throughout the paper, ψ^{-1} will denote the conformal map from \mathbb{H} onto \mathbb{D} defined by $\psi^{-1}(z) = (z - i)/(z + i)$, which maps i to the origin and infinity to 1.

For any pair of *distinct* points a and b on $\partial\mathbb{D}$, we define the hyperbolic line (ab) in \mathbb{D} to be the circular chord in \mathbb{D} that crosses $\partial\mathbb{D}$ orthogonally at both a and b (when $a = -b$, this ‘‘circular chord’’ is in fact a diameter line). In order to avoid confusion, we will use $[a, b]$ to denote straight Euclidean segments.

If we consider three distinct points a, b and c on $\partial\mathbb{D}$, we can define a hyperbolic triangle as the middle open connected component of $\mathbb{D} \setminus ((ab) \cup (bc) \cup (ca))$ (in other words, the connected component of this set that has a, b and c on its boundary). A triangle is thus identified with the unordered set of its three apexes $\{a, b, c\}$. The set of all hyperbolic triangles will be denoted by \mathcal{T} .

We will denote by \mathcal{T}_o the set of all *marked* hyperbolic triangles (it can be viewed as the set of ordered triplets (a, b, c) of distinct boundary points that are ordered anti-clockwise on $\partial\mathbb{D}$). Each hyperbolic triangle corresponds to three marked triangles (one just has to distinguish one apex in order to mark the triangle).

Let us stress that this is a slight abuse of terminology, as our hyperbolic triangles always have their apexes on the boundary of \mathbb{D} (these triangles are called *ideal* in hyperbolic geometry, but since we will not use any other triangles in the present paper, we will simply omit to specify that we always mean ideal triangles). Notice also that with our definition, any triangle is open and has non-empty interior.

Clearly, it is possible to identify the set \mathcal{T}_o of all marked triangles with the group \mathcal{M} of all Möbius transformations (hyperbolic isometries) of the unit disk, i.e., the group of transformations of the type

$$\phi_{z_0, \theta} : z \mapsto \frac{e^{i\theta}(z - z_0)}{\bar{z}_0 z - 1}, \quad \text{where } z_0 \in \mathbb{D} \text{ and } \theta \in [0, 2\pi).$$

Indeed, for each $(a, b, c) \in \mathcal{T}_o$, there exists a unique $\phi \in \mathcal{M}$ (which we can therefore call $\phi_{a,b,c}$) such that $\phi((1, j, j^2)) = (a, b, c)$ (where $j = \exp(2i\pi/3)$ denotes the cubic root of unity). Furthermore, in this identification $\phi_{a,b,c} \leftrightarrow (a, b, c)$, left multiplication by an element ϕ of \mathcal{M} corresponds to taking the image of (a, b, c) under this map, i.e., $\phi \circ \phi_{(a,b,c)} \leftrightarrow (\phi(a), \phi(b), \phi(c))$.

Let us now describe natural measures that one can define on these different sets. Recall first that the hyperbolic metric on \mathbb{D} is defined by

$$\mu = \frac{4 \, dx \, dy}{\pi(1 - x^2 - y^2)^2}$$

which is (up to a multiplicative constant) the unique measure on \mathbb{D} that is invariant under the group \mathcal{M} . Note that all triangles are equivalent up to hyperbolic isometry, so that they

all have the same hyperbolic area. It is easy to check that this area is finite, and we have normalized μ in such a way that the common area of all triangles is equal to 1.

The identification of \mathcal{T}_\circ with the locally compact Lie group \mathcal{M} immediately shows that, up to a multiplicative constant, there exists a unique Haar measure on \mathcal{T}_\circ that is invariant under the group \mathcal{M} (i.e. corresponding to the measures on \mathcal{M} invariant under left multiplication). In other words, there exists a unique Möbius-invariant measure ν_\circ on \mathcal{T}_\circ (up to a multiplicative constant). Recall that \mathcal{M} is unimodular, so that ν_\circ is also invariant under right multiplication. Here are a couple of simple explicit constructions of ν_\circ :

- Consider the product measure $\mu \otimes \lambda$ on $\mathbb{D} \times [0, 2\pi)$, where μ is the hyperbolic measure in \mathbb{D} and λ is the uniform probability measure on $[0, 2\pi)$. Each pair (z_0, θ) in this set defines the isometry $\phi_{z_0, \theta}$ in \mathcal{M} , and it is easy to check that the image measure of $\mu \otimes \lambda$ in \mathcal{M} is invariant under right multiplication. In other words, one can view ν_\circ as the image of $\mu \otimes \lambda$ under the map $(z_0, \theta) \mapsto \phi_{z_0, \theta}((1, j, j^2))$ (note that with this construction, the point $e^{i\theta} z_0$ is the “hyperbolic” center of the triangle); it is easy to check that indeed

$$\mu \otimes \lambda(\{(z_0, \theta) : 0 \in \phi_{z_0, \theta}((1, j, j^2))\}) = 1.$$

- Another way of constructing the Möbius-invariant measure on \mathcal{T}_\circ goes as follows: Define on \mathbb{R}^3 the measure

$$\frac{du \, dv \, dw}{|u - v| |v - w| |w - u|},$$

where we restrict ourselves to the triplets (u, v, w) that are ordered anti-clockwise around $\partial\mathbb{H}$ (i.e. $u < v < w, v < w < u$ or $w < u < v$). Clearly this measure is invariant under the transformations $z \mapsto -1/z, z \mapsto \lambda z$ and $z \mapsto z + z_0$ for $z_0 \in \mathbb{R}$ and $\lambda > 0$. Hence, the image of this measure under ψ^{-1} is a measure on $(\partial\mathbb{D})^3$ (or rather on \mathcal{T}_\circ) that is Möbius-invariant. It is therefore necessarily equal to a multiple of ν_\circ . In fact an explicit computation shows that the multiplicative constant is π^2 .

Similarly, if a measure η on \mathcal{T} is Möbius-invariant, we can observe that the measure on marked triangles obtained by marking one corner uniformly at random among the three, is an invariant measure on \mathcal{T}_\circ , and therefore a multiple of ν_\circ . It follows that η is a multiple of the measure ν obtained from ν_\circ by the natural projection from \mathcal{T}_\circ onto \mathcal{T} .

Recall that the ν_\circ -mass (and therefore the ν -mass also) of the set of all triangles that have the origin in their interior is equal to one. By Möbius invariance, the same holds for the set of all triangles that have a given point z in their interior. In the following, P_z (respectively P_z°) will denote the probability measure on \mathcal{T} (resp. \mathcal{T}_\circ) that is equal to ν (resp. ν_\circ) restricted to those triangles that contain z . The probability measure P_0 will be used repeatedly.

2.2. Möbius-invariant triangulations

A (hyperbolic) triangulation \mathbf{T} of \mathbb{D} is a disjoint collection of hyperbolic triangles of \mathbb{D} . Since every triangle has non-empty interior, such a collection is finite or countable. If \mathbf{T} is a triangulation and $z \in \mathbb{D}$, we define $T(z)$ to be the (unmarked) triangle of \mathbf{T} that

contains z if it exists, and $T(z) = \emptyset$ otherwise. Clearly, if we choose a fixed countable dense family $(z_q)_{q \in Q}$ in \mathbb{D} , then the family $(T(z_q), q \in Q)$ fully describes \mathbf{T} . This gives a way to define a natural sigma-field on the set of all triangulations of \mathbb{D} which we will implicitly use from now on. Note that this sigma-field in fact does not depend on the choice of the dense family (z_q) —indeed, if we order Q and identify Q with \mathbb{N} , then for all $z \in \mathbb{D}$, one has

$$T(z) = \bigcup_{j \geq 0} T(z_j) 1_{z \in T(z_j) \text{ and } z \notin T(z_1) \cup \dots \cup T(z_{j-1})}$$

(here and throughout, $T(z_j) 1_A$ is equal to the triangle $T(z_j)$ when the event A holds and to the empty set otherwise). It is also easy to check that this sigma-field coincides with that associated with the Hausdorff topology on \mathbb{D} (but we will not use this fact).

We say that a triangulation $\mathbf{T} = (T_j)$ is *complete* if the hyperbolic area of $\mathbb{D} \setminus \bigcup_j T_j$ is zero. Most triangulations that we will consider in this paper will be complete. Let us make two side remarks here (they will not be useful in this paper so that we just mention them, leaving the details to the interested reader):

Remark 1. Note that a complete triangulation \mathbf{T} is dense, in the sense that the union of the triangles of \mathbf{T} is dense in \mathbb{D} . However there exist dense triangulations that are not complete (an analogy that one can keep in mind is that there exist open subsets O of $[-1, 1]$ that are dense in $[-1, 1]$, but with Lebesgue measure strictly smaller than 2—such an open subset O then loosely speaking corresponds to the intersection of the interiors of the triangles with $[-1, 1]$).

Remark 2. A *lamination* is a closed subset of \mathbb{D} that can be written as a disjoint union of hyperbolic lines. One says that a lamination is *maximal* if it is maximal for inclusion among laminations. It is not hard to see that the complement of a maximal lamination is composed of disjoint open (ideal) triangles and thus is a hyperbolic triangulation (see [5] for more details on hyperbolic laminations).

We say that a random complete triangulation \mathbf{T} is *Möbius-invariant* if it is invariant (in law) under the action of each Möbius transformation of the unit disk. In words, it is Möbius-invariant if, for any conformal map ϕ from \mathbb{D} onto \mathbb{D} , $\phi(\mathbf{T})$ and \mathbf{T} have the same law.

Suppose now that \mathbf{T} is such a Möbius-invariant random complete triangulation.

By \mathcal{M} -invariance, the quantity $P(T(z) \neq \emptyset)$ is independent of $z \in \mathbb{D}$ and must be equal to 1 by completeness. We can also associate with \mathbf{T} an infinite “counting” measure η on \mathcal{T} as follows: For any measurable set A of triangles in \mathcal{T} , we define $\eta(A)$ to be the expected value of the number of triangles $T_j \in \mathbf{T}$ that fall in A . By \mathcal{M} -invariance of \mathbf{T} it follows that η is a Möbius-invariant measure on \mathcal{T} . Note also that $\eta(0 \in T)$ is equal to the mean number of triangles of \mathbf{T} that contain the origin, which is equal to 1 since \mathbf{T} is almost surely complete. Consequently, one has $\eta = \nu$, and for every $z \in \mathbb{D}$, $T(z)$ is distributed according to P_z .

Of course, it is worth checking if non-trivial Möbius-invariant triangulations exist at all. Here is a construction of the simplest one of all, based on the standard Farey–Ford tiling of \mathbb{D} . Suppose that τ is a given (unmarked) hyperbolic triangle. We construct




Fig. 2. Construction of $\text{Ref}(\tau)$ started from some τ .

deterministically a triangulation $\text{Ref}(\tau)$ containing τ by reflections: It is the only triangulation \mathfrak{T} with the property that for any triangle $T \in \mathfrak{T}$, if ϕ denotes any one of the three Möbius transformations that map $\{1, j, j^2\}$ onto T , then T has exactly three adjacent triangles in \mathfrak{T} that are $\phi(\{1, \omega^1, \omega^2\})$, $\phi(\{\omega^2, \omega^3, \omega^4\})$ and $\phi(\{\omega^4, \omega^5, 1\})$ where $\omega = \exp(i\pi/3)$. It is elementary to check that $\text{Ref}(\tau)$ is well-defined and is a complete hyperbolic triangulation. The triangulation $\mathbf{F} := \text{Ref}(\{1, j, j^2\})$ is called the *Farey–Ford tiling* (see [6, Chapter 8]). One can identify the set of all marked triangles in \mathbf{F} with the discrete subgroup G of \mathcal{M} that leaves \mathbf{F} invariant. In this way, the set of all triangles of $\text{Ref}(\tau)$ is nothing else than the family of all $g(\tau)$, where g spans G (and because we have been using marked triangles to define G , each triangle of $\text{Ref}(\tau)$ appears three times in this list).

Proposition 1. *If T_0 is distributed according to P_0 , then $\text{Ref}(T_0)$ is Möbius-invariant.*

Proof. Note that the knowledge of any triangle in $\text{Ref}(T_0)$ characterizes the entire triangulation. It therefore suffices to prove that if $T(z)$ is the triangle that contains z in this triangulation, it is distributed according to P_z (as this will imply that the law of the triangulation is invariant under any Möbius transformation that maps 0 onto z , because P_z is the image of P_0 under such a hyperbolic isometry).

As ν is invariant under \mathcal{M} , it follows that for each $g \in G$, the measures M_g and M'_g defined on the set of pairs of triangles by

$$M_g(A) = \nu(\{T \in \mathcal{T} : (T, gT) \in A\}) \quad \text{and} \quad M'_g(A) = \nu(\{T' \in \mathcal{T} : (g^{-1}T', T') \in A\})$$

are identical. It follows of course that $\sum_g M_g = \sum_g M'_g$. But, if one restricts ν to those triangles that contain the origin, one obtains P_0 , and furthermore, almost surely, only one unmarked triangle in $\text{Ref}(T_0)$ does contain z , i.e. $T(z) = gT_0$ for exactly three g 's in G , and no other gT_0 's contain z . Hence, it follows that

$$\begin{aligned} 3P(T(z) \in A) &= \sum_g \nu(\{T : 0 \in T, z \in gT, gT \in A\}) \\ &= \sum_g \nu(\{T' : 0 \in g^{-1}T', z \in T', T' \in A\}) = 3P_z(A). \quad \square \end{aligned}$$

2.3. Markovian triangulations

Let us now define the additional Markovian property that we will require of our random triangulations. A first rather weak assumption would be that, conditionally on $T(0)$, the intersections of \mathbf{T} with the three connected components of $\mathbb{D} \setminus T(0)$ are independent. Note that the previous randomized Farey–Ford example has this property (indeed, conditionally on $T(0)$, all other triangles are deterministic).

Our Markovian condition will be stronger. Suppose that we denote the three apexes of $T(0)$ by u_1, u_2 and u_3 ordered anti-clockwise, and the three connected components of $\mathbb{D} \setminus T(0)$ by O_1, O_2 and O_3 in such a way that $u_j \notin \partial O_j$ (see Figure 3). In order to define which of the three apexes is denoted by u_1 , we can for instance just choose it at random among the three.

We will say that a random complete triangulation is *Markovian* if

$$\text{conditionally on } (u_2, u_3), \mathbf{T} \cap O_1 \text{ is independent of } (\mathbf{T} \cap O_2, \mathbf{T} \cap O_3).$$

Note that (because the triangulation is complete), one can recover u_1 from $\mathbf{T} \cap O_2$ and $\mathbf{T} \cap O_3$, so that $\mathbf{T} \cap O_1$ is conditionally independent of u_1 , given (u_2, u_3) .

Note also that if \mathbf{T} is also Möbius-invariant, then the same statement holds for the restriction of \mathbf{T} to the three connected components of the complement of $T(z)$ (for any given fixed point z in \mathbb{D}). We are now ready to state our main result:

Theorem 2. *There exists exactly one (law of a) Markovian Möbius-invariant complete triangulation in \mathbb{D} .*

Until the rest of this section, \mathbf{T} will denote a random Markovian Möbius-invariant complete triangulation, and we will start to study its properties. Let us make a first observation. Define for each $j \in \{1, 2, 3\}$ a conformal transformation ψ_j from \mathbb{D} onto \mathbb{H} that maps O_j onto the domain

$$\mathbb{H}^+ := \{z \in \mathbb{H} : |z| > 1\}$$

(see Fig. 3). We choose a way to define ψ_j that is a deterministic function of O_j (so that ψ_1 does not depend on u_1 etc.)—let us for instance pick ψ_1 so that $\psi_1(u_3) = -1$, $\psi_1(u_2) = 1$ and $|\psi_1'(u_2)| = 1$. In this way, each $\tilde{\mathbf{T}}_j := \psi_j(\mathbf{T} \cap O_j)$ is a triangulation of \mathbb{H}^+ .

The following statement is a consequence (i.e., a reformulation) of our Markovian assumption for Möbius-invariant complete triangulations.

Lemma 3. *The variables $T(0), \tilde{\mathbf{T}}_1, \tilde{\mathbf{T}}_2, \tilde{\mathbf{T}}_3$ are independent, and the latter three have the same distribution. Furthermore, this common distribution σ is invariant under the one-dimensional group of all Möbius transformations ϕ of \mathbb{H} such that $\phi(\mathbb{H}^+) = \mathbb{H}^+$.*

Proof. Let us (for notational convenience) decide that each triangle of \mathbf{T} has been marked at random and independently (this defines $u_1(z), u_2(z), u_3(z), O_1(z), \dots$ for almost all $z \in \mathbb{D}$ because the triangulation is complete in such a way that if $T(z) = T(z')$ then $u_i(z) = u_i(z'), O_i(z) = O_i(z'), \dots$). We will denote by \mathbf{T}_j^z the triangulation \mathbf{T} restricted to $O_j(z), \tilde{\mathbf{T}}_j^z$ its image in \mathbb{H}^+ , and $\theta_1(z)$ will denote the harmonic measure at z of the




Fig. 3. The three maps ψ_1, ψ_2 and ψ_3 .

part of $\partial\mathbb{D}$ between $u_2 = u_2(z)$ and $u_3 = u_3(z)$ that does not contain $u_1 = u_1(z)$. In particular, $2\pi\theta_1(0) \in (0, \pi)$ is simply the angle at the origin of the triangle u_2Ou_3 .

Note that $\tilde{\mathbf{T}}_1, \tilde{\mathbf{T}}_2, \tilde{\mathbf{T}}_3$ are conditionally independent given $T(0)$ (because each $\tilde{\mathbf{T}}_j$ is a deterministic function of $\mathbf{T} \cap O_j$). In order to derive the full independence, it therefore suffices to check that (for each given j), $\tilde{\mathbf{T}}_j$ and $T(0)$ are independent. By symmetry, it is sufficient to consider the case $j = 1$. As $\tilde{\mathbf{T}}_1$ and u_1 are conditionally independent given (u_2, u_3) , it is enough to show that $\tilde{\mathbf{T}}_1$ and (u_2, u_3) are independent. Because of rotational invariance (and because $\tilde{\mathbf{T}}_1$ does not change if one rotates \mathbf{T}), it finally suffices to show that $\tilde{\mathbf{T}}_1$ and $\theta_1(0)$ are independent.

Let F denote a measurable bounded real-valued function on the set of triangulations, and h a measurable bounded function on \mathbb{R} . Then

$$\begin{aligned} E[F(\tilde{\mathbf{T}}_1)h(\theta_1(0))] &= E\left[F(\tilde{\mathbf{T}}_1)h(\theta_1(0)) \times \int_{\mathbb{D}} d\mu(z) 1_{z \in T(0)}\right] \\ &= \int_{\mathbb{D}} d\mu(z) E[F(\tilde{\mathbf{T}}_1^z)h(\theta_1(0)) 1_{0 \in T(z)}] \\ &= \int_{\mathbb{D}} d\mu(z) E[F(\tilde{\mathbf{T}}_1^0)h(\theta_1(z)) 1_{z \in T(0)}] \\ &= E[F(\tilde{\mathbf{T}}_1)] \times \int_{\mathbb{D}} h(\theta_1(z)) 1_{z \in T(0)} d\mu(z), \end{aligned}$$

where we have used the facts that the μ -area of $T(0)$ is one, that $0 \in T(z)$ if and only if $z \in T(0)$, that the triangulation is Möbius-invariant (in particular under the hyperbolic isometry that interchanges z and 0), and finally that $\int_{T(0)} h(\theta_1(z)) d\mu(z)$ is a constant that does not depend on the triangle $T(0)$ (because of Möbius invariance of all quantities involved). This completes the proof of the fact that $\tilde{\mathbf{T}}_1, \tilde{\mathbf{T}}_2$ and $\tilde{\mathbf{T}}_3$ are independent and independent of $T(0)$. They clearly have the same law that we denote by σ .

It remains to show that σ is invariant under all Möbius transformations of \mathbb{H} that map \mathbb{H}^+ onto itself. Recall from Section 2.2 that we know explicitly the distribution

of $T(0)$ which has a smooth density with respect to the Lebesgue measure on $(\partial\mathbb{D})^3$, and we have just seen that $T(0)$ is independent of $\tilde{\mathbf{T}}_1$. Hence, we can say that for any given triangle (a, b, c) that contains the origin, the conditional distribution of $\tilde{\mathbf{T}}_1$ given $T(0) = (u_1, u_2, u_3) = (a, b, c)$ is σ . Suppose that $\Phi : \mathbb{D} \rightarrow \mathbb{D}$ is some Möbius transformation, and define $\mathbf{T}' = \Phi(\mathbf{T})$. If we combine the previous decomposition with the Möbius invariance of \mathbf{T} , we see that for any a, b and c , the conditional distribution of $\tilde{\mathbf{T}}'_1$ given $T'(0) = (\Phi(a), \Phi(b), \Phi(c))$ is still σ as long as this new triangle contains 0.

Let ϕ be a fixed Möbius transformation of \mathbb{H} onto itself such that $\phi(1) = 1$ and $\phi(-1) = -1$. By choosing a, b and c appropriately (for instance $a = -1, b = -i \exp(i\varepsilon)$ and $c = \bar{b}$ with ε very small so that the latter two points are very close to $-i$ and i), we can make sure that if we define $\Phi = \psi_1^{-1} \circ \phi \circ \psi_1$ (where ψ_1 is the isometry defined deterministically using b, c only, that maps these two points onto 1 and -1 , chosen with the same rule as the one we used to define $\tilde{\mathbf{T}}_1$) then both triangles (a, b, c) and $(\Phi(a), \Phi(b), \Phi(c))$ contain the origin. Furthermore, we see that for this particular triple, when $T(0) = (a, b, c)$, the conditional law of $\tilde{\mathbf{T}}'_1$ is that of $\phi(\tilde{\mathbf{T}}_1)$. It follows that the law of $\tilde{\mathbf{T}}_1$ is indeed invariant under ϕ . □

We note that this proves in particular that we can define (in terms of σ) the conditional distribution of $T(z)$ given $T(0)$. We can also note that by Möbius invariance of \mathbf{T} , for each given z , the previous lemma also yields (using the conformal map that swaps 0 and z) a description of the conditional law of the three triangulations corresponding to \mathbf{T} restricted to each of the three connected components of the complement of $T(z)$, in terms of σ .

The next lemma shows that in order to prove uniqueness (in law) of Möbius-invariant complete Markovian triangulations, it suffices to prove that all their two-dimensional marginals are uniquely determined:

Lemma 4. *If for each $z, z' \in \mathbb{D}$ we know the joint law of $(T(z), T(z'))$, then we know the law of the entire triangulation \mathbf{T} .*

Proof. The law of $\mathbf{T} = (T(z), z \in \mathbb{D})$ is characterized by the law of its finite-dimensional marginals, i.e. by the law of $T(Z) := \{T(z_1), \dots, T(z_n)\}$ for all finite sets of points $Z = \{z_1, \dots, z_n\}$ in \mathbb{D} with rational coordinates (see the discussion on sigma-fields at the beginning of Section 2.2).

We say that a finite collection of disjoint triangles in the unit disk is *good* if each connected component of the complement of the union of these triangles in the disk has at most two neighboring triangles in this collection. The left picture of Figure 4 represents a set $T(Z)$ that is not good, because the shaded component neighbors three different triangles of $T(Z)$.

However, because \mathbf{T} is almost surely complete, for each given Z , with probability one, it is possible to add to $T(Z)$ finitely many triangles of \mathbf{T} in order to turn it into a good collection. In fact, there is a minimal way to do this, and we denote by $\tilde{T}(Z)$ the corresponding finite collection of triangles of \mathbf{T} . Note that (for each given Z) $T(Z)$ is a deterministic function of $\tilde{T}(Z)$ (it consists of those triangles in $\tilde{T}(Z)$ that contain a point of Z), so that the distribution of $\tilde{T}(Z)$ contains all the information about the distribution of $T(Z)$.

When Z' is a finite set of points in \mathbb{D} , we say that a finite collection \mathfrak{T} of triangles is in $\mathcal{A}(Z')$ if \mathfrak{T} is good, and if each triangle in \mathfrak{T} corresponds to exactly one point of Z' (i.e. each triangle of \mathfrak{T} contains exactly one point of Z' and each point of Z' is in a triangle of \mathfrak{T}). In particular, if \mathbf{T} is our random triangulation, the event $\{T(Z') \in \mathcal{A}(Z')\}$ holds iff $T(Z')$ is good and if each triangle of $T(Z')$ contains exactly one point of Z' .

Note that if Z is some other given finite family of points, by looking at $T(Z')$ only, we can see whether $\tilde{T}(Z) = T(Z')$. Similarly, we can also check whether this event holds or not by looking at $\tilde{T}(Z)$ only (it suffices to check that $\tilde{T}(Z) \in \mathcal{A}(Z')$). Suppose that for a given Z' , we know the law of $T(Z') \mathbb{1}_{T(Z') \in \mathcal{A}(Z')}$. Then, clearly, for each given Z , we know the law of

$$T(Z') \mathbb{1}_{T(Z') \in \mathcal{A}(Z') \text{ and } T(Z') = \tilde{T}(Z)} = \tilde{T}(Z) \mathbb{1}_{\tilde{T}(Z) \in \mathcal{A}(Z')}.$$

But if I denotes the family of finite sets Z' with rational coordinates, then for any given Z ,

$$P\left(\tilde{T}(Z) \in \bigcup_{Z' \in I} \mathcal{A}(Z')\right) = 1$$

(because each triangle of $\tilde{T}(Z)$ contains some point with rational coordinates). Hence, it follows that if, for all Z' , one knows the law of $T(Z') \mathbb{1}_{T(Z') \in \mathcal{A}(Z')}$, then one can reconstruct the law of $\tilde{T}(Z)$ and therefore that of $T(Z)$.

Finally, it remains to prove that for each finite Z (we use Z instead of Z' now), the law of $T(Z) \mathbb{1}_{T(Z) \in \mathcal{A}(Z)}$ is fully determined by the knowledge of all two-dimensional marginal distributions of $(T(z), T(z'))$ for $z, z' \in \mathbb{D}$. We are going to do one further reduction step: Let us now suppose that for some finite Z we have $T(Z) \in \mathcal{A}(Z)$ (recall that this means that $T(Z)$ is good, and that each triangle of $T(Z)$ corresponds to exactly one point of Z). This defines naturally a connected tree structure G on Z , where each z_j has one, two or three neighbors in the graph (see Figure 4). We can therefore decompose the event $\{T \in \mathcal{A}(Z)\}$ according to the tree structure that $T(Z)$ induces on Z . Hence, it suffices to describe for each Z and each possible tree structure Γ on Z , the law of $T(Z) \mathbb{1}_{T \in \mathcal{A}(Z, \Gamma)}$, where

$$\mathcal{A}(Z, \Gamma) = \mathcal{A}(Z) \cap \{G = \Gamma\}.$$




Fig. 4. A configuration $T(Z)$, its completed configuration $\tilde{T}(Z)$ and a configuration in $\mathcal{A}(Z')$.

We are going to proceed by induction on the number of points in Z . Suppose that we know the law of all two-dimensional marginals $(T(z), T(z'))$, and that for each Z with no more than n points, and for each tree structure Γ on Z , we know the law of $T(Z) \mathbb{1}_{T \in \mathcal{A}(Z, \Gamma)}$. Let us show that, we then know it also for all $Z = \{z_1, \dots, z_{n+1}\}$ with $n + 1$ points and all tree structures Γ on Z . Let us choose such a Z with $n + 1$ points and a tree structure Γ on Z . Consider a leaf-point (i.e., a point in Z with just one Γ -neighbor)—by relabeling the points, we can assume that this leaf is z_1 and that its Γ -neighbor is z_2 . Our assumptions and previous results show that we know:

- The distribution of $(T(z_2), \dots, T(z_{n+1}))$, when restricted to the event that it defines the tree structure obtained by removing the leaf z_1 from Γ .
- The fact that conditionally on the event that $T(z_2)$ separates z_1 from the other $n - 1$ points, $T(z_1)$ is independent of $(T(z_3), \dots, T(z_n))$.
- The joint distribution of $(T(z_1), T(z_2))$ (and therefore also the conditional distribution of $T(z_1)$ given $T(z_2)$).

This readily shows that we know the distribution of $T(Z) \mathbb{1}_{T \in \mathcal{A}(Z, \Gamma)}$: Indeed, first sample $(T(z_2), \dots, T(z_{n+1}))$, look if it is compatible with $\mathcal{A}(Z, \Gamma)$, and then sample $T(z_1)$ according to the conditional distribution given $T(z_2)$.

Hence, we have proved our claim by induction n , which provides a characterization of the law of all $T(Z) \mathbb{1}_{T \in \mathcal{A}(Z, \Gamma)}$, and therefore by our previous arguments, of \mathbf{T} itself. \square

3. Uniqueness

3.1. Warm-up

In order to help those readers who are not well acquainted with the theory of regenerative sets, we briefly review some very classical facts on this topic (we refer to [3, 4] for details). Those readers who are familiar with these objects can safely skip this subsection.

Suppose that we are given a random non-empty closed subset F of \mathbb{R}_+ such that almost surely, $0 \in F$, F is not bounded, and the Lebesgue measure of F is 0. Suppose furthermore that it has the following regenerative property: For any $t \geq 0$, if we define $X_t = \min[t, \infty) \cap F$, then the law of $F_t := (F \cap [X_t, \infty)) - X_t$ is equal to that of F . We also assume that F_t is independent of $(X_t, F \cap [0, t])$. In the standard terminology, this means that F is a “light” (because its Lebesgue measure is 0) regenerative subset of \mathbb{R}_+ .

Then, for each given small positive x , we can find intervals of length greater than x in $\mathbb{R}_+ \setminus F$ from left to right. This defines (at least for small enough x) a sequence $\chi^x := (\chi_1^x, \chi_2^x, \dots)$ of lengths. The previous assumptions readily imply that this is a sequence of independent identically distributed random variables that have some common law ρ_x . Furthermore, when $x' < x$, the fact that $\chi^{x'}$ is almost surely a subsequence of χ^x implies that $\rho_x = \rho_{x'}(\cdot \mid [x, \infty))$. Hence, we can define a measure ρ on all of $(0, \infty)$ with the property that for all small enough x ,

$$\rho_x(A) = \frac{\rho(A \cap [x, \infty))}{\rho([x, \infty))}.$$

The measure ρ is unique up to a multiplicative constant and in a way describes the relative likelihood of appearance of intervals of a certain length in the complement of F . Note that it can happen that the total mass of ρ is infinite, which corresponds to the fact that there can be infinitely many (small) intervals in the complement of $F \cap [0, 1]$ say.

Now, it turns out that the measure ρ completely characterizes the law of the random set F . For instance, we can define simultaneously for each $x > 0$, a sample of χ^x in such a way that they are all compatible (i.e. χ^x is almost surely a subsequence of $\chi^{x'}$ when $x' < x$). Then the left endpoint of the interval corresponding to χ_n^x will be the sum of all intervals (of arbitrary length) that have appeared before it, which can be recovered from the knowledge of all $\chi^{x'}$ for $x' < x$.

One convenient way to express this is to use a Poisson point process: This is a random countable collection $\mathcal{P} := (t_i, x_i)_{i \in I}$ in $\mathbb{R}_+ \times \mathbb{R}_+$ where we introduced an artificial time parametrization at which the intervals appear. Intuitively, the existence of the point (t_i, x_i) in \mathcal{P} means that at time t_i , an interval of length x_i appears. If for a given x , we write down the sequence of lengths x_i greater than x in their order (with respect to time) of appearance, one gets a sample of χ^x . Then, the position of the left endpoint of the interval corresponding to i_0 can be recovered from \mathcal{P} as it is equal to $\sum_{i \in I: t_i < t_{i_0}} x_i$.

Note that if we were looking at the set $\tilde{F} := \{\exp(t) : t \in F\}$ instead of F , we would have had a set with similar properties to F : With obvious notation, the law of \tilde{F}_t/\tilde{X}_t is equal to that of \tilde{F}_0 , and in order to recover \tilde{F} from ρ , one has to replace the sum of all jumps x_i by the product of all $\exp(x_i)$. We shall use rather natural generalizations of these ideas in the next subsection.

3.2. Accordion and Poisson point process

Let \mathbf{T} be a complete triangulation of \mathbb{D} . For $x \neq y$ in $\overline{\mathbb{D}}$, we define the *accordion* between x and y in \mathbf{T} as the collection of all triangles $T \in \mathbf{T}$ intersecting the part of the hyperbolic line between x and y that goes through these two points, and denote it by $\text{Acc}_{\mathbf{T}}(x, y)$ (see Fig. 5).




Fig. 5. An accordion.

Suppose now that \mathbf{T} is Möbius-invariant complete and Markovian. Clearly, if we know the distribution of the accordion $\text{Acc}_{\mathbf{T}}(0, 1)$, this will characterize the law of $(T(0), T(u))$ for all $u \in (0, 1)$ and therefore (by Möbius invariance and Lemma 4) also the distribution

of \mathbf{T} . The goal of the present section will be to show that there is (at most) one possible law for this accordion. The basic idea will be to see that it necessarily corresponds to some particular subordinator (when one “discovers” the triangles of the accordion from 0 towards 1). In the next section, we shall check that the random triangulation defined using these random accordions is indeed Möbius-invariant complete and Markovian.

For notational convenience, we will now choose to work in the upper half-plane \mathbb{H} instead of the unit disk \mathbb{D} . For the remainder of this section, \mathbf{T} will denote a random Möbius-invariant complete Markovian triangulation of the upper half-plane (i.e. the image under ψ of such a triangulation of \mathbb{D}). Note that almost surely, ∞ and 0 are on the boundary of none of the triangles of \mathbf{T} (this follows from rotational invariance and the fact that the set of triangles is countable—one can also just look at the second characterization of the measure ν in the preliminaries). In other words, all triangles are bounded and bounded away from the origin. We are going to focus on the accordion between i and ∞ in \mathbf{T} , which will be denoted by $\text{Acc}_{\mathbf{T}}(i, \infty)$ (we will omit to specify that we are working in \mathbb{H}).




Fig. 6. Sample of (part of) our triangulation in \mathbb{H} .

Almost surely, for all positive rational y , one of the boundaries of $T(iy)$ does separate iy from ∞ . We denote it by (ℓ_y, r_y) where $\ell_y < 0 < r_y$. Clearly, $-\ell_y$ and r_y are non-decreasing functions of y . We can therefore define (ℓ_y, r_y) for all positive y simultaneously (including those that are in no triangle) by choosing the right-continuous version of $y \mapsto (\ell_y, r_y)$.

Let us first outline the idea of the proof: If we discover the triangle $T(iy)$, the conditional law of the part of the triangulation that is “above” this triangle can be described via σ , and it is (modulo taking its image under the affine map that maps ℓ_1, r_1 to $-1, 1$) always the same. It follows that the closure of the set $\{(r_y - \ell_y)/(r_1 - \ell_1) : y > 1\}$ is the exponential of a regenerative set, just as \tilde{F} at the end of the warm-up subsection. It can therefore be described thanks to a Poisson point process—the fact that the triangula-

tion is complete will imply that the Lebesgue measure of this set \tilde{F} is 0. The set \tilde{F} does not, however, contain enough information to reconstruct the accordion because when a triangle appears, one needs to know which one of the two processes ℓ or r is jumping. We will therefore describe the accordion via a slightly enriched Poisson point process that contains this additional information.

For all positive y , define

$$\tilde{y} := \sup\{v \in (0, y) : T(iv) \neq T(iy)\}.$$

For all y such that $T(iy) \neq \emptyset$, the third vertex of $T(iy)$ (apart from ℓ_y and r_y) is necessarily one of the two points $\ell_{\tilde{y}-}$ or $r_{\tilde{y}-}$. Note also that the jumps of the process (ℓ, r) exactly correspond to the triangles of $\text{Acc}_{\mathbf{T}}(i, \infty)$ (i.e., the set \mathcal{J} of “jumping heights” is equal to $\{\tilde{y} : y > 0\}$). Note also that y can never be a simultaneous jumping height for ℓ and r (because almost surely, no $T(iy)$ is a quadrilateral).

For each positive y , we denote by φ_y the affine map that maps (ℓ_{y-}, r_{y-}) onto $(-1, 1)$. We can then describe the jumps of $T(iy)$ by defining for each $y \in \mathcal{J}$, $X(y)$ to be the image of the third apex of $T(iy)$ (apart from ℓ_{y-} and r_{y-}) under φ_y . In other words,

$$X(y) = \varepsilon(y) \cdot \left(2 \frac{r_y - \ell_y}{r_{y-} - \ell_{y-}} - 1 \right) \tag{1}$$

where

$$\varepsilon(y) = \mathbf{1}_{\{r_y \neq r_{y-}\}} - \mathbf{1}_{\{\ell_y \neq \ell_{y-}\}}.$$

When $y \notin \mathcal{J}$, we can declare $X(y)$ to be equal to an abstract cemetery point δ .

Note that for all $y_1 > 1$, and all $x > 1$, the number of jumps \tilde{y} in $[1, y_1]$ such that $|X(\tilde{y})| > x$ is finite. Hence, it follows that the collection $(X(\tilde{y}) \mathbf{1}_{|X(\tilde{y})| > x}, \tilde{y} \geq 1)$ almost surely defines an ordered discrete sequence $\zeta^x = (\zeta_1^x, \zeta_2^x, \dots)$ in $\mathbb{R} \setminus [-x, x]$. Note that when $x' > x$, the sequence $\zeta^{x'}$ is a deterministic subsequence of ζ^x . We define \mathcal{C} to be this nested family of sequences $(\zeta^x, x > 1)$ (we cannot view it as just one sequence, because infinitely many “small” jumps occur before any given jump).

An equivalent way to encode \mathcal{C} is to define it as the process of jumps $(X(\tilde{y}), \tilde{y} \geq 1)$ but defined modulo increasing time reparametrization, i.e., only the order of arrivals of the jumps matters.

In the following, for $x \geq 1$ we let $I_x = \mathbb{R} \setminus [-x, x]$. In particular $I_1 = \mathbb{R} \setminus [-1, 1]$.

Lemma 5. *The ordered (but unparametrized) set of jumps \mathcal{C} has the same distribution as the ordered family of jumps (modulo increasing time reparametrization) of a Poisson point process $\mathcal{P} = \{(t_i, x_i)_{i \in I}\}$ on $\mathbb{R}_+ \times \mathbb{R}$ with intensity $dt \otimes \rho$, where ρ is some sigma-finite measure on $\mathbb{R} \setminus [-1, 1]$.*

Proof. Let $t > 1$. By the spatial Markov property applied to the triangle $T(it)$ (i.e., the push-forward of Lemma 3 by ψ), one deduces that the ordered (but unparametrized) collection of jumps $(X(\tilde{y}), \tilde{y} \geq t)$ is independent of $(X(\tilde{y}), 1 \leq \tilde{y} \leq t)$ and has the same distribution as the ordered family of jumps $(X(\tilde{y}), \tilde{y} \geq 1)$.

Fix $x_0 > 1$ such that there almost surely exists a jump in I_{x_0} . We deduce from the above remark that for every $1 < x < x_0$, the discrete random sequence $(X(\tilde{y}) 1_{|X(\tilde{y})| \in I_x}, \tilde{y} > 1)$ has the same distribution as i.i.d. samples from a certain probability measure ρ_x on I_x . Furthermore, the distributions ρ_x satisfy the compatibility condition

$$\rho_{x'}(\cdot | I_x) = \rho_x$$

for all $1 < x' < x$. Consequently, one can uniquely define a sigma-finite measure ρ on $I_1 = \mathbb{R} \setminus [-1, 1]$ such that $\rho(\cdot \cap I_x) / \rho(I_x) = \rho_x$ and $\rho(I_{x_0}) = 1$. It is then easy to see that the jumps of \mathcal{C} have the same distribution as the unparametrized jumps of a Poisson point process $\mathcal{P} = \{(t_i, x_i)_{i \in I}\}$ on $\mathbb{R}_+ \times \mathbb{R}$ with intensity $dt \otimes \rho$ (see the warm-up section). The details are left to the reader.

The sigma-finite measure ρ has an arbitrary multiplicative normalization (but note that the multiplicative constant does not change the law of the ordered family of jumps, it just changes the time-parametrization). \square

We have now seen that if a Möbius-invariant complete Markovian triangulation \mathbf{T} exists, then one can associate with it a measure ρ that describes the law of the jumps of \mathcal{C} , and we have also seen that the distribution of $T(i)$ is necessarily the image of P_0 under ψ . Furthermore, the Markovian property (Lemma 3) shows that $T(0)$ and the jumps of \mathcal{C} are independent. The following lemma proves that conversely, one can recover the law of $\text{Acc}_{\mathbf{T}}(i, \infty)$ from ρ and $T(i)$:

Lemma 6. *The distributions of $T(i)$ and \mathcal{C} fully characterize the law of $\text{Acc}_{\mathbf{T}}(i, \infty)$.*

Proof. It is clear that $\text{Acc}_{\mathbf{T}}(i, \infty)$ can be recovered from the two processes $(\ell_y)_{y \geq 1}$ and $(r_y)_{y \geq 1}$ and the initial triangle $T(i)$. More precisely, instead of the full processes (ℓ) and (r) , it suffices to know (ℓ, r) up to time reparametrization to reconstruct the accor-dion. Indeed, only the range of (ℓ, r) matters in order to define $\text{Acc}_{\mathbf{T}}(i, \infty)$.

We first claim that the ranges of the processes (ℓ) and (r) are both of zero one-dimensional Lebesgue measure. Recall that the triangulation \mathbf{T} is almost surely complete, and Möbius-invariant, so that any given point in \mathbb{H} is almost surely in the interior of some triangle of \mathbf{T} . Hence, the (one-dimensional) Lebesgue measure of the intersection \mathcal{I} of the imaginary line with the closure of the union of all arches $(\ell_y r_y)$ is almost surely equal to 0.

Indeed, assume that the Lebesgue measure of the intersection of the range of (ℓ) with some interval $[-l_1, -l_2]$ is positive. Clearly, one can associate to each point ℓ_y of this range a point iy on the imaginary half-line in such a way that for any $y < y'$ we have $|y' - y| > |\ell_{y'} - \ell_y| / K(l_1)$ for some constant $K(l_1)$. Hence, it follows readily that the one-dimensional Lebesgue measure of \mathcal{I} is positive. Since this is prohibited, we conclude that the range of (ℓ) (and of (r) , by the same argument) is almost surely of zero (one-dimensional) Lebesgue measure.

As (ℓ) and (r) are monotone functions, the fact that their ranges are of zero Lebesgue measure implies that the range of (ℓ, r) is characterized by its jumps (which themselves are described by \mathcal{C}) and by its initial value (which is given by $T(i)$). Hence, we can recover, up to time reparametrization, the process (ℓ, r) from \mathcal{C} and $T(i)$. This is sufficient to reconstruct $\text{Acc}_{\mathbf{T}}(i, \infty)$. \square

3.3. Identification of the jump measure

It now remains to show that (up to a multiplicative constant) there is in fact at most one possibility for the measure ρ defined in Lemma 5. This will follow from the Möbius invariance of the measure σ as heuristically described in the introduction.

Let us suppose for the remainder of this section that \mathbf{T} is a Möbius-invariant complete Markovian triangulation, and that ρ and \mathcal{P} are defined as in Lemma 5 (and \mathcal{P} is coupled with \mathcal{C} in such a way that their ordered families of jumps are identical). If φ is a Möbius transformation of \mathbb{H} , the action of φ can be extended to the boundary $\partial\mathbb{H} = \mathbb{R}$. We will implicitly use this extension in what follows. Note that it is sufficient for an arch or a triangle to track down its apexes $\in \partial\mathbb{H}$ to know it entirely.

Lemma 7. *The image measure of ρ under any Möbius transformation of the upper half-plane into itself that fixes -1 and 1 is proportional to ρ .*

Proof. Fix $x_0 > 1$ in such a way that $\rho(\{x_0, -x_0\}) = 0$ and $\rho(I_{x_0}) > 0$. Define $\ell_1 < 0 < r_1$ as before so that (ℓ_1, r_1) is the top boundary of $T(i)$. We denote \tilde{x} the first jump of the point process \mathcal{P} such that $|\tilde{x}| > x_0$ and write $\tilde{\ell} < 0 < \tilde{r}$ for the feet of the bottom hyperbolic line of the triangle corresponding to the jump \tilde{x} . Let \tilde{a} denote the third apex of this triangle. See Figure 7.




Fig. 7. Setup of the proof.

For each small ε and ε' , we consider the events

$$A(\varepsilon, \varepsilon') := \{-1 - \varepsilon < \tilde{\ell} < \ell_1 < -1 \text{ and } 1 < r_1 < \tilde{r} < 1 + \varepsilon'\}.$$

By standard properties of Poisson point processes the event $A(\varepsilon, \varepsilon')$ is independent of \tilde{x} which is distributed according to the measure ρ_{x_0} . Thus for any Borel positive $f : \mathbb{R} \setminus [-1, 1] \rightarrow \mathbb{R}_+$ with compact support we have

$$E[f(\tilde{x}) \mid A(\varepsilon, \varepsilon')] = \rho(I_{x_0})^{-1} \int_{I_{x_0}} \rho(da) f(a). \tag{2}$$

We will now let $x_0 \rightarrow 1$. To avoid subsequent normalizations, we consider another positive measurable function $g : \mathbb{R} \setminus [-1, 1] \rightarrow \mathbb{R}_+$ with compact support: Using the last display and letting $x_0 \rightarrow 1$ we have

$$\frac{E[f(\tilde{x}) \mid A(\varepsilon, \varepsilon')]}{E[g(\tilde{x}) \mid A(\varepsilon, \varepsilon')]} \xrightarrow{\varepsilon, \varepsilon' \rightarrow 0} \frac{\int_{I_1} \rho(dx) f(x)}{\int_{I_1} \rho(dx) g(x)}. \tag{3}$$

On the other hand, \tilde{x} can be related to the geometric quantity \tilde{a} as follows. When ε and ε' are both very small (and $A(\varepsilon, \varepsilon')$ holds) then the jump \tilde{x} is necessarily very close to the first foot \tilde{a} of the accordion with absolute value larger than x_0 . Thanks to the above remark, the \tilde{x} can be replaced by the geometric \tilde{a} on the left-hand side of (3).

Let us now suppose that φ is a Möbius transformation that maps \mathbb{H} onto itself with $\varphi(-1) = -1$ and $\varphi(1) = 1$. Note in particular that since the semi-circle $(-1, 1)$ is preserved by φ , for every $\varepsilon, \varepsilon' > 0$ there exist $\delta, \delta' > 0$ such that if $A(\delta, \delta')$ is satisfied for \mathbf{T} then $A(\varepsilon, \varepsilon')$ holds for $\varphi(\mathbf{T})$ and furthermore $T(i) = T(\varphi(i))$. Since $\varphi(\mathbf{T})$ and \mathbf{T} are identically distributed it follows readily using the same arguments as before that

$$\frac{E[f(\varphi(\tilde{a})) \mid A(\varepsilon, \varepsilon')]}{E[g(\varphi(\tilde{a})) \mid A(\varepsilon, \varepsilon')]} \xrightarrow{\varepsilon, \varepsilon' \rightarrow 0} \frac{\int_{I_1} \rho(dx) f(x)}{\int_{I_1} \rho(dx) g(x)}. \tag{4}$$

Thus comparing (4) with (3) (with \tilde{x} replaced by \tilde{a}) we deduce that

$$\frac{\int_{I_1} \rho(dx) f(\varphi(x))}{\int_{I_1} \rho(dx) g(\varphi(x))} = \frac{\int_{I_1} \rho(dx) f(x)}{\int_{I_1} \rho(dx) g(x)}.$$

Hence the push-forward of ρ under the map φ is indeed a multiple of ρ . □

A natural candidate for the measure ρ is the measure ζ on $\mathbb{R} \setminus [-1, 1]$ defined by

$$\zeta(dx) = \frac{2 dx}{|x|^2 - 1} 1_{\{|x|>1\}}$$

as it is the only measure (up to a multiplicative constant) that is invariant under all Möbius transformations of \mathbb{H} that fix the boundary points -1 and 1 (note for instance that it is the image of the measure dx/x on \mathbb{R}_+ under the map $(1+x)/(1-x)$). Indeed:

Lemma 8. *The measure ρ defined in Lemma 5 is necessarily equal to a constant times the measure ζ .*

Proof. Let us consider $\varphi(x) = (1+x)/(1-x)$ as above. Clearly, if $\tilde{\rho}$ denotes the push-forward of ρ under φ^{-1} , this measure $\tilde{\rho}$ on \mathbb{R}_+ will have the property that the image of $\tilde{\rho}$ under any map $z \mapsto \lambda z$ for positive λ (these are the Möbius transformations of \mathbb{H} onto itself that fix 0 and ∞) is a multiple (which may depend on λ) of $\tilde{\rho}$. It follows that for some real α and some positive constant c ,

$$\tilde{\rho}(dx) = cx^{-\alpha} 1_{\{x>0\}} dx.$$

We want to show that α is necessarily equal to 1. Let us assume that $\alpha < 1$. Then $\tilde{\rho}[\varepsilon, \infty) = \infty$ while $\tilde{\rho}[0, \varepsilon) < \infty$ for any $\varepsilon > 0$. In terms of ρ , this implies in particular that $\rho[1, \infty) < \infty$. But the proof of Lemma 6 then tells us that the set of $(r_y)_{y \geq 1}$ has no accumulation points, i.e. all r_y 's are isolated. In particular, this implies that if E denotes the set of all corners of triangles in \mathbf{T} that separate 0 from ∞ in \mathbb{H} , then $E_+ := E \cap \mathbb{R}_+$ is almost surely discrete in the sense that for all $0 < a < b$, $E_+ \cap [a, b]$ is finite (here $[a, b]$ denotes the horizontal segment between a and b).

On the other hand, for any $\varepsilon > 0$, $\rho[-1 - \varepsilon, -1] = \infty$. This readily shows there are infinitely many jumps for (ℓ_y) while (r_y) only jumps finitely many times. In particular, we see that almost surely, there exist $b < a < 0$ such that there are infinitely many points in the intersection of $E_- := E \cap \mathbb{R}_-$ with the horizontal segment $[b, a]$.

Finally, because of invariance of the law of \mathbf{T} under the transformation $z \mapsto -1/z$, we note that E_+ has the same law as $\{-1/z, z \in E_-\}$, which contradicts the previous facts that we just proved for E_+ and E_- .

We therefore conclude that $\alpha \geq 1$. In exactly the same way, we can exclude the possibility that $\alpha > 1$ (because then $\rho[1, \infty) = \infty$ while $\rho(-\infty, -1] < \infty$). Hence, we see that ρ is a multiple of the image under φ of $x^{-1} dx 1_{\{x>0\}}$, i.e., a multiple of ζ . \square

The previous lemmas describe the joint law of $(T(i), T(iy))$ for any given $y > 1$. But, for any z and z' in \mathbb{D} , there exists some $y \geq 1$ and a Möbius transformation from \mathbb{D} onto \mathbb{H} that maps z onto i and z' onto iy ; by Möbius invariance, we can therefore describe the joint law of $(T(z), T(z'))$, and by Lemma 4, we have completed the proof of the uniqueness part of Theorem 2:

Proposition 9. *There exists at most one (law of a) complete Möbius-invariant Markovian triangulation.*

4. Existence

The goal of this section is to define the candidate for the random triangulation, and to check that it is complete, Markovian and Möbius-invariant.

4.1. The half-plane accordion

In order to define a random accordion in \mathbb{H}^+ (i.e., what will turn out to be our distribution σ), we start with a Poisson point process $\mathcal{P} = \{(t_i, x_i)_{i \in I}\}$ on $\mathbb{R}_+ \times (\mathbb{R} \setminus [-1, 1])$ with intensity $dt \otimes \zeta$.

We then construct two pure jump processes $(L_t)_{t \geq 0}$ (for left) and $(R_t)_{t \geq 0}$ (for right) that jump only at the jumping times of \mathcal{P} whose jumps (defined as in (1)) are the x_i 's. Set $L_0 = -1$ and $R_0 = 1$. The idea is that L is decreasing, R is increasing, and when a jump (t, x) occurs, then $R - L$ is multiplied by $(|x| + 1)/2$, and L jumps only if $x < -1$, while R jumps only if $x > 1$.

More precisely, if we set $\mathcal{P}_t := \{(t_i, x_i) \in \mathcal{P} : t_i \leq t\}$, then we can first define

$$\Delta_t := (R_0 - L_0) \prod_{(t_i, x_i) \in \mathcal{P}_t} \frac{|x_i| + 1}{2},$$

which is the exponential of the pure jump process with intensity given by the image of ζ under the mapping $x \mapsto \log \frac{|x|+1}{2}$. It is easily checked that this subordinator is well-defined (does not blow up), using the explicit expression for its jump measure. Then, we simply set

$$R_t := R_0 + \sum_{(t_i, x_i) \in \mathcal{P}_t} (\Delta_{t_i} - \Delta_{t_i-}) 1_{x_i > 1}, \tag{5}$$

$$L_t := L_0 - \sum_{(t_i, x_i) \in \mathcal{P}_t} (\Delta_{t_i} - \Delta_{t_i-}) 1_{x_i < -1}. \tag{6}$$

As $\Delta_t = R_t - L_t$ is almost surely finite for all t , the two processes $(L_t)_{t \geq 0}$ and $(R_t)_{t \geq 0}$ are well-defined as well (note that L is non-increasing and R is non-decreasing). Note also that $R_t \rightarrow \infty$ and $L_t \rightarrow -\infty$ almost surely as $t \rightarrow \infty$.

Equivalently, for each $i \in I$, we can write

$$(L_{t_i}, R_{t_i}) = \varphi_{t_i-}^{-1}((-1, x_i)) 1_{x_i > 1} + \varphi_{t_i-}^{-1}(x_i, 1) 1_{x_i < -1}$$

where φ_t^{-1} denotes the affine map that maps $(-1, 1)$ onto (L_{t-}, R_{t-}) .

We are now ready to define our accordion in \mathbb{H}^+ . For each $(t_i, x_i) \in \mathcal{P}$, define the hyperbolic triangle in \mathbb{H} with three corners given by $(L_{t_i}, R_{t_i-}, R_{t_i})$ if $x_i > 1$, and by $(L_{t_i}, L_{t_i-}, R_{t_i})$ if $x_i < -1$. The definition clearly ensures that each of these triangles separates the semi-circle $\{z \in \mathbb{H} : |z| = 1\}$ from ∞ in \mathbb{H}^+ and that these triangles are disjoint.




Fig. 8. Sample of a piece of $\mathcal{A}(\mathcal{P})$.

In fact, to indicate that this accordion is from the semi-circle $(-1, 1)$ to ∞ in \mathbb{H} , we will denote it by $\mathcal{A}(\mathcal{P})_{[(-1, 1) \rightarrow \infty; \mathbb{H}]}$ (and omit the \mathbb{H} when it is clear that we are working in \mathbb{H} , and simply write $\mathcal{A}(\mathcal{P})_{(-1, 1) \rightarrow \infty}$).

We can immediately extend our construction to the case where the initial position $(L_0, R_0) = (l_0, r_0)$ (for $r_0 > l_0$) is different from $(-1, 1)$, and this defines the accordion

$\mathcal{A}(\mathcal{P})_{(l_0, r_0) \rightarrow \infty}$. It is easy to check that that this new accordion has the same distribution as the image of $\mathcal{A}(\mathcal{P})_{(-1, 1) \rightarrow \infty}$ under the linear map that maps -1 to l_0 and 1 to r_0 .

In other words, we have in fact defined (L_t, R_t) as a Markov process on $\{(l, r) : l < 0 < r\}$ with translation-invariant and scale-invariant transition kernel (the process started from $(-1 + x, 1 + x)$ has the same law as $(x + L_t, x + R_t)_{t \geq 0}$ when $L_0 = -1$ and $R_0 = 1$, and on the other hand, the process started from $(-r, r)$ has the same law as $(rL_t, rR_t)_{t \geq 0}$).

4.2. Towards completeness

Let us now prove the following statement:

Lemma 10. *Almost surely, the ranges of $(L_t)_{t \geq 0}$ and $(R_t)_{t \geq 0}$ restricted to any compact interval of \mathbb{R} are both of box-counting dimension 0.*

Proof. Consider an auxiliary subordinator $(Q_t, t \geq 0)$ defined by

$$Q_t = \sum_{(t_i, x_i) \in \mathcal{P}_t} \frac{x_i - 1}{2} 1_{x_i > 1}.$$

Let $T > 0$. It is clear from its construction that the process $(R_t, t \geq 0)$ jumps exactly when the process $(Q_t, t \geq 0)$ jumps, and that up to time T , the size of a jump of R is less than the corresponding jump of Q multiplied by Δ_T : indeed, if $t < T$ is a jump time for R we have

$$R_t - R_{t-} = \Delta_t - \Delta_{t-} = \Delta_{t-} \frac{x_i - 1}{2} = \Delta_{t-} (Q_t - Q_{t-}).$$

Thus, the box-counting dimension of the range $R[0, T]$ is almost surely not larger than that of $Q[0, T]$ (because the former set is the image of the latter under a Lipschitz map). But the box-counting dimension of $Q[0, T]$ is easily seen to be almost surely equal to zero (see [4, Chapter 5.1.1] or [3, Chapter III.5], and use the behavior of $\zeta(dx)$ near $x = 1$). As the process L has the same law as $-R$, the lemma follows. \square

Similarly to the case of Lemma 6, we will translate the previous result on the range of L and R into a property of the set of points of the accordion that are on the imaginary axis. More precisely, let us define the set \mathcal{I}_∞ of points of the type iy for $y \geq 1$ that are not inside a triangle of $\mathcal{A}(\mathcal{P})_{(-1, 1) \rightarrow \infty}$. Then we have:

Corollary 11. *The (one-dimensional) Lebesgue measure of \mathcal{I}_∞ is almost surely equal to zero.*

Proof. For any $k > 1$, let us define the set J_k of points iy for $y \geq 1$ that are in the closure of the union of the semi-circles $((L_t, R_t), t < \sigma_k)$, where σ_k is the first time at which $\max(R_t, -L_t) \geq k$. Clearly, it is sufficient to prove that for any given k , this set J_k has almost surely zero Lebesgue measure. For each positive ε , define N_ε to be the minimal number of intervals of length ε that are needed to cover $R[0, \sigma_k] \cup L[0, \sigma_k]$. Lemma 10 in particular implies that almost surely, $N_\varepsilon / \varepsilon^{-1/3}$ vanishes as ε tends to 0.




Fig. 9. Sketch of the covering.

Suppose that iy (for $y > 1$) is in no triangle of $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow \infty}$. This means that one can find one of the intervals of length ε covering the range of L (call it I_l), and one of the intervals of length ε covering the range of R (that we call I_r) such that iy is in one of the semi-circles joining a point in I_l to a point in I_r . See Fig. 9. But, for a given k , and any two such intervals, the length of the set of points on the imaginary axis that can be reached in this way is bounded by a constant $C = C(k)$ times ε . Thus we have

$$\text{Leb}_1(J_k) \leq C\varepsilon \times N_\varepsilon^2.$$

The right-hand side goes almost surely to 0 as ε vanishes, which concludes the proof of the corollary. \square

4.3. Target-independence

Let us now recall a simple classical lemma (see for instance [3, Section O.5]—it can be viewed as a direct consequence of the “compensation formula”) that roughly states that if we start with a Poisson point process, and modify it in a way that preserves both the independence and the intensity measure, then the law of the modified point process is still the same:

Lemma 12 (Modification of Poisson point processes). *Let $\mathcal{P} = \{(t_i, x_i), i \in I\}$ be a Poisson point process on $\mathbb{R}_+ \times \mathbb{R}$ of intensity $dt \otimes \rho$ (where ρ denotes some measure on \mathbb{R}). Let $(H_t)_{t \geq 0}$ be a predictable process taking values in the space of non-negative measurable functions $\mathbb{R} \rightarrow \mathbb{R}$ such that almost surely, for every $t \geq 0$ the push-forward of the measure ρ by H_t is ρ . Then $\mathcal{P}' := \{(t_i, H_{t_i}(x_i)), i \in I\}$ has the same law as \mathcal{P} .*

We will use this lemma in order to derive a target-independence property of our accordion $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow \infty}$.

Fix $a \in \mathbb{R} \setminus [-1, 1]$ and let us define $\psi_{a,0}$ to be the Möbius map from \mathbb{H} onto itself that maps $(-1, 1, \infty)$ onto $(-1, 1, a)$. We define the accordion $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow a}$ in \mathbb{H} to be the image of $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow \infty}$ under $\psi_{a,0}$. We finally denote by $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow \infty \wedge a}$ the sub-accordion of $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow \infty}$ whose triangles intersect the line $(0a)$, and similarly we denote by $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow a \wedge \infty}$ the sub-accordion of $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow a}$ whose triangles intersect the line (0∞) .

Proposition 13 (Target independence). *For any $a \in \mathbb{R} \setminus [-1, 1]$, the two accordions $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow a \wedge \infty}$ and $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow \infty \wedge a}$ have the same law.*

Proof. The idea is to decompose the global action of the composition of an accordion with a Möbius transformation into an iteration of infinitesimal transformations of the jumps by (predictable) functions. Consider (L) and (R) , the two functions associated with a standard accordion $[(-1, 1) \rightarrow \infty; \mathbb{H}]$, and introduce the disconnection time of a and ∞ :

$$\theta_a := \inf\{t \geq 0 : L_t < a < R_t\}.$$

For each jumping time $t_i \leq \theta_a$, the jump x_i in the accordion $[(-1, 1) \rightarrow \infty]$ corresponds to the image of L_{t_i} or R_{t_i} under the affine map that sends $(L_{t_i-}, R_{t_i-}, \infty)$ onto $(-1, 1, \infty)$. Let us now see what is the corresponding jump $x_i^{(a)}$ in the image of $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow \infty}$ under $\psi_{a,0}$ that we consider as an accordion growing towards ∞ , at least as long as $t_i \leq \theta_a$. The jump $x_i^{(a)}$ corresponding to x_i in $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow a}$ (which is $\psi_{a,0}(\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow \infty})$ by definition) is the image of $\psi_{a,0}(L_{t_i})$ or $\psi_{a,0}(R_{t_i})$ under the hyperbolic isometry that sends $(\psi_{a,0}(L_{t_i-}), \psi_{a,0}(R_{t_i-}), \infty)$ onto $(-1, 1, \infty)$. We deduce that $x_i^{(a)}$ is the image of x_i by the hyperbolic isometry

$$\psi_{a,t_i} := \varphi_{(\psi_{a,0}(L_{t_i-}), \psi_{a,0}(R_{t_i-}), \infty) \rightarrow (-1, 1, \infty)} \circ \psi_{a,0} \circ \varphi_{(-1, 1, \infty) \rightarrow (L_{t_i-}, R_{t_i-}, \infty)}.$$

Note that the measure ζ is invariant under ψ_{a,t_i} and that this is a predictable function (with respect to the natural filtration defined by the Poisson point process). When $t > \theta_a$, we simply define $\psi_{a,t}$ to be the identity. Hence, we deduce from Lemma 12 that the two ordered but unparametrized families $\{x_i, i \in I\}$ and $\{\psi_{a,t_i}(x_i), i \in I\}$ have the same law. This, together with the fact that the jumps characterize the accordion (Lemma 6), tells us precisely that up to the first time at which one disconnects ∞ from a , the two accordions $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow a}$ and $\mathcal{A}(\mathcal{P})_{(-1,1) \rightarrow \infty}$ are identically distributed. Note that the final “jump” (i.e. the triangle that disconnects a from ∞) is also included in this description. \square

It is therefore possible to couple the two accordions aiming at $a \in \mathbb{R} \setminus [-1, 1]$ and ∞ , in such a way that they coincide up to the triangle disconnecting a and ∞ . This compatibility shows that it is in fact possible to couple accordions (all based on $(-1, 1)$) aiming at all points (with rational coordinates, say) in $\mathbb{R} \setminus [-1, 1]$ in such a way that any two of them coincide up to the first triangle that disconnects their two targets. We define by $\mathcal{A}_{(-1,1),\mathbb{H}}$ the union of all the triangles in this “accordion tree”. Then:

- The distribution of $\mathcal{A}_{(-1,1),\mathbb{H}}$ is invariant under the one-dimensional family of conformal maps from \mathbb{H} onto itself that fix $(-1, 1)$. This follows just from the definition of the accordion targeting other points than ∞ via Möbius invariance.
- This triangulation $\mathcal{A}_{(-1,1),\mathbb{H}}$ is almost surely complete, i.e. almost surely, the two-dimensional Lebesgue measure of the complement in \mathbb{H}^+ of the union of all the triangles of $\mathcal{A}_{(-1,1),\mathbb{H}}$ is zero. This is just due to the fact that for any $a \in \mathbb{R} \setminus \{0\}$, the intersection of this set with the hyperbolic line joining 0 to a has almost surely zero (one-dimensional) Lebesgue measure (which again follows from the result for $a = \infty$, i.e. from Corollary 11, via Möbius invariance).

This target independence is reminiscent of the “locality property” of SLE₆ [9].

4.4. Reversibility of the accordion

On top of being invariant under Möbius transformations of \mathbb{H} that leave \mathbb{H}^+ invariant, the measure ζ has another property that will yield “reversibility” of the accordion. Recall from the end of Section 4.1 that we can view $M_t = M_t^+ = (R_t, L_t)_{t \geq 0}$ as a pure jump Markov process on the space $\{(r, l) \in \mathbb{R}^2 : l < 0 < r\}$.

We will also use the Markov process $M_t^- = (L_t^-, R_t^-)$ defined on the same state space, but aiming at 0 instead of ∞ . It is defined exactly as M_t except that the measure ζ is replaced by the measure $\zeta^-(dx) = 2 dx / (1 - x^2)$ with support in $[-1, 1]$. Note that ζ^- is the image of ζ under the map $x \mapsto -1/x$, so that it follows (using the same arguments as in the proof of Proposition 13) that if $M_0 = M_0^- = (-1, 1)$, then the two processes $(L_t^-, R_t^-)_{t \geq 0}$ and $(-1/R_t, -1/L_t)_{t \geq 0}$ have the same law.

Let us define, for any $u < v$ in \mathbb{R} , the two measures $\zeta_{[u,v]}$ and $\zeta_{[v,u]}$ that are supported respectively on $[u, v]$ and $\mathbb{R} \setminus [u, v]$ with respective densities

$$\zeta_{[u,v]}(dw) = \frac{(v - u) dw}{(v - w)(w - u)} \quad \text{and} \quad \zeta_{[v,u]}(dw) = \frac{(v - u) dw}{(w - v)(w - u)}.$$

Note that these two measures are invariant under any Möbius transformation of \mathbb{H} that fixes v and w , and that they are the push-forward of the measure dx/x on \mathbb{R}_+ by any Möbius transformation from \mathbb{H} onto itself that maps 0 and ∞ to u and v (in that order for $\zeta_{[u,v]}$ and in the reverse order for $\zeta_{[v,u]}$; in fact, the measure dx/x on \mathbb{R}_+ can be interpreted as $\zeta_{[0,\infty]}$). Hence, all these measures are images of each other under some hyperbolic isometry. Note also that $\zeta_{[1,-1]}$ is exactly our measure ζ and that $\zeta_{[-1,1]} = \zeta^-$.

The definition of our Markov process M_t shows that $\zeta_{[v,u]}$ describes its jump intensity measure (i.e. the location of the new point after the jump when $(L, R) = (u, v)$) and similarly $\zeta_{[u,v]}$ describes the jump intensity measure for M_t^- .

Here is a simple observation: Suppose that the pair (u, v) is defined under the infinite measure

$$\pi(du dv) = \frac{du dv}{(v - u)^2} 1_{\{u < 0 < v\}}$$

on the set $\mathbb{R}_- \times \mathbb{R}_+$ (it is important for what follows that we restrict ourselves to this set!). We can then define w under the measure $\zeta_{[v,u]}$ so that the triple (u, v, w) is defined on the

set $\{u < 0 < v < w\} \cup \{w < u < 0 < v\}$ by the measure with density

$$\nu(du \, dv \, dw) = \frac{du \, dv \, dw}{(w - v)(v - u)(w - u)}.$$

We recognize here (a multiple of) the Haar measure on unmarked hyperbolic triangles in \mathbb{H} (see Section 2.1), restricted to those triangles that separate 0 from ∞ . Note also that when one sees this triangle without knowing which is u or which point is v , one can recover it immediately: These triangles always have at least one point on \mathbb{R}_+ and one point on \mathbb{R}_- . If there are two points in \mathbb{R}_+ , then $u < 0 < v < w$, and if there are two points in \mathbb{R}_- , then $w < u < 0 < v$.

We can use the same procedure in the other direction. Let us first define (α, β) under the same measure π and then γ under $\zeta_{[\alpha, \beta]}$ (mind that this time, $\gamma \in [\alpha, \beta]$) so that one obtains the triple (α, β, γ) defined on the set $\{\alpha < \gamma < 0 < \beta\} \cup \{\alpha < 0 < \gamma < \beta\}$ under the measure with intensity $d\alpha \, d\beta \, d\gamma / ((\beta - \alpha)(\gamma - \alpha)(\beta - \gamma))$. In this way, we get exactly the same measure as before, and we can also recover from the unmarked triangle which apexes are α, β and γ .

We have thus proved the following by looking at the properties of the jump measures of the two processes M and M^- (see for instance [7] for background on duality for Markov processes):

Proposition 14. *The measure π is an invariant measure for the Markov transition kernel of M and the dual kernel is that of M^- .*

In plain words, suppose that for a given positive t :

- We define (u, v) according to π , then sample a Markov process M starting from $M_0 = (u, v)$; this defines an infinite measure on quadrilaterals (u, v, L_t, R_t) .
- We define (u', v') according to π , then sample a Markov process M^- starting from $M_0^- = (u', v')$; this defines an infinite measure on quadrilaterals (L_t^-, R_t^-, u', v') .

Then these two measures on quadrilaterals are the same.

Suppose now that $s = iy$ is some fixed point on the vertical line with $y > 1$. Let us first sample a triangle T_i according to $P_i^{\mathbb{H}}$ (this is the probability measure obtained by restricting the Haar measure on triangles in \mathbb{H} to those that contain i). Denote its three apexes by a, l, r in such a way that $l < a < r$. The arc of T_i that separates i from infinity is therefore (l, r) . Note that one way to sample T_i is to define l, r under some universal constant c_0 times π , and then a under $\zeta_{[l, r]}$ and to finally restrict the resulting measure on (l, a, r) to those triangles that contain i .

We now define a triangle $T(s)$ that contains s . If $s \in T_i$, we take $T(s) = T_i$. Otherwise, we sample an accordion $\mathcal{A}_{(l, r) \rightarrow \infty}$. Then, almost surely, s is in one of the triangles of the accordion, which we call $T(s)$. Our goal in this subsection is to show the following reversibility property of the joint law of $(T_i, T(s))$:

Lemma 15. *If $\varphi : z \mapsto -y/z$ denotes the Möbius transformation in \mathbb{H} that interchanges i and s , then the law of $(T_i, T(s))$ is equal to that of $(\varphi(T(s)), \varphi(T_i))$.*

An equivalent possibly clearer way to phrase this reversibility is as follows: Define a random triangle T'_s that contains s according to $P_s^{\mathbb{H}}$, and then define the triangle $T'(i)$ that contains i , obtained by letting an accordion grow from T'_s towards 0. Since $(T'(i), T'_s)$ is distributed as $(\varphi(T(s)), \varphi(T_i))$, the lemma says that $(T_i, T(s))$ and $(T'(i), T'_s)$ are identically distributed.

Proof of Lemma 15. Let us first notice that when restricted to the event that the two triangles are equal (i.e., $s \in T_i$ for the first pair, and $i \in T'_s$ for the second), the laws of $(T_i, T(s))$ and $(T'(i), T'_s)$ are equal (they are both described via the Haar measure on triangles, restricted to the triangles that contain both i and s). We can therefore focus on the event where the two triangles are different.

Suppose now that T_i (and therefore (l, r)) has been sampled and does not contain s . When one grows the accordion from the arc (l, r) towards ∞ , we use the Poisson point process described in previous subsections. Note that we are interested in the law of the first “time” (in the Poisson point process parametrization) at which one discovers a triangle (i.e., a jump in our pure jump process) that swallows the point s . The classical theory of Poisson point processes (see for instance the “master formula” in [10]) shows that it is possible to decompose the law of $T(s)$ according to the time at which the jump over s occurs. More precisely, let us grow the accordion from $M_0 = (L_0, R_0) := (l, r)$ towards infinity, and let $M_t = (L_t, R_t)$ denote the top boundary arc of the accordion at time t .

For each given t , we can sample M_t and then define a point w on $\mathbb{R} \setminus [L_t, R_t]$ according to the measure $\zeta_{[R_t, L_t]}$. This is an infinite measure, but the mass m_t of the event that s is in the triangle (L_t, R_t, w) is finite, and furthermore, the definition of the Poisson point process, together with the fact that the triangulation is complete, ensures that $E(\int_0^\infty m_t dt) = 1$ (just because there is a.s. exactly one triangle in the accordion that contains s), given the fact that $s \notin T_i$.

Then (conditionally on (l, r) and on the fact that s lies above this arc), the distribution of $T(s)$ is described by

$$E(g(T(s)) \mid (l, r)) = \int_0^t dt E_{(L_0, R_0)=(l, r)} \left(\int \zeta_{[R_t, L_t]}(dw) g(\{L_t, R_t, w\}) 1_{\{s \in (L_t, R_t, w)\}} \right)$$

for any measurable bounded function g . For convenience, we are now going to assume that $g(T) = 0$ as soon as $s \notin T$, as this will enable us to incorporate the indicator function in g . As we anyway restrict ourselves to the case where $i \notin T(s)$, we assume as well that $g(T) = 0$ as soon as $i \in T$. Similarly, we will consider a measurable function f on the set of unmarked triangles, such that $f(T) = 0$ as soon as $i \notin T$ or $s \in T$.

If we now combine this with the description of the law of T_i , we get

$$E(g(T(s))f(T_i)) = \int_0^\infty dt c_0 \pi(d\alpha d\beta) \times E_{(L_0, R_0)=(\alpha, \beta)} \left(\int \zeta_{[\alpha, \beta]}(d\gamma) \zeta_{[R_t, L_t]}(dw) f(\{\alpha, \beta, \gamma\}) g(\{L_t, R_t, w\}) \right).$$

It now suffices to apply the reversibility of our Markov process, i.e. Proposition 14, which implies that this quantity is equal to

$$\int_0^\infty dt c_0 \pi(du dv) \times E_{(L_0^-, R_0^-)=(u, v)} \left(\int \zeta_{[L_t^-, R_t^-]}(d\gamma) \zeta_{[v, u]}(dw) f(\{\gamma, L_t^-, R_t^-\}) g(\{u, v, w\}) \right)$$

and to note that (just in the same way as before), this is precisely $E(g(T'_s)f(T'(i)))$, which completes the proof. \square

4.5. End of the proof of Theorem 2

Let us construct a triangulation of \mathbf{T} in the unit disk as follows. First sample $T(0)$ according to P_0 . In each of the three remaining domains O_1, O_2 and O_3 (see Fig. 3) sample independently an accordion tree started from $(u_2u_3), (u_3u_1)$ and (u_1u_2) respectively. This clearly defines a random complete triangulation \mathbf{T} . Furthermore, this construction ensures directly that \mathbf{T} is Markovian. What remains to be checked is its Möbius invariance.

As the definition also yields invariance under rotations around the origin, it suffices to check that for any given z_0 in $\mathbb{D} \setminus \{0\}$, the law of \mathbf{T} is invariant under the hyperbolic isometry φ_{z_0} in \mathbb{D} that interchanges 0 and z_0 .

In other words, define another random triangulation as follows. First sample $T'(z_0)$ according to P_{z_0} . Then, in each of the three remaining domains, sample independently an accordion tree. This defines a random triangulation \mathbf{T}' (and by construction, its law is that of $\varphi_{z_0}(\mathbf{T})$ because the image measure of P_0 under φ_{z_0} is P_{z_0} and by invariance of the accordion tree under Möbius transformations). Our goal is to prove that \mathbf{T} and \mathbf{T}' are identically distributed.

As in the proof of Lemma 4, in order to prove that \mathbf{T} and \mathbf{T}' are identically distributed, it suffices to see that for all $Z = \{0, z_0, z_1, \dots, z_n\}$ and for any tree structure Γ on Z , one has identity in law between $T(Z) 1_{T(Z) \in \mathcal{A}(Z, \Gamma)}$ and $T'(Z) 1_{T'(Z) \in \mathcal{A}(Z, \Gamma)}$.

But Lemma 15 shows readily that this is certainly the case as soon as 0 and z_0 are neighbors in Γ . Indeed, we can note that for both triangulations the marginal law of $(T(0), T(z_0))$ is the same (this is just the image in \mathbb{D} of Lemma 15), and that (on the event that the remaining points lie in the four outer connected components of the complement of these two triangles) the conditional distribution of $T(Z)$ given these two triangles is also the same by construction.

Suppose now that in Γ , 0 is at distance 2 of z_0 , i.e. both 0 and z_0 are neighbors of some z_l . Then we can use another random triangulation \mathbf{T}'' that is constructed just as \mathbf{T} and \mathbf{T}' , except that it is based at z_l , i.e., one starts defining $T''(z_l)$ using P_{z_l} etc. We have just showed that

- $T(Z) 1_{T(Z) \in \mathcal{A}(Z, \Gamma)}$ and $T''(Z) 1_{T''(Z) \in \mathcal{A}(Z, \Gamma)}$ are identically distributed,
- $T''(Z) 1_{T''(Z) \in \mathcal{A}(Z, \Gamma)}$ and $T'(Z) 1_{T'(Z) \in \mathcal{A}(Z, \Gamma)}$ are identically distributed,

which clearly settles the case when the graph-distance between 0 and z_0 is 2. The very same argument works without any problem when the distance is larger. \square

5. The Markovian hyperbolic quadrangulations, pentangulations etc.

It is of course natural to wonder if our results are specific to hyperbolic triangulations, or if they do have extensions to random partitions into other hyperbolic polygons. The answer is that, while triangulations are of course in some way special, there exist analogs to our Markovian triangulation when one replaces triangles by other polygons. In order to avoid any notational mess, the discussion in the present section will deliberately remain on a rather descriptive and maybe wordier style, and we will leave mathematical details to the interested reader.

Let us first focus on the case of *regular polygons*: We say that the sequence of distinct points x_1, \dots, x_n on the unit circle that is ordered anti-clockwise is a *hyperbolic n -gon* (and the n hyperbolic lines (x_1x_2) etc. are its boundary edges). We say that it is a *regular n -gon* if there exists a Möbius transformation ϕ such that $\phi(x_j) = e^{2\pi j/n}$ for all j in $\{1, \dots, n\}$. For instance, any (anti-clockwise ordered) triple x_1, x_2, x_3 is a regular 3-gon, but then, only one possible x_4 turns x_1, x_2, x_3, x_4 into a regular 4-gon (we will from now on call regular 4-gons *hyperbolic squares*). Note that any n -gon is obtained by glueing together $n - 2$ adjacent triangles (and triangles are all equivalent up to hyperbolic isometry). The μ -hyperbolic area of any n -gon is therefore always equal to n (with our normalization of the hyperbolic measure μ). A hyperbolic square is just the glueing of two “conformally symmetric” adjacent triangles and it has μ -area 2. Note also that an (unmarked) hyperbolic n -gon corresponds to n different possible marked n -gons (one has to choose which one of the corners is x_1).

It is trivial to extend the definition of the Markovian property to random tilings of \mathbb{D} into n -gons. The first extension of Theorem 2 goes as follows:

Theorem 16. *For any $n \geq 3$, there exists exactly one (law of a) Markovian Möbius-invariant complete partition of \mathbb{D} into regular n -gons.*

The case $n = 3$ is exactly Theorem 2, whereas when $n = 4$, the statement is that there exists a unique Markovian Möbius-invariant partition of \mathbb{D} into hyperbolic squares.

Sketch of the proof. The proof goes along similar lines to that of Theorem 2. Let us focus here on the case $n = 4$ (the proof of the other cases is quasi-identical and involves essentially no new idea) and just highlight the main differences with the proof of the case $n = 3$.

The origin square. First of all, note that there is a natural measure $\nu_{\circ}^{(4)}$ on marked hyperbolic squares: It is obtained from the measure ν_{\circ} on marked triangles (a, b, c) by looking at (a, b, c, d) where d is the symmetric image of b with respect to (ac) , i.e., the only point d such that (a, b, c, d) is a hyperbolic square. Clearly, this measure is invariant under circular relabeling (i.e., under $(a, b, c, d) \mapsto (b, c, d, a)$), and the marginal measure of any of its four halves (i.e. the triangles obtained by dropping one of the four points) is ν_{\circ} . It follows immediately that if a random tiling \mathbf{S} of the disk into hyperbolic squares is Möbius-invariant, then the measure on the marked square $S(0)$ containing the origin (if one marks it by choosing one of its corners uniformly at random among the four) is a multiple of $\nu_{\circ}^{(4)}$ restricted to those squares that contain the origin (the multiplicative constant

being chosen in such a way that the $\nu_o^{(4)}$ -mass of the set of marked squares containing the origin is equal to 1).

Uniqueness. Let us now consider the accordion of squares in the quadrangulation \mathbf{S} between 0 and 1 in \mathbb{D} . There is still no problem to define this object; it corresponds to the set of squares of \mathbf{S} that intersect the segment $[0, 1]$.

In the uniqueness part of the proof, we assume that \mathbf{S} is a Möbius-invariant Markovian decomposition into hyperbolic squares. One then first checks using the Markovian property and Möbius invariance that the corresponding half-plane accordion $\text{Acc}_{\mathbf{S}}(i, \infty)$ can be described via a Poisson point process of intensity given by some measure $\rho^{(4)}$ on hyperbolic squares in \mathbb{H} that have two adjacent corners at exactly -1 and 1 . The measure $\rho^{(4)}$ plays the same role as ρ in Section 3; it can be viewed as a measure on the set of pairs of points (x_1, x_2) in $\mathbb{R} \setminus [-1, 1]$ such that $-1, 1, x_1, x_2$ are ordered anti-clockwise on $\partial\mathbb{H}$, and $-1, 1, x_1, x_2$ is always a hyperbolic square (see Fig. 10). The latter means that

$$\frac{x_2 - 1}{x_2 + 1} = 2 \cdot \frac{x_1 - 1}{x_1 + 1}. \tag{7}$$




Fig. 10. Sketch of squares of type I_1 , II and I_2 .

Just as for ρ (using Möbius invariance), one checks that the image of this measure under any Möbius transformation that preserves -1 and 1 is a multiple of $\rho^{(4)}$. If one now defines $\rho_1^{(4)}$ to be the image measure of $\rho^{(4)}$ under the mapping $(x_1, x_2) \mapsto x_1$, it follows readily (using analogous arguments to those in Section 3 in order to show that ρ is in fact invariant under those transformations) that $\rho_1^{(4)}$ is equal to ρ itself (or to a constant multiple of ρ) and that it is also equal to the image of $\rho^{(4)}$ under $(x_1, x_2) \mapsto x_2$. Hence the measure $\rho^{(4)}$ is totally described by (7) and the fact that $\rho_1^{(4)} = \rho$ (up to a multiplicative constant).

The rest of the uniqueness part of the proof (the fact that this $\rho^{(4)}$ characterizes the law of the accordion, and that this law characterizes the law of \mathbf{S}) follows exactly the same arguments as in the case $n = 3$ in Section 3.

Existence. We can decompose the set of squares $(-1, 1, x_1, x_2)$ into three parts: the part I_2 where $1 < x_1 < x_2$ (and on this part, the only relevant information in order to construct the “future” of the accordion is the value of x_2 , which is distributed just as ρ on $[1, \infty)$), the part I_1 where $x_1 < x_2 < -1$ (and here, the only relevant information for the future of the accordion is x_1 , which is distributed just as ρ on $(-\infty, -1]$), and the third

part that we will call II with $x_2 < -1 < 1 < x_1$ (here, one needs to know both x_1 and x_2 to construct the future of the accordion).

We then consider a Poisson point process of squares with intensity $\rho^{(4)}$. For a square S in the upper half-plane (with four apexes on the real line), denote its leftmost apex by $a^l(S)$ and its rightmost one $a^r(S)$. This defines a Poisson point process $\{(a_t^l, a_t^r)\}$ (the square corresponding to (a_t^l, a_t^r) is of type I or II depending on whether one of the values $-a_t^l$ and a_t^r is 1 or not). Then, as in the case of the triangulations, one defines two pure jump processes (L_t, R_t) by putting

$$(L_t, R_t) = \varphi_{t-}^{-1}((a_t^l, a_t^r))$$

at all jump times (where φ_{t-}^{-1} is the affine map that maps $(-1, 1)$ onto (L_{t-}, R_{t-})). In contrast to the triangulation case, L and R can jump simultaneously (if the corresponding square is of type II). Note that I_1 and I_2 have infinite $\rho^{(4)}$ -mass. The key observation is that the $\rho^{(4)}$ -mass of type-II squares is finite (this can be for instance seen from the fact that this mass is equal to $\rho(\{4 \leq x_1\})$). Hence, when one constructs the half-plane accordion, the “times” (in the Poisson point process) at which one discovers a square where the two sides separating the square from 0 and from ∞ are not adjacent, form a discrete locally finite set. In other words, the accordion will be quite similar to the accordion with triangles except that:

- For each triangle (abc) in the accordion, one adds the fourth point d in order to turn it into a square, in such a way that d lies on the same side of the triangle (abc) as neither 0 nor infinity.
- One has squeezed in, in a Poissonian way, a discrete locally finite family of hyperbolic squares where the sides that separate 0 from infinity are not adjacent.

It then suffices to use the fact that the “times” at which one adds those type-II squares are in fact exactly the same (modulo time reversal) when one looks at the accordion in \mathbb{D} from -1 to 1 or from 1 to -1 (here, one uses the fact that the “time” parametrizations of the forward and backward accordions are the same, except for the time reversal). This will then indeed allow one to prove the Möbius invariance of the random decomposition into squares that is defined in this way. □

Let us now list further possible extensions:

- The previous proof shows that it is also very easy to construct a tiling of \mathbb{D} into a mixture of triangles and squares. The Möbius-invariant measure would then be described by a parameter $p \in [0, 1]$ (each value of p would correspond to a different distribution on tilings) that is equal to the probability that (in the tiling) the origin is in a triangle (and not in a square). More generally, for any distribution P on $\{3, 4, \dots, n_0\}$, one can define a random tiling of \mathbb{D} into regular n -gons (with varying n 's) in such a way that the probability that the origin lies in a regular n -gon is equal to $P(n)$. Conversely, these tilings will be the only complete Markovian tilings of \mathbb{D} into regular n -gons for $n \leq n_0$.
- So far, we have been dealing with tilings by regular n -gons. When $n = 3$, any triangle is a regular 3-gon, so that this was not a restriction, but for $n \geq 4$, it is. It raises the additional question of the existence and characterization of Markovian complete tilings of \mathbb{D}

into general n -gons. It turns out to be very easy to construct such tilings by non-symmetric tiles. For instance, we could be looking for tiles that are 4-gons with a prescribed conformal structure, i.e. such that the four ordered boundary points (x_1, x_2, x_3, x_4) can be mapped by some Möbius transformation onto one of the two 4-gons $(i, -1, -i, e^{i\theta})$ or $(i, -1, -i, e^{-i\theta})$ for some given $\theta \in (0, \pi/2)$ (note that the condition also has to be satisfied by (x_2, x_3, x_4, x_1)). This is then the unique Markovian tiling of the disk into hyperbolic rectangles of prescribed aspect ratio. Loosely speaking, the only main difference with the previously described case of tiling by squares is now that in the accordion, one tosses a fair coin for each 4-gon in order to choose between θ or $-\theta$ (i.e. if one discovers one of its “long” sides or one of its “short” sides).

The general statement about Markovian tilings by n -gons (of not necessarily prescribed hyperbolic structure) would then go along the following lines: Suppose that J is a Möbius-invariant measure on the set of ordered polygons P_n where

$$P_n = \{(x_1, \dots, x_n) \in \partial\mathbb{D}, \text{ ordered clockwise}\}$$

such that

1. The image measure of J via the projection $(x_1, \dots, x_n) \mapsto (x_1, x_2, x_3)$ of P_n onto P_3 is a multiple of ν_\circ .
2. The J -mass of the set of n -gons that contain the origin is 1 (this is just a matter of normalization).
3. J is invariant under $(x_1, x_2, \dots, x_n) \mapsto (x_2, x_3, \dots, x_n, x_1)$.

Then there exists a complete Möbius-invariant Markovian tiling of \mathbb{D} such that the n -gon containing 0 (if one marks it by choosing x_1 uniformly at random among the corners of the n -gon) is distributed according to the restriction of J to those n -gons that contain the origin. Conversely, this construction basically describes all possible complete Markovian Möbius-invariant tilings by n -gons.

Furthermore, all such measures J can be obtained via a product measure $\nu_\circ \otimes P$, where ν_\circ chooses the first three points x_1, x_2, x_3 , and P the hyperbolic position of the other $n - 3$ points with respect to (x_1, x_2, x_3) .

We leave out the details of the proofs, as well as generalizations to tilings into mixtures of n -gons for varying n 's, to the interested reader.

6. Concluding remarks

We conclude the paper with some remarks and open questions.

Hausdorff dimension. One property of our Markovian triangulation that we have collected on the way (we safely leave the details to the reader, recall Lemma 10 and Corollary 11) is:

Proposition 17. *The Hausdorff dimension of the closure of the union of all triangle boundaries in our Markovian hyperbolic triangulation is almost surely equal to 1.*

In other words, the “dimension” of the Markovian triangulation is not larger than 1, as opposed to other natural random triangulations that are “fatter” (see for instance [1, 8]).

Completeness. For non-complete Möbius-invariant triangulations, one can still make sense of the definition of the Markovian property in the following way: First note that the probability that $z \in \mathbb{D}$ is in some triangle of the triangulation is equal to some constant p_0 that does not depend on $z \in \mathbb{D}$ because of Möbius invariance, and if we assume that \mathbf{T} is not almost surely empty, p_0 is strictly positive. Then, we can condition on the event that $T(0)$ is not empty, and define the Markovian property as in Section 2.3.

It is easy to define non-trivial Möbius-invariant Markovian triangulations that are not complete. Here is an example: Pick $p \in [0, 1)$ and consider our random complete triangulation \mathbf{T} as defined in Theorem 2. Conditionally on \mathbf{T} , let $(d_T)_{T \in \mathbf{T}}$ be independent Bernoulli variables of parameter p indexed by the triangles of \mathbf{T} , and define

$$\mathbf{T}^{(p)} = \{T \in \mathbf{T} : d_T = 1\}.$$

In other words, we keep each triangle of \mathbf{T} with probability p . Clearly, $\mathbf{T}^{(p)}$ is a non-complete Möbius-invariant Markovian triangulation.

It is nevertheless possible to strengthen Theorem 2 replacing the completeness assumption by a density assumption. We recall that a triangulation \mathbf{T} is dense if the union of the triangles of \mathbf{T} is dense in \mathbb{D} . The statement then becomes: *There is a unique (law of a) dense Möbius-invariant triangulation of \mathbb{D} that has the spatial Markov property.*

In other words, any dense Möbius-invariant Markovian triangulation is in fact complete (see Remark 1). For expository reasons we chose to focus on Theorem 2 and its proof in the present paper, and we therefore do not include the proof of this last statement, but let us nevertheless give some brief hints for the interested reader: Most of the analysis goes along similar lines to that of Section 3. Density makes it possible to still define properly the processes (ℓ) , (r) and the jumps X . The fact that the triangulation is not complete however a priori allows the possibility to include a drift part in the subordinator-like evolution of r and ℓ . But it turns out that there is no Möbius-invariant way to define a non-zero drift term, as applying a hyperbolic isometry would change the relative speed of the drift near r and near ℓ (which is the usual problem when one tries to define, in a Möbius-invariant way, processes growing simultaneously at two different points).

Open questions. We conclude with three natural open questions:

1. It would be nice and enlightening to have an alternative construction of our Markovian Möbius-invariant triangulation (for instance via an auxiliary Poissonian model, some allocation idea or some statistical physics arguments) that would explain “directly” why it exists.
2. What are the natural discrete models that one could think of, and that would give rise to such Markovian hyperbolic triangulations in the scaling limit? Note that the Markovian property is somewhat reminiscent of Omer Angel’s exploration of percolation interfaces in random triangulations [2] (which however gives rise to a different scaling limit). This could also provide another (heuristic or rigorous) justification to the existence of our triangulation (just as the discrete percolation model “explains” the locality and reversibility properties of SLE₆ [9]).

3. What are the natural and Möbius-invariant dynamics (if they exist) on the set of triangulations that leave our measure invariant? One would for instance like to have a “continuous” evolution (in some appropriate topology) and such that the evolution of triangles that are far apart de-correlate fast with this distance. Each “individual” triangle could for instance follow some Brownian motion in the three-dimensional Lie group of Möbius transformations.

Acknowledgments. We thank Frédéric Paulin, Pierre Pansu and—of course—Itai Benjamini for stimulating discussions. We also thank the referees for their careful reading and useful comments. The authors acknowledge support of the Fondation Cino del Duca and of the Agence Nationale pour la Recherche (projet MAC2).

References

- [1] Aldous, D.: Triangulating the circle, at random. *Amer. Math. Monthly* **101**, 223–233 (1994) [Zbl 0804.52011](#) [MR 1264002](#)
- [2] Angel, O.: Growth and percolation on the uniform infinite planar triangulation. *Geom. Funct. Anal.* **13**, 935–974 (2003) [Zbl 1039.60085](#) [MR 2024412](#)
- [3] Bertoin, J.: *Lévy Processes*. Cambridge Tracts in Math. 121, Cambridge Univ. Press, Cambridge (1996) [Zbl 0861.60003](#) [MR 1406564](#)
- [4] Bertoin, J.: Subordinators: examples and applications. In: *Lectures on Probability Theory and Statistics (Saint-Flour, 1997)*, Lecture Notes in Math. 1717, Springer, Berlin, 1–91 (1999) [Zbl 0955.60046](#) [MR 1746300](#)
- [5] Bonahon, F.: Geodesic laminations on surfaces. In: *Laminations and Foliations in Dynamics, Geometry and Topology (Stony Brook, NY, 1998)*, Contemp. Math. 269, Amer. Math. Soc., Providence, RI, 1–37 (2001) [Zbl 0996.53029](#) [MR 1810534](#)
- [6] Bonahon, F.: *Low-Dimensional Geometry*. Student Math. Library 49, Amer. Math. Soc., Providence, RI (2009) [Zbl 1176.57001](#) [MR 2522946](#)
- [7] Chung, K.-L., Walsh, J. B.: *Markov Processes, Brownian Motion, and Time Symmetry*. 2nd ed., Springer, New York (2005) [Zbl 1082.60001](#) [MR 2152573](#)
- [8] Curien, N., Le Gall, J.-F.: Random recursive triangulations of the disk via fragmentation theory. *Ann. Probab.* **39**, 2224–2270 (2011) [Zbl 1252.60016](#) [MR 2932668](#)
- [9] Lawler, G., Schramm, O., Werner, W.: Values of Brownian intersection exponents I: Half-plane exponents. *Acta Math.* **187**, 237–273 (2001) [Zbl 1005.60097](#) [MR 1879850](#)
- [10] Revuz, D., Yor, M.: *Continuous Martingales and Brownian Motion*. Springer, Berlin (1991) [Zbl 0731.60002](#) [MR 1083357](#)
- [11] Schramm, O.: Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.* **118**, 221–288 (2000) [Zbl 0968.60093](#) [MR 1776084](#)
- [12] Sleator, D. D., Tarjan, R. E., Thurston, W. P.: Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.* **1**, 647–681 (1988) [Zbl 0653.51017](#) [MR 0928904](#)