

Supercritical causal maps : geodesics and simple random walk

Thomas BUDZINSKI *

July 7, 2019

Abstract

We study the random planar maps obtained from supercritical Galton–Watson trees by adding the horizontal connections between successive vertices at each level. These are the hyperbolic analog of the maps studied by Curien, Hutchcroft and Nachmias in [15], and a natural model of random hyperbolic geometry. We first establish metric hyperbolicity properties of these maps: we show that they admit bi-infinite geodesics and satisfy a weak version of Gromov-hyperbolicity. We also study the simple random walk on these maps: we identify their Poisson boundary and, in the case where the underlying tree has no leaf, we prove that the random walk has positive speed. Some of the methods used here are robust, and allow us to obtain more general results about planar maps containing a supercritical Galton–Watson tree.

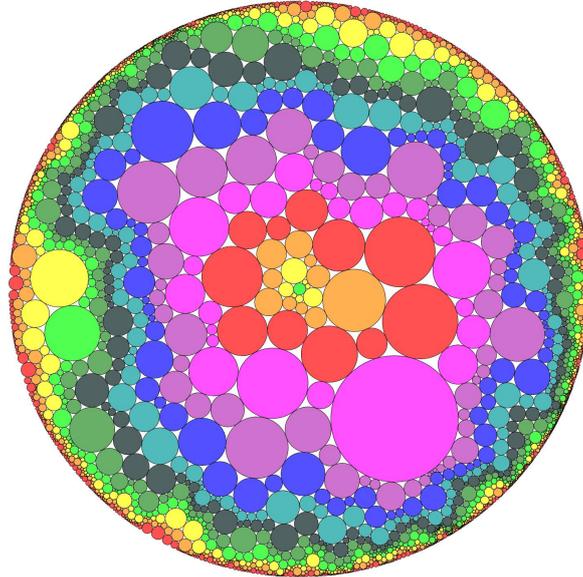


Figure 1: The circle packing of a causal triangulation constructed from a Galton–Watson tree with geometric offspring distribution of mean $3/2$. This was made with the help of the software CirclePack by Ken Stephenson.

*ENS Paris and Université Paris-Saclay, thomas.budzinski@ens.fr

Introduction

Causal maps and random hyperbolic geometry. *Causal triangulations* were introduced by theoretical physicists Ambjørn and Loll [2], and have been the object of a lot of numerical investigations. However, their rigorous study is quite recent [18, 15]. They are a discrete model of Lorentzian quantum gravity with one time and one space dimension where, in contrast with uniform random planar maps, time and space play asymmetric roles.

Here is the definition of the model. For any (finite or infinite) plane tree t , we denote by $\mathcal{C}(t)$ the planar map obtained from t by adding at each level the horizontal connections between consecutive vertices, as on Figure 2 (this includes an horizontal connection between the leftmost and rightmost vertices at each level). Our goal here is to study the graph $\mathcal{C}(T)$, where T is a supercritical Galton–Watson tree conditioned to survive.

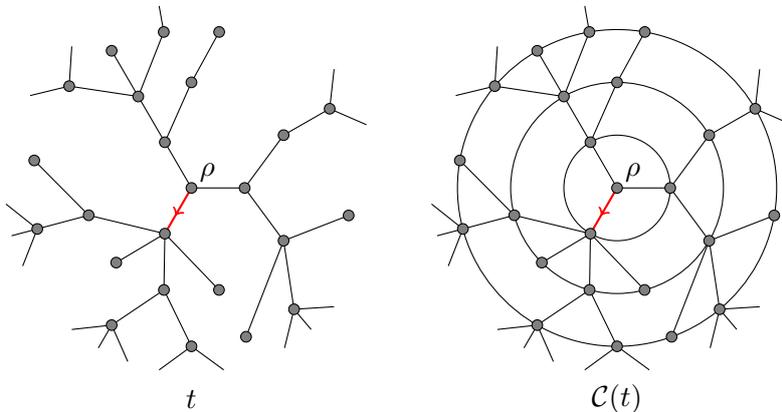


Figure 2: An infinite plane tree t and the associated causal map $\mathcal{C}(t)$. The edge in red joins the root vertex to its leftmost child.

This defines a new model of random "hyperbolic" graph. Several other such models have been investigated so far, such as supercritical Galton–Watson trees [24], Poisson–Voronoi tessellations of the hyperbolic plane [9], or the Planar Stochastic Hyperbolic Infinite Triangulations (PSHIT) of [13]. Many notions appearing in the study of these models are adapted from the study of Cayley graphs of nonamenable groups, and an important idea is to find more general versions of the useful properties of these Cayley graphs. Let us mention two such tools.

- For example, the three aforementioned models are all *stationary*, which means their distribution is invariant under rerooting along the simple random walk¹. This property generalizes the transitivity of Cayley graphs, and is a key tool to prove positive speed for the simple random walk on supercritical Galton–Watson trees [24, 1] or on the PSHIT [13]. More generally, in the context of stationary random graphs, general relations are known between the exponential growth rate, the speed of the random walk, and its asymptotic entropy, which is itself related to the Poisson boundary and the Liouville property. See [8, Proposition 3.6], which adapts classical results about Cayley graphs [20]. On the other hand, supercritical causal maps are not stationary, and it seems hard to find a stationary

¹This is not exactly true for Galton–Watson trees, but it is true for the closely related *augmented Galton–Watson trees*.

environment for the simple random walk². In absence of stationarity, we will be forced to use other properties of our graphs such as the independence properties given by the structure of Galton–Watson trees.

- Another important property in the study of random hyperbolic graphs is anchored expansion, which is a weaker version of nonamenability, and may be thought of as a natural generalization of nonamenability to random graphs. It is known to imply positive speed and heat kernel decay bounds of the form $\exp(-n^{1/3})$ for bounded-degree graphs [29]. This property also played an important role in the study of non-bounded-degree graphs such as Poisson-Voronoi tessellations of the hyperbolic plane [9], and the half-planar versions of the PSHIT [5]. However, we have not been able to establish this property for causal maps, and need once again to use other methods.

Supercritical causal maps. In all that follows, we fix an offspring distribution μ with $\sum_{i=0}^{\infty} i\mu(i) > 1$. Note that we do not require the mean number of children to be finite. We denote by T a Galton–Watson tree with offspring distribution μ conditioned to survive. The goal of this work is to study the maps $\mathcal{C}(T)$. We will study both large-scale metric properties of $\mathcal{C}(T)$, and the simple random walk on this map. All the results that we will prove show that $\mathcal{C}(T)$ has a hyperbolic flavour, which is also true for the tree T .

Metric hyperbolicity properties. The first goal of this work is to establish two metric hyperbolicity properties of $\mathcal{C}(T)$. We recall that a graph G is hyperbolic in the sense of Gromov if there is a constant $k \geq 0$ such that all the triangles are k -thin in the following sense. Let x, y and z be three vertices of G and $\gamma_{xy}, \gamma_{yz}, \gamma_{zx}$ be geodesics from x to y , from y to z and from z to x . Then for any vertex v on γ_{xy} , the graph distance between v and $\gamma_{yz} \cup \gamma_{zx}$ is at most k . However, such a strong, uniform statement usually cannot hold for random graphs. For example, if $\mu(1) > 0$, then $\mathcal{C}(T)$ contains arbitrarily large portions of the square lattice, which is not hyperbolic. Therefore, we suggest a weaker, "anchored" definition³.

Definition 1. Let M be a rooted planar map. We say that M is *weakly anchored hyperbolic* if there is a constant $k \geq 0$ such that the following holds. Let x, y and z be three vertices of M and γ_{xy} (resp. γ_{yz}, γ_{zx}) be a geodesic from x to y (resp. y to z, z to x). Assume the triangle formed by γ_{xy}, γ_{yz} and γ_{zx} surrounds the root vertex ρ . Then

$$d_M(\rho, \gamma_{xy} \cup \gamma_{yz} \cup \gamma_{zx}) \leq k.$$

Theorem 1 (Metric hyperbolicity of $\mathcal{C}(T)$). Let T be a supercritical Galton–Watson tree conditioned to survive, and let $\mathcal{C}(T)$ be the associated causal map.

1. The map $\mathcal{C}(T)$ is a.s. weakly anchored hyperbolic.

²See for example [27] for the particular case where the tree is the complete binary tree: the existence and uniqueness of a stationary environment are proved, but it is very difficult to say anything explicit about the distribution of this environment.

³The most natural definition would be to require that any geodesic triangle surrounding the root is k -thin, but this is still too strong (consider the triangle formed by root vertex and two vertices x, y in a large portion of square lattice).

2. The map $\mathcal{C}(T)$ a.s. admits bi-infinite geodesics, i.e. paths $(\gamma(i))_{i \in \mathbb{Z}}$ such that for any i and j , the graph distance between $\gamma(i)$ and $\gamma(j)$ is exactly $|i - j|$.

These two results are very robust and hold in a much more general setting that includes the PSHIT. In particular, the second point for the PSHIT answers a question of Benjamini and Tessera [11]. More general results are discussed in the end of this introduction.

Poisson boundary. The second goal of this work is to study the simple random walk on $\mathcal{C}(T)$ and to identify its Poisson boundary. First note that $\mathcal{C}(T)$ contains as a subgraph the supercritical Galton–Watson tree T , which is transient, so $\mathcal{C}(T)$ is transient as well. We recall the general definition of the Poisson boundary. Let G be an infinite, locally finite graph, and let $G \cup \partial G$ be a compactification of G , i.e. a compact metric space in which G is dense. Let also (X_n) be the simple random walk on G started from ρ . We say that ∂G is a *realization of the Poisson boundary* of G if the following two properties hold:

- (X_n) converges a.s. to a point $X_\infty \in \partial G$,
- every bounded harmonic function h on G can be written in the form

$$h(x) = \mathbb{E}_x [g(X_\infty)],$$

where g is a bounded measurable function from ∂G to \mathbb{R} .

We denote by ∂T the space of infinite rays of T . If $\gamma, \gamma' \in \partial T$, we write $\gamma \sim \gamma'$ if $\gamma = \gamma'$ or if γ and γ' are two "consecutive" rays in the sense that there is no ray between them. Then \sim is a.s. an equivalence relation for which countably many equivalence classes have cardinal 2 and all the others have cardinal 1. We write $\widehat{\partial T} = \partial T / \sim$. There is a natural way to equip $\mathcal{C}(T) \cup \widehat{\partial T}$ with a topology that makes it a compact space. We refer to Section 3.1 for the construction of this topology, but we mention right now that $\widehat{\partial T}$ is homeomorphic to the circle, whereas ∂T is homeomorphic to a Cantor set. The space $\mathcal{C}(T) \cup \widehat{\partial T}$ can be seen as a compactification of the infinite graph $\mathcal{C}(T)$. We show that this is a realization of its Poisson boundary.

Theorem 2 (Poisson boundary of $\mathcal{C}(T)$). Almost surely:

1. the limit $\lim(X_n) = X_\infty$ exists and its distribution has full support and no atoms in $\widehat{\partial T}$,
2. $\widehat{\partial T}$ is a realization of the Poisson boundary of $\mathcal{C}(T)$.

Note that, by a result of Hutchcroft and Peres [19], the second point will follow from the first one.

Positive speed. A natural and strong property shared by many models of hyperbolic graphs is the positive speed of the simple random walk. See for example [24] for supercritical Galton–Watson trees, and [13, 5] for the PSHIT or their half-planar analogs. The third goal of this work is to prove that the simple random walk on $\mathcal{C}(T)$ has a.s. positive speed. Unfortunately, we have only been able to prove it in the case where $\mu(0) = 0$, i.e. when the tree T has no leaf. We recall that (X_n) is the simple random walk on $\mathcal{C}(T)$, and denote by $d_{\mathcal{C}(T)}$ the graph distance on $\mathcal{C}(T)$.

Theorem 3 (Positive speed on $\mathcal{C}(T)$). If $\mu(0) = 0$ and $\mu(1) < 1$, then there is $v_\mu > 0$ such that

$$\frac{d_{\mathcal{C}(T)}(\rho, X_n)}{n} \xrightarrow[n \rightarrow +\infty]{\text{a.s.}} v_\mu.$$

However, we expect to still have positive speed if $\mu(0) = 0$. As mentioned above, this result is not obvious because of the lack of stationarity (for stationary graphs, the results of [8] show that positive speed is equivalent to being non-Liouville under some mild assumptions).

The critical case. We note that similar properties have been studied in the critical case in [15]. The results of [15] show that the geometric properties of causal maps are closer to those of uniform random maps than to those of the trees from which they were built. This contrasts sharply with the supercritical case, where the properties of the causal map are very close to those of the associated tree. More precisely, in the finite variance case, the distance between vertices at some fixed height r is $o(r)$, but $r^{1-o(1)}$. Moreover, the exponents describing the behaviour of the simple random walk are the same as for the square lattice, and different from the exponents we would obtain in a tree.

Robustness of the results and applications to other models. Another motivation to study causal maps is that many other models of random planar maps can be obtained by adding connections (and, in some cases, vertices) to a random tree. For example, the UIPT [6] or its hyperbolic variants the PSHIT [13] can be constructed from a reverse Galton–Watson tree or forest via the Krikun decomposition [21, 16, 12]. Among all the maps that can be obtained from a tree t in such a way that the branches of the tree remain geodesics, the causal map is the one with the "closest" connections, which makes it a useful toy model. The causal map may even provide general bounds for any map obtained from a fixed tree (we will see such applications in this paper, see also [14] for applications to uniform planar maps via the Krikun decomposition).

Here, the causal maps $\mathcal{C}(T)$ fit in a more general framework. We define a *strip* as an infinite, one-ended planar map s with exactly one infinite face, such that the infinite face has a simple boundary ∂s , and equipped with a root vertex on the boundary on the infinite face. If t is an infinite tree with no leaf and $(s_i)_{i \in \mathbb{N}}$ is a sequence of strips, let $\mathcal{M}(t, (s_i))$ be the map obtained by filling the (infinite) faces of t with the strips s_i (see Section 1 for a more careful construction). Some of our results can be generalized to random maps of the form $\mathcal{M}(\mathbf{T}, (s_i))$, where \mathbf{T} is a supercritical Galton–Watson tree with no leaf, and the s_i are strips (which may depend on \mathbf{T}).

By the backbone decomposition for supercritical Galton–Watson trees (that we recall in Section 1), the maps of the form $\mathcal{C}(T)$, where T is a supercritical Galton–Watson tree with leaves, are a particular case of this construction. The results of [12] prove that the PSHIT \mathbb{T}_λ can also be obtained by this construction: the tree \mathbf{T} is then the tree of infinite leftmost geodesics of \mathbb{T}_λ and has geometric offspring distribution.

We will show that Theorem 1 is very robust and applies to this general context, see Theorem 4. A particular case of interest are the PSHIT. In particular, point 2 of Theorem 1 for the PSHIT answers a question of Benjamini and Tessler [11].

As for causal maps, any map of the form $\mathcal{M}(\mathbf{T}, (s_i))$ contains the transient graph \mathbf{T} , so it is transient itself. Most of our proof of Theorem 2 can also be adapted to the general setting where the strips s_i are i.i.d. and independent of \mathbf{T} . However, Theorem 2 cannot be true if the strips s_i are too large (for example if themselves have a non-trivial Poisson boundary). On the other

	non-Liouville	$\widehat{\partial}\mathbf{T}$ is the Poisson boundary	positive speed
T	✓	✓	✓
$\mathcal{C}(T)$ (if $\mu(0) = 0$)	✓	✓	✓
$\mathcal{C}(T)$ (if $\mu(0) > 0$)	✓	✓	?
PSHIT	✓	✓	✓
\mathcal{M} with (S_i) i.i.d., recurrent, bounded-degree	✓	✓	✗
\mathcal{M} with (S_i) i.i.d., recurrent	✓	?	✗
\mathcal{M} with (S_i) i.i.d.	✓	✗	✗
general \mathcal{M}	✗	✗	✗

Figure 3: The symbol ✓ means that the property is proved in an earlier work or in this one. The symbol ? indicates properties that we believe to be true but did not prove in this paper, and the symbol ✗ means the property is false in general. See Section 5 for a quick description of some counterexamples.

hand, we can still show that the Poisson boundary is non-trivial. See Theorem 4 for a precise statement, and Figure 3 for a summary of the results proved in this paper and the results left to prove.

As we will see later, Theorem 4 is not strictly speaking more general than Theorem 2, since the strips used to construct $\mathcal{C}(T)$ from the backbone of T are not completely independent. On the other hand, once again, the PSHIT satisfy these assumptions (up to a root transformation, since the strip containing the root has a slightly different distribution). However, it was already known that the PSHIT are non-Liouville (see [13], or [3] for another identification of the Poisson boundary via circle packings). We also prove in [12], by a specific argument based on the peeling process, that $\widehat{\partial}\mathbf{T}$ is indeed a realization of the Poisson boundary in the case of the PSHIT.

Structure of the paper. The paper is structured as follows. In Section 1, we fix some definitions and notations that will be used in all the rest of this work, and recall the backbone decomposition of supercritical Galton–Watson trees. In Section 2, we investigate metric properties and establish Theorem 4, of which Theorem 1 is a particular case. Section 3 is devoted to the study of the Poisson boundary and to the proof of Theorems 2 and 4. In Section 4, we prove Theorem 3 about positive speed. Finally, in Section 5, we discuss some counterexamples related to Figure 3, and state a few conjectures.

Acknowledgments: I thank Nicolas Curien for his comments on earlier versions of this work, Arvind Singh for explanations about renewal theory, and Itai Benjamini for providing the reference [7]. I am grateful to the anonymous referee for pointing out that the third item of Theorem 4 could also be proved. I acknowledge the support of ANR Liouville (ANR-15-CE40-0013), ANR GRAAL (ANR-14-CE25-0014) and ERC GeoBrown (740943).

Contents

1 General framework and the backbone decomposition

7

2	Metric hyperbolicity properties	9
2.1	A hyperbolicity result about slices	10
2.2	Weak anchored hyperbolicity and bi-infinite geodesics	14
3	Poisson boundary	17
3.1	General setting	17
3.2	Transience away from the boundary in causal slices	18
3.3	Consequences on the Poisson boundary	23
3.4	Robustness of Proposition 6	25
4	Positive speed of the simple random walk	29
4.1	Sketch of the proof and definition of the half-plane model \mathcal{H}	29
4.2	An exploration method of \mathcal{H}	30
4.3	Quasi-positive speed in \mathcal{H}	39
4.4	Positive speed in \mathcal{H} via regeneration times	42
5	Counterexamples and open questions	44
A	The regeneration structure	46

1 General framework and the backbone decomposition

The goal of this first section is to give definitions and notations, and to make a few useful remarks that will be needed in all the paper. All our constructions will be based on infinite, locally finite plane trees. We insist that the plane tree structure is important to define the associated causal map. We will use normal letters to denote general infinite trees, and bold letters like \mathbf{t} for trees with no leaf. All the trees will be rooted at a vertex ρ . If v is a vertex of a tree t , we denote by $h(v)$ its distance to the root, which we will sometimes call its *height*. A *ray* in a tree t is an infinite sequence $(\gamma(i))_{i \geq 0}$ of vertices such that $\gamma(0) = \rho$, and $\gamma(i+1)$ is a child of $\gamma(i)$ for every $i \geq 0$. If t is an infinite tree, the *backbone* of t is the union of its rays, i.e. the set of the vertices of t that have infinitely many descendants. We will denote it by $\mathbf{B}(t)$, and we note that $\mathbf{B}(t)$ is always an infinite tree with no leaf.

We recall that if t is an infinite plane tree, then $\mathcal{C}(t)$ is the map obtained from t by adding horizontal edges at every height between consecutive vertices. We also define the *causal slice* $\mathcal{S}(t)$ associated to t , which will be used a lot in all that follows. Let γ_ℓ (resp. γ_r) be the leftmost (resp. rightmost) infinite ray of $\mathbf{B}(t)$. Then $\mathcal{S}(t)$ is the map obtained from t by deleting all the vertices on the left of γ_ℓ and on the right of γ_r , and by adding the same horizontal edges as for $\mathcal{C}(t)$ between the remaining vertices, except the edge between γ_ℓ and γ_r at each level (cf. Figure 4). The union of γ_ℓ and γ_r is the *boundary* of \mathcal{S} , and is written $\partial\mathcal{S}$.

In all this work, μ will denote a supercritical offspring distribution, i.e. a distribution satisfying $\sum_{i \geq 0} i\mu(i) > 1$, and T will be a Galton–Watson tree with offspring distribution μ conditioned to survive. For every $n \geq 0$, we will denote by Z_n the number of vertices of T at height n . We will write \mathcal{C} for $\mathcal{C}(T)$ and \mathcal{S} for $\mathcal{S}(T)$, unless stated otherwise.

If v is a vertex of $\mathbf{B}(T)$, we will denote by $T[v]$ the tree of descendants of v in T , and by $\mathcal{S}[v]$ the causal slice associated to $T[v]$. An important consequence of the backbone decomposition stated below is that for each $v \in T$, conditionally on $v \in \mathbf{B}(T)$, the slice $\mathcal{S}[v]$ has the same

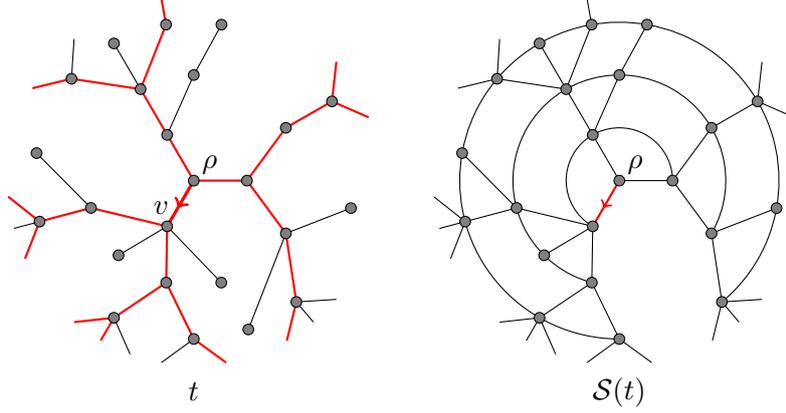


Figure 4: The same infinite tree t as on Figure 2 and the associated causal slice $\mathcal{S}(t)$. Note that two vertices have been deleted. On the left part, the backbone of t is in red. We have $c(v) = 4$, $c_{\mathbf{B}}(v) = 2$ and $A_v = \{2, 4\}$.

distribution as \mathcal{S} . Moreover, these slices are independent for base points that are not ancestors of each other.

If G is a graph rooted at a vertex ρ , we will denote by d_G its graph distance and by $B_r(G)$ (resp. $\partial B_r(G)$) the set of vertices of G at distance at most r (resp. exactly r) from ρ . Note that the vertices of $B_r(T)$ and of $B_r(\mathcal{C})$ are the same. For any vertex v of T , we also denote by $c_T(v)$ (or by $c(v)$ when there is no ambiguity) the number of children of v in T . Note that the degree of v in \mathcal{C} is equal to $c_T(v) + 3$ if $v \neq \rho$, and to $c_T(v)$ if $v = \rho$. For every graph G and every vertex v of G , we will denote by $P_{G,v}$ the distribution of the simple random walk on G started from v .

We now recall the backbone decomposition for supercritical Galton–Watson trees conditioned to survive, as it appears e.g. in [23]. Let f be the generating function of μ , i.e. $f(x) = \sum_{i \geq 0} \mu(i)x^i$. Let also q be the extinction probability of a Galton–Watson tree with offspring distribution μ , i.e. the smallest fixed point of f in $[0, 1]$. We define \mathbf{f} and \tilde{f} by

$$\mathbf{f}(s) = \frac{f(q + (1 - q)s) - q}{1 - q} \quad \text{and} \quad \tilde{f}(s) = \frac{f(qs)}{q} \quad (1)$$

for $s \in [0, 1]$. Then \mathbf{f} is the generating function of a supercritical offspring distribution μ with $\mu(0) = 0$, and \tilde{f} is the generating function of a subcritical offspring distribution $\tilde{\mu}$. For every vertex x of $\mathbf{B}(T)$, we denote by A_x the set of indices $1 \leq i \leq c(x)$ such that the i -th child of x is in the backbone, and we write $c_{\mathbf{B}}(x) = |A_x|$. Finally, let $\mathbf{B}'(T)$ be the set of vertices y of T such that the parent of y is in $\mathbf{B}(T)$, but y is not. The following result characterizes entirely the distribution of T .

Theorem 4. 1. The tree $\mathbf{B}(T)$ is a Galton–Watson tree with offspring distribution μ .
 2. Conditionally on $\mathbf{B}(T)$, the variables $c(x) - c_{\mathbf{B}}(x)$ are independent, with distribution characterized by

$$\mathbb{E} \left[s^{c(x) - c_{\mathbf{B}}(x)} \mid \mathbf{B}(T) \right] = \frac{f^{(c_{\mathbf{B}}(x))}(qs)}{f^{(c_{\mathbf{B}}(x))}(q)}$$

for every $s \in [0, 1]$, where $f^{(k)}$ stands for the k -th derivative of f .

3. Conditionally on $\mathbf{B}(T)$ and the variables $c(x)$ for $x \in \mathbf{B}(T)$, the sets A_x for $x \in \mathbf{B}(T)$ are independent and, for every x , the set A_x is uniformly distributed among all the subsets of $\{1, 2, \dots, c(x)\}$ with $c_{\mathbf{B}}(x)$ elements.
4. Conditionally on everything above, the trees $T[y]$ for $y \in \mathbf{B}'(T)$ are independent Galton–Watson trees with offspring distribution $\tilde{\mu}$.

In particular, this decomposition implies that, for every $h \geq 1$, conditionally on $B_h(T)$ and on the set $\partial B_h(T) \cap \mathbf{B}(T)$, the trees $T[x]$ for $x \in \partial B_h(T) \cap \mathbf{B}(T)$ are i.i.d. copies of T . Therefore, the slices $\mathcal{S}[x]$ for $x \in \partial B_h(T) \cup \mathbf{B}(T)$ are i.i.d. copies of \mathcal{S} . This "self-similarity" property of \mathcal{S} will be used a lot later on.

We end this section by adapting these notions to the more general setting of strips glued in the faces of a tree with no leaf. We recall that a *strip* is an infinite, one-ended planar map with an infinite, simple boundary, such that all the faces except the outer face have finite degree. A strip is also rooted at a root vertex on its boundary. Let \mathbf{t} be an infinite plane tree with no leaf. We draw \mathbf{t} in the plane in such a way that its edges do not intersect (except at a common endpoint), and every compact subset of the plane intersects finitely many vertices and edges. Then \mathbf{t} separates the plane into a countable family $(f_i)_{i \geq 0}$ of faces, where f_0 is the face delimited by the leftmost and the rightmost rays of \mathbf{t} , and the other faces are enumerated in a deterministic fashion. For every index $i \geq 0$, we denote by ρ_i the lowest vertex of \mathbf{t} adjacent to f_i , and by h_i its height. Note that this vertex is always unique. On the other hand, for every vertex v of \mathbf{t} , there are exactly $c_{\mathbf{t}}(v) - 1$ faces f_i such that $\rho_i = v$.

Let $(s_i)_{i \geq 0}$ be a family of random strips. We denote by $\mathcal{M}(\mathbf{t}, (s_i)_{i \geq 0})$ the infinite planar map obtained by gluing s_i in the face f_i for every $i \geq 0$, in such a way that the root vertex of s_i coincides with ρ_i for every i . We also denote by $\mathcal{S}(\mathbf{t}, (s_i)_{i \geq 0})$ the map obtained by gluing s_i in the face f_i for every $i > 0$ (this is a map with an infinite boundary analog to the slice \mathcal{S}). If v is a vertex of \mathbf{t} , we also define the "slice of descendants" of v as the map enclosed between the leftmost and the rightmost rays of \mathbf{t} started from v . We denote it by $\mathcal{S}(\mathbf{t}, (s_i)) [v]$.

We note that causal maps are a particular case of this construction. This is trivial for supercritical Galton–Watson trees with no leaf. Thanks to Theorem 4, this can be extended to the case $\mu(0) > 0$, with $\mathbf{B}(T)$ playing the role of \mathbf{t} . This time, however, the strips are random, but they are not independent. Indeed, if v is a vertex of $\mathbf{B}(T)$ and w one of its children in $\mathbf{B}(T)$, the children of v on the left of w and the children on the right of w belong to different strips. However, by points 2 and 3 of Theorem 4, the numbers of children on the left and on the right are not independent, except in some very particular cases (for example if μ is geometric).

In what follows, we will study maps of the form $\mathcal{M}(\mathbf{T}, (s_i)_{i \geq 0})$, where \mathbf{T} is a supercritical Galton–Watson tree with no leaf. We notice right now that if the strips s_i are random and i.i.d., then the slice $\mathcal{S}(\mathbf{T}, (s_i))$ has the same self-similarity property as the causal slices of the form $\mathcal{S}(T)$. Let $h > 0$. We condition on $B_h(\mathbf{T})$ and on all the strips s_i such that $h_i \leq h - 1$. Then the trees $\mathbf{T}[v]$ for $v \in \partial B_h(\mathbf{T})$ are independent copies of \mathbf{T} , so the slices $\mathcal{S}(\mathbf{T}, (s_i)) [v]$ for $v \in \partial B_h(\mathbf{T})$ are i.i.d. copies of $\mathcal{S}(\mathbf{T}, (s_i))$. This will be useful in Section 3.

2 Metric hyperbolicity properties

The goal of this section is to prove the following result, of which Theorem 1 is a particular case. Note that we make no assumption about the strips s_i below. In particular, they may be

deterministic or random, and may depend on the tree \mathbf{T} .

Theorem 4. Let \mathbf{T} be a supercritical Galton–Watson tree with no leaf, and let (s_i) be a sequence of strips. Then:

1. the map $\mathcal{M}(\mathbf{T}, (s_i))$ is a.s. weakly anchored hyperbolic,
2. the map $\mathcal{M}(\mathbf{T}, (s_i))$ a.s. admits bi-infinite geodesics.

In all this section, we will only deal with the general case of $\mathcal{M}(\mathbf{T}, (s_i))$ where \mathbf{T} is a supercritical Galton–Watson tree with no leaf and the s_i are strips. We will write \mathcal{M} for $\mathcal{M}(\mathbf{T}, (s_i))$ and \mathcal{S} for $\mathcal{S}(\mathbf{T}, (s_i))$. Our main tool will be the forthcoming Proposition 2, which roughly shows that \mathcal{S} is hard to cross horizontally at large heights.

2.1 A hyperbolicity result about slices

We call γ_ℓ and γ_r the left and right boundaries of \mathcal{S} , and ρ its root (note that γ_ℓ and γ_r may have an initial segment in common near ρ). Both points of Theorem 4 will be consequences of the following hyperbolicity result about \mathcal{S} .

Proposition 2. There is a (random) $K \geq 0$ such that any geodesic in \mathcal{S} from a point on γ_ℓ to a point on γ_r contains a point at distance at most K from ρ .

We first give a very short proof of this proposition in the particular case of a causal slice of the form $\mathcal{S}(T)$. Let $i, j > 0$ and let γ be a geodesic in $\mathcal{S}(T)$ from $\gamma_\ell(i)$ to $\gamma_r(j)$. Let v_0 be the lowest point of γ , and let h_0 be the height of v_0 . By the structure of $\mathcal{S}(T)$, each step of γ is either horizontal or vertical. Since the height varies by at most 1 at each vertical step, we need at least $(i - h_0) + (j - h_0)$ vertical steps. Moreover, for every $h \geq 0$, let $Z_h^{\mathbf{B}}$ be the number of vertices of $\mathbf{B}(T)$ at height h . Then γ needs to cross all the trees $T[x]$ for $x \in \mathbf{B}(T)$ at height h_0 , so γ contains at least $Z_{h_0}^{\mathbf{B}}$ horizontal steps. On the other hand, γ is a geodesic so it is shorter than the "obvious" path following γ_ℓ from $\gamma_\ell(i)$ to ρ and then γ_r until $\gamma_r(j)$. Therefore, we have

$$i + j - 2h_0 + Z_{h_0}^{\mathbf{B}} \geq |\gamma| \geq i + j,$$

so $Z_{h_0}^{\mathbf{B}} \leq 2h_0$. However, $Z^{\mathbf{B}}$ has a.s. exponential growth, so this inequality only holds for finitely many values of h_0 , so h_0 is bounded independently of i and j .

To generalize this proof, there are two obstacles: first, the branches of the tree are no longer geodesics in \mathcal{S} in the general case, so a single step may change the height by more than one. Second, an edge of \mathcal{S} can play the role both of a vertical and a horizontal step if it crosses a strip and joins two vertices of \mathbf{T} at different heights. However, we can still cross at most one strip per step in this way.

In order to prove the general Proposition 2, we first state a lemma showing roughly that if a path in \mathbf{T} with nondecreasing height stays at height h during a time subexponential in h , it cannot cross \mathcal{S} from left to right.

More precisely, we fix a sequence of positive integers $(u_i)_{i \geq 0}$ and a height $k \geq 0$. Let $x \in \partial B_k(\mathbf{T})$. We define by induction two sequences $(y_i)_{i \geq k}$ and $(z_i)_{i \geq k}$ of vertices of \mathbf{T} with $y_i, z_i \in \partial B_i(\mathbf{T})$ as follows (see Figure 5 for an example):

- (i) $y_k = x$,

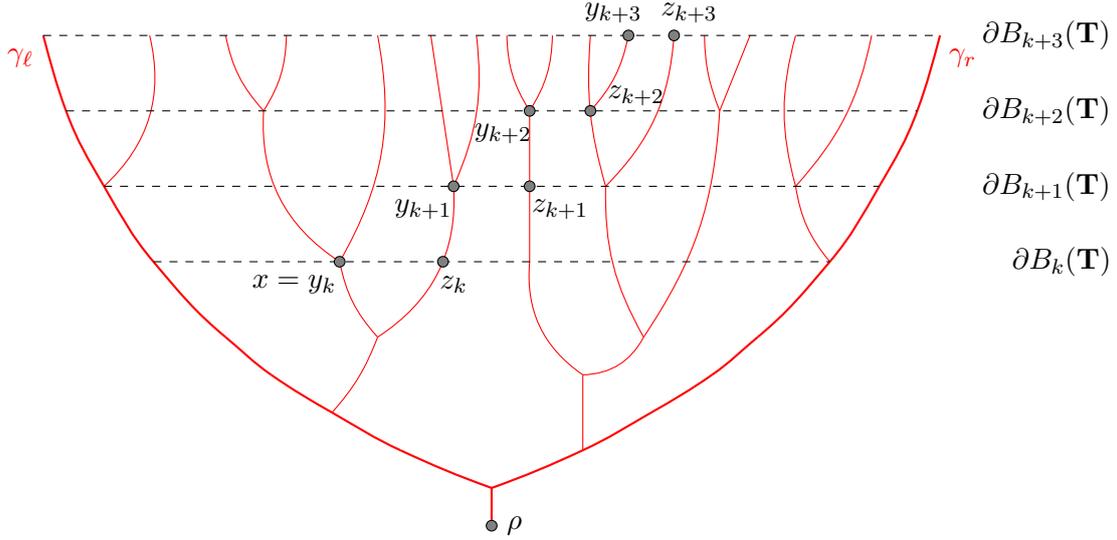


Figure 5: The sequences $(y_i)_{i \geq k}$ and $(z_i)_{i \geq k}$. Here we have taken $u_i = 1$ for every i . The tree \mathbf{T} is in red.

- (ii) for every $i \geq k$, if there are at least u_i vertices on the right of y_i on $\partial B_i(\mathbf{T})$ (y_i excluded), then z_i is the u_i -th such vertex,
- (iii) if there are less than u_i vertices of $\partial B_i(\mathbf{T})$ on the right of y_i , the sequences (y_i) and (z_i) are killed at time i ,
- (iv) for every $i \geq k$, if $z_i \notin \gamma_r$, the vertex y_{i+1} is the rightmost child of z_i in \mathbf{T} . If $z_i \in \gamma_r$, both sequences are killed.

We call $(y_i)_{i \geq k}$ and $(z_i)_{i \geq k}$ the *sequences escaping from x on the right*. We say that x is *u -far from γ_r* if the sequences (y_i) and (z_i) survive, that is, y_i and z_i are well-defined and do not hit γ_r for all $i \geq k$.

Note that being u -far from γ_r is a monotonic property: if we shift the point x to the left, then the points y_i and z_i are also shifted to the left. Hence, if a point $x \in \partial B_k(\mathbf{T})$ is u -far from γ_r and $x' \in \partial B_k(\mathbf{T})$ lies on the left of x , then x' is also u -far from γ_r . We can similarly define the sequences escaping on the left, and a vertex u -far from γ_ℓ .

Lemma 3. Assume u is subexponential, i.e. $u_i = o(c^i)$ for every $c > 1$. Then there is a (random) K such that for any $k \geq K$, the vertex $\gamma_\ell(k)$ is u -far from γ_r and the vertex $\gamma_r(k)$ is u -far from γ_ℓ .

Proof. The idea is to reduce the proof to the study of a supercritical Galton–Watson process where u_i individuals are killed at generation i . It is enough to show that $\gamma_\ell(k)$ is u -far from γ_r for k large enough. Note that if $\gamma_\ell(k)$ is u -far from γ_r and $(y_i)_{i \geq k}$ is its sequence escaping on the right, then $(y_i)_{i \geq k+1}$ is the sequence escaping on the right of y_{k+1} , so y_{k+1} is also u -far from γ_r and, by monotonicity, so is $\gamma_\ell(k+1)$. Therefore, it is enough to show that there is $k \geq 0$ such that $\gamma_\ell(k)$ is u -far from γ_r .

Let $k \geq 0$ and $x = \gamma_\ell(k)$. Let $(y_i)_{i \geq k}$ and $(z_i)_{i \geq k}$ be the sequences escaping from x on the right. We also denote by Z_i^k the number of vertices of $\partial B_i(\mathbf{T})$ lying (strictly) on the right of z_i . We first remark that the evolution of the process Z^k can be described explicitly. We have $Z_k^k = Z_k - 1$. Moreover, we recall that μ is the offspring distribution of \mathbf{T} . Conditionally on $(Z_k^k, Z_{k+1}^k, \dots, Z_i^k)$, the variable Z_{i+1}^k has the same distribution as

$$\left(\left(\sum_{j=1}^{Z_i^k} X_{i,j} \right) - u_{i+1} \right)^+,$$

where the $X_{i,j}$ are i.i.d. with distribution μ . Moreover, the process Z^k is killed when it hits 0. To prove our lemma, it is enough to show that $\mathbb{P}(\forall i \geq k, Z_i^k > 0)$ goes to 1 as k goes to $+\infty$. Since the process Z describing the number of individuals at each generation is a supercritical Galton–Watson process, there is a constant $c > 1$ such that

$$\mathbb{P}(\partial B_k(\mathbf{T}) > c^k) \xrightarrow[k \rightarrow +\infty]{} 1.$$

Therefore, by a monotonicity argument, it is enough to prove that

$$\mathbb{P}(\forall i \geq k, Z_i^k > 0 | Z_k^k = \lfloor c^k \rfloor) \xrightarrow[k \rightarrow +\infty]{} 1.$$

To prove this, we show that Z^k dominates a supercritical Galton–Watson process. Let $\delta > 0$ be such that $(1 - \delta) \sum_{i \geq 0} i\mu(i) > 1$. Let also Z^* be the Markov chain defined by

$$Z_k^* = \lfloor c^k \rfloor \quad \text{and} \quad Z_{i+1}^* = \left(\sum_{j=1}^{Z_i^*} X_{i,j} \right) - N_{i+1},$$

where the $X_{i,j}$ are i.i.d. with distribution μ and, conditionally on $(X_{i,j})_{i,j \geq 0}$ and $(N_j)_{1 \leq j \leq i}$, the variable N_{i+1} has binomial distribution with parameters δ and $\sum_{j=1}^{Z_i^*} X_{i,j}$. In other words, Z^* is a Galton–Watson process in which at every generation, right after reproduction, every individual is killed with probability δ . By our choice of δ , the process Z^* is a supercritical Galton–Watson process and, on survival, grows exponentially. By an easy large deviation argument, we have

$$\mathbb{P}(\forall i \geq k, N_i \geq u_i | Z_k^* = \lfloor c^k \rfloor) \xrightarrow[k \rightarrow +\infty]{} 1.$$

If this occurs, then $Z_i^* \leq Z_i^k$ for every $i \geq k$ (by an easy induction on i), so

$$\mathbb{P}(\forall i \geq k, Z_i^k \geq Z_i^* > 0 | Z_k^k = \lfloor c^k \rfloor) \xrightarrow[k \rightarrow +\infty]{} 1$$

and the lemma follows. \square

The proof of Proposition 2 given Lemma 3 only relies on deterministic considerations.

Proof of Proposition 2. We note that for any $a \in \mathbf{T}$, the slice $\mathcal{S}[a]$ is of the form $\mathcal{S}(\mathbf{T}[a], (s'_i))$ where $\mathbf{T}[a]$ is a supercritical Galton–Watson tree with no leaf, so we can apply Lemma 3 to $\mathcal{S}[a]$.

Now let a, a' be two vertices of \mathbf{T} , neither of which is an ancestor of the other, as on Figure 6. Then $\mathcal{S}[a]$ and $\mathcal{S}[a']$ are disjoint. Without loss of generality, we may assume that $\mathcal{S}[a]$ lies on the right of $\mathcal{S}[a']$. Let K (resp. K') be given by the conclusion of Lemma 3 for $\mathcal{S}[a]$ (resp. $\mathcal{S}[a']$) and

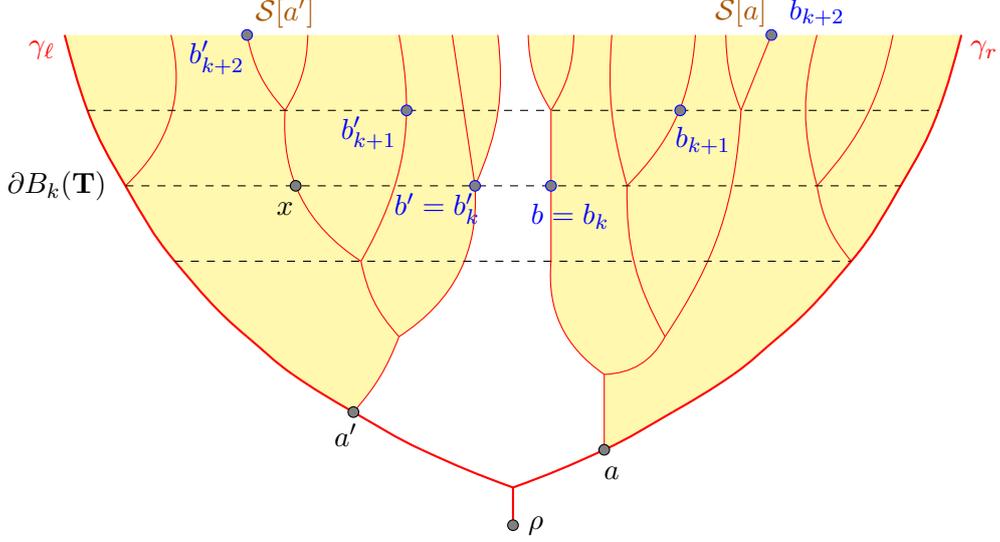


Figure 6: Proof of Proposition 2: the point b is u -far from γ_r and b' is u -far from γ_ℓ , so any point on $\partial B_k(\mathbf{T})$ is u -far from either γ_r or γ_ℓ . Here we have taken $u_i = 1$.

$u_i = 2(i + \max(d(\rho, a), d(\rho, a'))) + 1$. We take $K'' = \max(d(\rho, a) + K + 1, d(\rho, a') + K' + 1)$. We consider a geodesic γ from a vertex $\gamma_\ell(m)$ to a vertex $\gamma_r(n)$ in \mathcal{S} . Let k be the minimal height of $\gamma \cap \mathbf{T}$. We assume $k > K''$, and we will get a contradiction.

Let b be the leftmost vertex of $\mathcal{S}[a] \cap \mathbf{T}$ that lies at height k , and let b' be the rightmost vertex of $\mathcal{S}[a'] \cap \mathbf{T}$ that lies at height k (cf. Figure 6). By Lemma 3 and our choice of K'' , the point b is \tilde{u} -far from γ_r in \mathcal{S} and b' is \tilde{u} -far from γ_ℓ in \mathcal{S} , where $\tilde{u}_i = 2i + 1$ (the change of the sequence u to the sequence \tilde{u} is due to the fact that the distances to the root are not the same in $\mathcal{S}[a]$ and \mathcal{S} , so the sequence needs to be shifted). But any vertex of $\partial B_k(\mathbf{T})$ lies either on the left of b or on the right of b' , so it is either \tilde{u} -far from γ_ℓ or from γ_r (see Figure 6). In particular, let x be the first point of γ lying on \mathbf{T} at height k . We may assume that x is \tilde{u} -far from γ_r , the other case can be treated in the same way.

Recall that n is the height of the endpoint of γ . For every $k \leq i \leq n$, let

$$j_i = \max \{j \in \llbracket 0, |\gamma| \rrbracket \mid \gamma(j) \in \mathbf{T} \text{ and } h(\gamma(j)) \leq i\},$$

and let $x_i = \gamma(j_i)$. Note that we have $h(x_i) \leq i$, but since the height can increase by more than 1 in one step, the inequality may be strict.

Let also $(y_i)_{i \geq k}$ and $(z_i)_{i \geq k}$ be the sequences escaping from x on the right in \mathcal{S} (for $\tilde{u}_i = 2i + 1$). By our assumption that x is \tilde{u} -far from γ_r , these sequences are well-defined and do not hit γ_r . Moreover, for every $k \leq i \leq n$, the vertices x_i and $z_{h(x_i)}$ are both in \mathbf{T} and at the same height $h(x_i)$. We claim that for every $i \geq k$, the vertex x_i lies strictly on the left of the vertex $z_{h(x_i)}$. This is enough to prove the proposition, since then x_n cannot lie on γ_r .

We show this claim by induction on i , and we start with the case $i = k$. Let j_* be the index such that $x = \gamma(j_*)$. The vertices $x = \gamma(j_*)$ and $x_k = \gamma(j_k)$ both lie on $\partial B_k(\mathbf{T})$, so the distance in \mathcal{S} between them is at most $2k$. Hence, since γ is a geodesic, we have $|j_k - j_*| \leq 2k$. Now, we consider the slices $\mathcal{S}[v]$ for $v \in \partial B_k(\mathbf{T})$. These slices are disjoint and, by definition of k , the

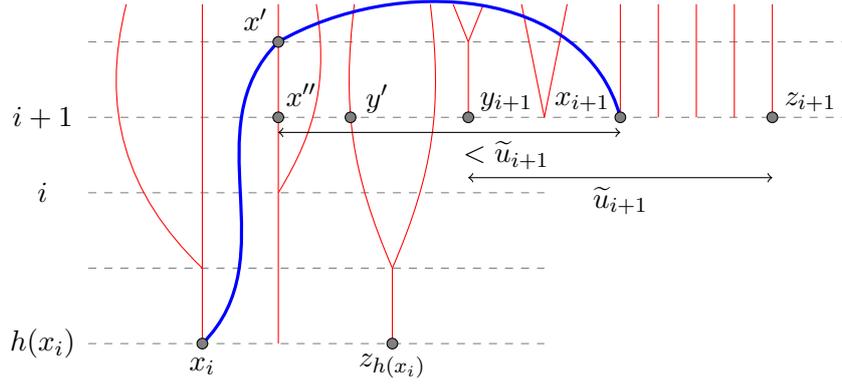


Figure 7: The induction step in the proof of Proposition 2. The path γ between x_i and x_{i+1} is in blue. Branches of \mathbf{T} are in red. The dashed lines are not edges of \mathcal{S} , but indicate the height. We see that x'' is on the left of y' which is on the left of y_{i+1} , so x_{i+1} is on the left of z_{i+1} .

path $(\gamma(j))_{j_* \leq j \leq j_k}$ does not cross \mathbf{T} below height k , so it cannot intersect more than $2k < \tilde{u}_k$ of these slices, which implies that x_k lies on the left on z_k .

We now move on to the induction step. We assume x_i lies strictly on the left of $z_{h(x_i)}$. We recall that $h(x_{i+1}) \leq i + 1$, and split the proof in two cases.

- If $h(x_{i+1}) < i + 1$, then x_{i+1} is the last point of γ at height at most $i + 1$, so it is also the last point of γ at height at most i , so $x_{i+1} = x_i$. In particular, it is strictly on the left of $z_{h(x_{i+1})} = z_{h(x_i)}$.
- If $h(x_{i+1}) = i + 1$, we need to introduce some more notation that is summed up on Figure 7. We denote by x' the first point of γ after x_i that belongs to \mathbf{T} (note that $h(x') \geq i + 1$ by definition of x_i), and by x'' the ancestor of x' at height $i + 1$. Let also y' be the leftmost descendant of $z_{h(x_i)}$ at height $i + 1$. Note that by construction of the sequences (y_i) and (z_i) , the vertex y' is on the left of y_{i+1} . We know that between x_i and x' , the path γ does not cross \mathbf{T} , so x_i and x' must be adjacent to the same strip. Since x_i is strictly on the left of $z_{h(x_i)}$, it implies that x' lies on the left of any descendant of $z_{h(x_i)}$, so x'' is on the left of y' , and therefore on the left of y_{i+1} . Moreover, by the definition of x_i , all the vertices of γ between x_i and x_{i+1} that belong to \mathbf{T} have height at least $i + 1$. By the same argument as before, the length of the part of γ between x_i and x_{i+1} is at most $h(x_i) + h(x_{i+1}) \leq 2(i + 1) < \tilde{u}_{i+1}$. Hence, this part cannot cross \tilde{u}_{i+1} of the slices $\mathcal{S}[v]$ with $v \in \partial B_{i+1}(\mathbf{T})$, so the distance between x'' and x_{i+1} along $\partial B_{i+1}(\mathbf{T})$ is less than \tilde{u}_{i+1} . Since x'' is on the left of y_{i+1} and z_{i+1} is at distance \tilde{u}_{i+1} on the right of y_{i+1} , it follows (see again Figure 7) that x_{i+1} is strictly on the left of z_{i+1} , which concludes the induction and the proof of the proposition. □

2.2 Weak anchored hyperbolicity and bi-infinite geodesics

We can now deduce Theorem 4 from Proposition 2.

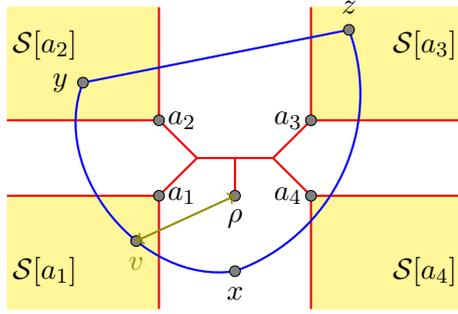


Figure 8: Illustration of the proof of point 1 of Theorem 4. Here $\mathcal{S}[a_1]$ contains none of the vertices x , y and z , so it is crossed by the geodesic from x to y , and contains a point v at bounded distance from ρ .

Proof of point 1 of Theorem 4. Let $(a_i)_{1 \leq i \leq 4}$ be four points of \mathbf{T} , neither of which is an ancestor of another. The slices $\mathcal{S}[a_i]$ are disjoint and satisfy the assumptions of Proposition 2. For every $1 \leq i \leq 4$, let K_i be given by Proposition 2 for $\mathcal{S}[a_i]$. Now consider three vertices x, y, z of \mathcal{M} and three geodesics γ_{xy} (resp. γ_{yz}, γ_{zx}) from x to y (resp. y to z, z to x) that surround ρ . There is an index $1 \leq i \leq 4$ such that $\mathcal{S}[a_i]$ contains none of the points x, y and z . Assume it is $\mathcal{S}[a_1]$. Since the triangle formed by γ_{xy}, γ_{yz} and γ_{zx} surrounds ρ , one of these three geodesics must either intersect the path in \mathbf{T} from ρ to a_1 , or cross the slice $\mathcal{S}[a_1]$, as on Figure 8. We assume this geodesic is γ_{xy} . In the first case, γ_{xy} contains a point at distance at most $h(a_1)$ from ρ . In the second case, assume γ_{xy} crosses $\mathcal{S}[a_1]$ from left to right. Let γ_ℓ and γ_r be respectively the left and right boundaries of $\mathcal{S}[a_1]$. Let v be the last point of γ_{xy} that lies on γ_ℓ and let w be the first point of γ_{xy} after v that lies on γ_r . Then the portion of γ_{xy} between v and w is a geodesic in \mathcal{M} so it is also a geodesic in $\mathcal{S}[a_1]$ that crosses $\mathcal{S}[a_1]$. Hence, it contains a point z such that $d(z, a_1) \leq K_1$. This concludes the proof by taking

$$K = \max_{1 \leq i \leq 4} (K_i + h(a_i)).$$

□

Proof of point 2 of Theorem 4. Let $a_1, a_2 \in \mathbf{T}$, neither of which is an ancestor of the other, so that $\mathcal{S}[a_1]$ and $\mathcal{S}[a_2]$ are disjoint. Let γ_ℓ and γ_r be the left and right boundaries of $\mathcal{S}[a_1]$. The idea of our construction is the following: we first "approximate" the paths γ_ℓ and γ_r by two infinite geodesics $\tilde{\gamma}_\ell$ and $\tilde{\gamma}_r$, and we then try to connect $\tilde{\gamma}_\ell$ to $\tilde{\gamma}_r$ in the shortest possible way. Note that the first step is needed because γ_r and γ_ℓ may not be geodesics in the general case. Before making this construction explicit, we need to reinforce slightly Proposition 2.

By Proposition 2, we know that any geodesic γ in $\mathcal{S}[a_1]$ between a point of γ_ℓ and a point of γ_r contains a vertex at bounded distance from ρ . We claim that this is also the case if we consider geodesics in \mathcal{M} instead of $\mathcal{S}[a_1]$. Indeed, let K_1 (resp. K_2) be given by Proposition 2 for $\mathcal{S}[a_1]$ (resp. $\mathcal{S}[a_2]$). Let $i, j \geq 0$ and let γ be a geodesic from $\gamma_\ell(i)$ to $\gamma_r(j)$ in \mathcal{M} . We are in one of the three following cases (cf. Figure 9) :

- (i) γ intersects the path in \mathbf{T} from ρ to a_1 or from ρ to a_2 ,
- (ii) γ crosses $\mathcal{S}[a_1]$,

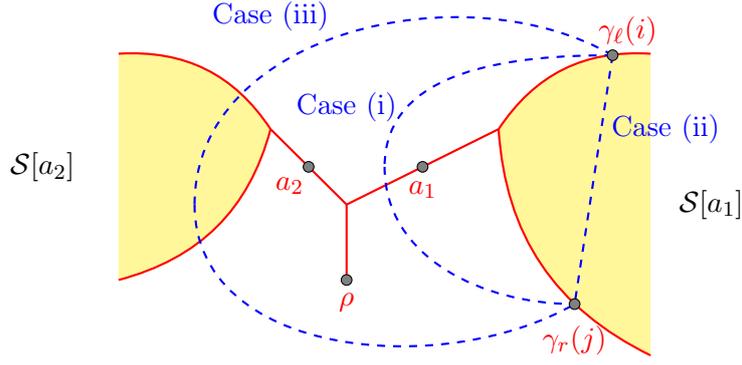


Figure 9: Reinforcement of Proposition 2. In blue, the geodesic γ . It must intersect a geodesic from ρ to a_1 or a_2 , or cross $\mathcal{S}[a_1]$ or $\mathcal{S}[a_2]$.

(iii) γ crosses $\mathcal{S}[a_2]$.

In all three cases, γ contains a point at distance at most K from ρ , where

$$K = \max(h(a_1) + K_1, h(a_2) + K_2). \quad (2)$$

We can now build our infinite geodesics $\tilde{\gamma}_\ell$ and $\tilde{\gamma}_r$. For every n , let γ_ℓ^n be a geodesic from ρ to $\gamma_\ell(n)$. By an easy compactness argument, there is an infinite geodesic $\tilde{\gamma}_\ell$ such that, for every $i \geq 0$, there are infinitely many n such that $\tilde{\gamma}_\ell(i) = \gamma_\ell^n(i)$. We build an infinite geodesic $\tilde{\gamma}_r$ from γ_r in a similar way.

For every $i, j \geq 0$, we define

$$a_{i,j} = i + j - d_{\mathcal{M}}(\tilde{\gamma}_\ell(i), \tilde{\gamma}_r(j)).$$

This quantity measures "by how much" the concatenation of $\tilde{\gamma}_\ell$ and $\tilde{\gamma}_r$ between $\tilde{\gamma}_\ell(i)$ and $\tilde{\gamma}_r(j)$ is not a geodesic. We note that for any $i, j \geq 0$, we have

$$d_{\mathcal{M}}(\tilde{\gamma}_\ell(i+1), \tilde{\gamma}_r(j)) \leq d_{\mathcal{M}}(\tilde{\gamma}_\ell(i+1), \tilde{\gamma}_\ell(i)) + d_{\mathcal{M}}(\tilde{\gamma}_\ell(i), \tilde{\gamma}_r(j)) = 1 + d_{\mathcal{M}}(\tilde{\gamma}_\ell(i), \tilde{\gamma}_r(j)),$$

so $a_{i+1,j} \geq a_{i,j}$, so $a_{i,j}$ is nondecreasing in i . Similarly, it is nondecreasing in j . We claim the following.

Lemma 4. Almost surely, $(a_{i,j})_{i,j \geq 0}$ is bounded.

Proof of Lemma 4. Let $i, j \geq 0$. By the definition of $\tilde{\gamma}_\ell$ and $\tilde{\gamma}_r$, there are two indices m and n such that $\tilde{\gamma}_\ell(i)$ lies on a geodesic from ρ to $\gamma_\ell(m)$ and $\tilde{\gamma}_r(j)$ lies on a geodesic from ρ to $\gamma_r(n)$. Therefore, we have

$$\begin{aligned} d_{\mathcal{M}}(\gamma_\ell(m), \gamma_r(n)) &\leq d_{\mathcal{M}}(\gamma_\ell(m), \tilde{\gamma}_\ell(i)) + d_{\mathcal{M}}(\tilde{\gamma}_\ell(i), \tilde{\gamma}_r(j)) + d_{\mathcal{M}}(\tilde{\gamma}_r(j), \gamma_r(n)) \\ &= d_{\mathcal{M}}(\rho, \gamma_\ell(m)) - i + d_{\mathcal{M}}(\tilde{\gamma}_\ell(i), \tilde{\gamma}_r(j)) + d_{\mathcal{M}}(\rho, \gamma_r(n)) - j \\ &= d_{\mathcal{M}}(\rho, \gamma_\ell(m)) + d_{\mathcal{M}}(\rho, \gamma_r(n)) - a_{i,j}. \end{aligned}$$

On the other hand, we know that for any $m, n \geq 0$, any geodesic from $\gamma_\ell(m)$ to $\gamma_r(n)$ contains a vertex v_0 at distance at most K from ρ , where K is given by (2). Therefore, we have

$$\begin{aligned} d_{\mathcal{M}}(\gamma_\ell(m), \gamma_r(n)) &= d_{\mathcal{M}}(\gamma_\ell(m), v_0) + d_{\mathcal{M}}(v_0, \gamma_r(n)) \\ &\geq d_{\mathcal{M}}(\gamma_\ell(m), \rho) + d_{\mathcal{M}}(\gamma_r(n), \rho) - 2d_{\mathcal{M}}(\rho, v_0) \\ &\geq d_{\mathcal{M}}(\gamma_\ell(m), \rho) + d_{\mathcal{M}}(\gamma_r(n), \rho) - 2K. \end{aligned}$$

By combining the last two equations, we obtain $a_{i,j} \leq 2K$ for every $i, j \geq 0$. \square

The construction of a bi-infinite geodesic is now easy. Let i_0, j_0 be two indices such that $a_{i_0, j_0} = \sup\{a_{i,j} \mid i, j \geq 0\}$, and let $d = d_{\mathcal{M}}(\tilde{\gamma}_\ell(i_0), \tilde{\gamma}_r(j_0))$. Let also $\hat{\gamma}$ be a geodesic from $\tilde{\gamma}_\ell(i_0)$ to $\tilde{\gamma}_r(j_0)$ in \mathcal{M} . We define a bi-infinite path γ as follows :

$$\gamma(i) = \begin{cases} \tilde{\gamma}_\ell(i_0 - i) & \text{if } i \leq 0, \\ \hat{\gamma}(i) & \text{if } 0 \leq i \leq d, \\ \tilde{\gamma}_r(i - d + j_0) & \text{if } i \geq d. \end{cases}$$

We finally check that this is indeed a bi-infinite geodesic. Let $i \geq i_0$ and $j \geq j_0$. We have $a_{i,j} = a_{i_0, j_0}$, so

$$d_{\mathcal{M}}(\tilde{\gamma}_\ell(i), \tilde{\gamma}_r(j)) = d + (i - i_0) + (j - j_0),$$

so $d(\gamma(i_0 - i), \gamma(d + j - j_0)) = (d + j - j_0) - (i_0 - i)$. Therefore, we have $d(\gamma(i'), \gamma(j')) = j' - i'$ for $i' \leq 0$ small enough and $j' \geq 0$ large enough, so γ is a bi-infinite geodesic. \square

3 Poisson boundary

3.1 General setting

The goal of this subsection is to build a compactification of maps that, as we will later prove, is under some assumptions a realization of their Poisson boundary. We will perform this construction directly in the general framework $\mathcal{M} = \mathcal{M}(\mathbf{T}, (s_i))$. This construction is exactly the same as the construction performed for the PSHIT in Section 3.1 of [12].

We recall that $\partial\mathbf{T}$ is the set of infinite rays from ρ in \mathbf{T} . If $\gamma, \gamma' \in \partial\mathbf{T}$, we write $\gamma \sim \gamma'$ if $\gamma = \gamma'$ or if γ and γ' are "consecutive" in the sense that there is no ray between them (in particular, if γ_ℓ and γ_r are the leftmost and rightmost rays of \mathbf{T} , then $\gamma_\ell \sim \gamma_r$). It is equivalent to saying that γ and γ' are the left and right boundaries of some strip s_i in the map \mathcal{M} . Note that a.s., every ray of \mathbf{T} contains infinitely many branching points, so no ray is equivalent to two distinct other rays. It follows that \sim is a.s. an equivalence relation for which countably many equivalence classes have cardinal 2, and all the others have cardinal 1. We write $\hat{\partial}\mathbf{T} = \partial\mathbf{T} / \sim$ and we denote by $\gamma \rightarrow \hat{\gamma}$ the canonical projection from $\partial\mathbf{T}$ to $\hat{\partial}\mathbf{T}$. Finally, for every strip s_i , the left and right boundaries of s_i correspond to the same point of $\hat{\partial}\mathbf{T}$, that we denote by $\hat{\gamma}_i$.

Our goal is now to define a topology on $\mathcal{M} \cup \hat{\partial}\mathbf{T}$. It should be possible to define it by an explicit distance, but such a distance would be tedious to write down, so we prefer to give an "abstract" construction. Let s_i and s_j be two distinct strips of \mathcal{M} , and fix $h > 0$ such that both s_i and s_j both intersect $B_h(\mathbf{T})$. Then $\mathcal{M} \setminus (B_h(\mathcal{M}) \cup s_i \cup s_j)$ has two infinite connected

components, that we denote by (s_i, s_j) and (s_j, s_i) (the vertices on the boundaries of s_i and s_j do not belong to (s_i, s_j) and (s_j, s_i)). We also write

$$\widehat{\partial}(s_i, s_j) = \{\widehat{\gamma} \mid \gamma \text{ is a ray of } \mathbf{T} \text{ such that } \gamma(k) \in (s_i, s_j) \text{ for } k \text{ large enough}\}.$$

We define $\widehat{\partial}(s_j, s_i)$ similarly. Note that $\widehat{\partial}(s_i, s_j)$ and $\widehat{\partial}(s_j, s_i)$ are disjoint subsets of $\widehat{\partial}\mathbf{T}$, and their union is $\widehat{\partial}\mathbf{T} \setminus \{\widehat{\gamma}_i, \widehat{\gamma}_j\}$.

We can now equip the set $\mathcal{M} \cup \widehat{\partial}\mathbf{T}$ with the topology generated by the following open sets:

- the singletons $\{v\}$, where v is a vertex of \mathcal{M} ,
- the sets $(s_i, s_j) \cup \widehat{\partial}(s_i, s_j)$, where s_i and s_j are two distinct strips of \mathcal{M} .

This topology is separated (if $\widehat{\gamma}_1 \neq \widehat{\gamma}_2$, then there are two strips separating γ_1 and γ_2) and has a countable basis, so it is induced by a distance. Moreover, any open set of our basis intersects \mathcal{M} , so \mathcal{M} is dense in $\mathcal{M} \cup \widehat{\partial}\mathbf{T}$. Finally, we state an intuitive result about the topology of $\mathcal{M} \cup \widehat{\partial}\mathbf{T}$. Its proof in the particular case of the PSHIT can be found in [12], and adapts without any change to the general case.

Lemma 5. The space $\mathcal{M} \cup \widehat{\partial}\mathbf{T}$ is compact, and $\widehat{\partial}\mathbf{T}$ is homeomorphic to the unit circle.

3.2 Transience away from the boundary in causal slices

The goal of this section is to prove Proposition 6, which is the main tool in the proof of Theorem 2. We recall that \mathcal{S} is the causal slice associated to a supercritical Galton–Watson tree T , and $\partial\mathcal{S}$ is the *boundary* of \mathcal{S} , i.e. the set of vertices of \mathcal{S} that are either the leftmost or the rightmost vertex of their generation. We also write $\tau_{\partial\mathcal{S}} = \min\{n \geq 0 \mid X_n \in \partial\mathcal{S}\}$, where (X_n) is the simple random walk on \mathcal{S} .

Proposition 6. Almost surely, there is a vertex $x \in \mathcal{S}$ such that

$$P_{\mathcal{S},x}(\tau_{\partial\mathcal{S}} = +\infty) > 0.$$

Note that if such a vertex x exists, then we have $P_{\mathcal{S},v}(X_n \notin \partial\mathcal{S} \text{ for } n \text{ large enough}) > 0$ for every vertex $v \in \mathcal{S}$. The proof of Proposition 6 is based on estimates of effective resistances. We will use the following inequality, that holds for every graph and every vertex x :

$$P_{\mathcal{S},x}(\tau_{\partial\mathcal{S}} < +\infty) \leq \frac{R_{\text{eff}}^{\mathcal{S}}(x \leftrightarrow \{\partial\mathcal{S}, \infty\})}{R_{\text{eff}}^{\mathcal{S}}(x \leftrightarrow \partial\mathcal{S})} \leq \frac{R_{\text{eff}}^{\mathcal{S}}(x \leftrightarrow \infty)}{R_{\text{eff}}^{\mathcal{S}}(x \leftrightarrow \partial\mathcal{S})}. \quad (3)$$

For example, this is a particular case of Exercise 2.36 of [26]. We will find a sequence (x_n) of vertices satisfying the following two properties:

1. we have $R_{\text{eff}}^{\mathcal{S}}(x_n \leftrightarrow \partial\mathcal{S}) \rightarrow +\infty$ a.s. when $n \rightarrow +\infty$,
2. for every $n \geq 0$, the resistance $R_{\text{eff}}^{\mathcal{S}}(x_n \leftrightarrow \infty)$ is stochastically dominated by $R_{\text{eff}}^{\mathcal{S}}(\rho \leftrightarrow \infty)$. In particular, a.s., $(R_{\text{eff}}^{\mathcal{S}}(x_n \leftrightarrow \infty))$ has a bounded subsequence.

By (3), this will guarantee that

$$P_{\mathcal{S},x_n}(\tau_{\partial\mathcal{S}} < +\infty) \xrightarrow[n \rightarrow +\infty]{} 0$$

along some subsequence, which is enough to prove Proposition 6.

We choose for the sequence (x_n) the nonbacktracking random walk on the backbone of T . More precisely, we take $x_0 = \rho$ and, for every $n \geq 0$, conditionally on \mathcal{S} and x_0, \dots, x_n , the vertex x_{n+1} is chosen uniformly among the children of x_n in $\mathbf{B}(T)$. We can give a "spinal decomposition" of $\mathbf{B}(T)$ along (x_n) . We recall that $\boldsymbol{\mu}$ is the offspring distribution of $\mathbf{B}(T)$, cf. (1). For every $n \geq 0$, let L_n (resp. R_n) be the number of children of x_n in $\mathbf{B}(T)$ on the left (resp. on the right) of x_{n+1} . A vertex v of $\mathbf{B}(T)$ will be called a *spine brother* if the parent of v is equal to x_n for some n but $v \neq x_{n+1}$. Then the pairs (L_n, R_n) are i.i.d. with distribution ν given by

$$\mathbb{P}(L_n = \ell, R_n = r) = \nu(\{(\ell, r)\}) = \frac{1}{r + \ell + 1} \boldsymbol{\mu}(r + \ell + 1). \quad (4)$$

Moreover, conditionally on (L_n) and (R_n) , the backbones of the trees of descendants of the spine brothers are i.i.d. Galton–Watson trees with offspring distribution $\boldsymbol{\mu}$. The distribution of T conditionally on this backbone is then given by Theorem 4. In particular, for every $n \geq 0$, the tree of descendants of x_n has the same distribution as T , so $\mathcal{S}[x_n]$ has the same distribution as \mathcal{S} .

Therefore, for every n , we have

$$R_{\text{eff}}^{\mathcal{S}}(x_n \leftrightarrow \infty) \leq R_{\text{eff}}^{\mathcal{S}[x_n]}(x_n \leftrightarrow \infty),$$

where $R_{\text{eff}}^{\mathcal{S}[x_n]}(x_n \leftrightarrow \infty)$ has the same distribution as $R_{\text{eff}}^{\mathcal{S}}(\rho \leftrightarrow \infty)$. This proves the second property that we wanted (x_n) to satisfy. Hence, it only remains to prove that $R_{\text{eff}}^{\mathcal{S}}(x_n \leftrightarrow \partial\mathcal{S})$ almost surely goes to $+\infty$, which is the goal of the next lemma.

Lemma 7. There are disjoint vertex sets $(A_k)_{k \geq 0}$, satisfying the following properties:

- (i) for any $k \geq 1$, the set A_k separates $\partial\mathcal{S}$ from all the sets A_i with $i > k$, and from x_n for n large enough,
- (ii) the parts of \mathcal{S} lying between A_{2k} and A_{2k+1} for $k \geq 0$ are i.i.d.,
- (iii) we have $R_{\text{eff}}^{\mathcal{S}}(A_{2k} \leftrightarrow A_{2k+1}) > 0$ a.s. for every $k \geq 0$.

Proof of Proposition 6 given Lemma 7. Fix k and choose n large enough, so that x_n is separated from $\partial\mathcal{S}$ by $A_0, A_1, \dots, A_{2k-1}$. Since A_0, \dots, A_{2k-1} are disjoint cutsets separating x_n from $\partial\mathcal{S}$, we have

$$\begin{aligned} R_{\text{eff}}^{\mathcal{S}}(x_n \leftrightarrow \partial\mathcal{S}) &\geq R_{\text{eff}}^{\mathcal{S}}(\partial\mathcal{S} \leftrightarrow A_0) + R_{\text{eff}}^{\mathcal{S}}(x_n \leftrightarrow A_{2k-1}) + \sum_{i=0}^{2k-2} R_{\text{eff}}^{\mathcal{S}}(A_i \leftrightarrow A_{i+1}) \\ &\geq \sum_{i=0}^{k-1} R_{\text{eff}}^{\mathcal{S}}(A_{2i} \leftrightarrow A_{2i+1}). \end{aligned}$$

Since the variables $R_{\text{eff}}^{\mathcal{S}}(A_{2i} \leftrightarrow A_{2i+1})$ for $0 \leq i \leq k-1$ are i.i.d. and a.s. positive, this goes to $+\infty$ as $k \rightarrow +\infty$, so we have $R_{\text{eff}}^{\mathcal{S}}(x_n \leftrightarrow \partial\mathcal{S}) \rightarrow +\infty$ a.s. when $n \rightarrow +\infty$, which ends the proof. \square

Remark 8. The law of large numbers even shows that there is a constant $c > 0$ such that, for n large enough, we have $R_{\text{eff}}^{\mathcal{S}}(x_n \leftrightarrow \partial\mathcal{S}) \geq cn$. This is quite similar to the resistance estimates

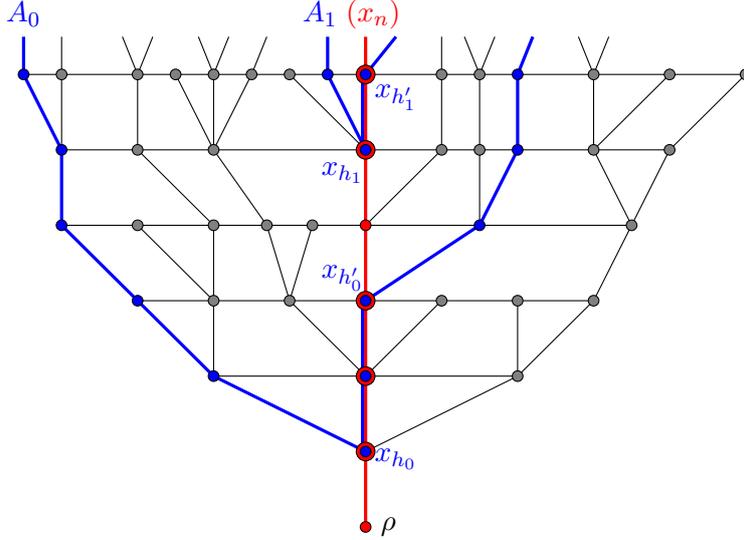


Figure 10: The slice \mathcal{S} with the sequences of vertices (x_n) (in red), and the separating sets A_k (in blue). Here we have $h_0 = 1$, $h'_0 = 3$, $h_1 = 5$ and $h'_1 = 6$.

proved in [7] in the case where T is the complete binary tree, and it might be interesting to apply our results to the study of the Gaussian free field on causal maps. Unfortunately, our estimates only hold in "typical" directions, and not uniformly for all the vertices. We also note that the idea to "cut" \mathcal{S} along the sets A_k was inspired by the proof of Lemma 1 of [7].

We now build the subsets A_k . We define by induction the heights h_k and h'_k for $k \geq 0$ by

$$\begin{aligned} h_0 &= \min\{n \geq 0 \mid L_n > 0\}, \\ h'_k &= \min\{n > h_k \mid R_n > 0\}, \\ h_{k+1} &= \min\{n > h'_k \mid L_n > 0\}. \end{aligned}$$

Note that the pairs (L_n, R_n) are i.i.d. with $\mathbb{P}(L_n > 0) > 0$ and $\mathbb{P}(R_n > 0) > 0$, so h_k and h'_k are a.s. well-defined for every k . We define A_k as the union of the leftmost ray of $\mathbf{B}(T)$ from x_{h_k} , the rightmost ray of $\mathbf{B}(T)$ from $x_{h'_k}$ and the vertices x_i with $h_k \leq i \leq h'_k$ (see Figure 10).

It is easy to see that the sets A_k are disjoint and that A_k separates $\partial\mathcal{S}$ from A_i for every $i > k$, and from x_n for n large enough, so they satisfy property (i) of Lemma 7.

For every $k \geq 0$, let \mathcal{U}_k be the sub-map of \mathcal{S} whose vertices are the vertices between A_k and A_{k+1} (the vertices of A_k and A_{k+1} are included), rooted at x_{h_k} . By using (4) and the backbone decomposition, it is quite straightforward to prove that the maps \mathcal{U}_{2k} for $k \geq 0$ are i.i.d.. We do not give a precise description of the distribution of \mathcal{U}_k , the only property of \mathcal{U}_k that we will need later is that it contains a copy of T on each side of the spine.

Remark 9. It is still true that the \mathcal{U}_k for $k \geq 0$ are identically distributed. However, \mathcal{U}_k and \mathcal{U}_{k+1} are not independent. Indeed, for $h_{k+1} \leq n < h'_{k+1}$, the variables L_n and R_n are not independent, and L_n affects \mathcal{U}_{k+1} , whereas R_n affects \mathcal{U}_k . This is why we restrict ourselves to even values of k .

It only remains to prove property (iii) in Lemma 7. Let \mathcal{U} be distributed as the \mathcal{U}_k , and let A_b (resp. A_t) be its bottom (resp. top) boundary, playing the same role as A_k (resp. A_{k+1}). The proof that $R_{\text{eff}}^{\mathcal{U}}(A_b \leftrightarrow A_t) > 0$ relies on a duality argument. We first recall a classical result about duality of resistances in planar maps. Let M be a finite planar map drawn in the plane, and let a and z be two vertices adjacent to its outer face. We draw two infinite half-lines from a and z that split the outer face in two faces a^* and z^* . Now let M^* be the dual planar map whose vertices are a^* , z^* and the internal faces of M . Then we have

$$R_{\text{eff}}^M(a \leftrightarrow z) = \left(R_{\text{eff}}^{M^*}(a^* \leftrightarrow z^*) \right)^{-1}. \quad (5)$$

In our case, the infinite graph \mathcal{U} has two ends ∞_ℓ (on the left of the spine) and ∞_r (on the right). Informally, we would like to write

$$R_{\text{eff}}^{\mathcal{U}}(A_b \leftrightarrow A_t) = \left(R_{\text{eff}}^{\mathcal{U}^*}(\infty_\ell^* \leftrightarrow \infty_r^*) \right)^{-1}, \quad (6)$$

which would reduce the problem to the proof of $R_{\text{eff}}^{\mathcal{U}^*}(\infty_\ell^* \leftrightarrow \infty_r^*) < +\infty$. Our first job will be to state and prove (a proper version of) (6).

More precisely, we denote by \mathcal{U}_ℓ (resp. \mathcal{U}_r) the part of \mathcal{U} lying on the left (resp. on the right) of the spine. We also define \mathcal{U}^* as the dual map of \mathcal{U} in the following sense: the vertices of \mathcal{U}^* are the finite faces of \mathcal{U} , and for every edge of \mathcal{U} that does not link two vertices of A_b or two vertices of A_t , we draw an edge e^* between the two faces adjacent to e . Let θ^* be a flow on \mathcal{U}^* with no source. We assume that θ^* is unitary in the sense that the mass of θ^* crossing the spine from left to right is equal to 1. We recall that the *energy* $\mathcal{E}(\theta^*)$ of θ^* is the sum over all edges e^* of \mathcal{U}^* of $\theta^*(e^*)^2$. For every $n \geq 0$, let $A_b(n)$ (resp. $A_t(n)$) be the set of vertices of A_b (resp. A_t) at height at most n . We consider the map $\mathcal{U}(n)$ obtained by cutting \mathcal{U} above height n . The restriction of θ to the dual of this map becomes a unitary flow crossing $\mathcal{U}(n)^*$ from left to right. Therefore, the dual resistance from left to right in $\mathcal{U}(n)^*$ is at most $\mathcal{E}(\theta^*)$ so, by (5), we obtain

$$R_{\text{eff}}^{\mathcal{U}}(A_b(n) \leftrightarrow A_t(n)) \geq \mathcal{E}(\theta^*)^{-1}$$

and, by letting $n \rightarrow +\infty$, we get $R_{\text{eff}}^{\mathcal{U}}(A_b \leftrightarrow A_t) \geq \mathcal{E}(\theta^*)^{-1}$. In particular, if there is such a flow θ^* with finite energy, then $R_{\text{eff}}^{\mathcal{U}}(A_b \leftrightarrow A_t) > 0$ and Lemma 7 is proved.

We now define \mathcal{U}_ℓ (resp. \mathcal{U}_r) as the part of \mathcal{U} lying on the left (resp. on the right) of the spine. Let f_ℓ and f_r be two faces of \mathcal{U} lying respectively on the left and on the right of the same edge of the spine. A simple way to construct a unitary flow θ^* with no sources is to concatenate a flow θ_ℓ^* from infinity to f_ℓ in \mathcal{U}_ℓ^* , a flow of mass 1 in the dual edge from f_ℓ to f_r and a flow θ_r^* from f_r to infinity in \mathcal{U}_r^* . For this flow to have finite energy, we need θ_ℓ^* and θ_r^* to have finite energy, so we need both \mathcal{U}_ℓ^* and \mathcal{U}_r^* to be transient.

We now define \mathcal{S}^* as the dual of the slice \mathcal{S} (as above, the vertices of \mathcal{S}^* are the inner faces of \mathcal{S}). We note that, for every vertex $v_0 \in \mathcal{U}_\ell \cap \mathbf{B}(T)$ that does not belong to the spine, the tree of descendants of v_0 has the same distribution as T and is entirely contained in \mathcal{U}_ℓ . Therefore, \mathcal{U}_ℓ contains a copy of \mathcal{S} , so \mathcal{U}_ℓ^* contains a copy of \mathcal{S}^* , and the same is true for \mathcal{U}_r^* . Hence, we have reduced the proof of Proposition 6 to the next result.

Lemma 10. The dual slice \mathcal{S}^* is a.s. transient.

Proof. We will show that we can embed a transient tree in \mathcal{S}^* . The idea will be to follow the branches of the tree T in the dual, to obtain a tree T^* that is similar to T . However, vertices

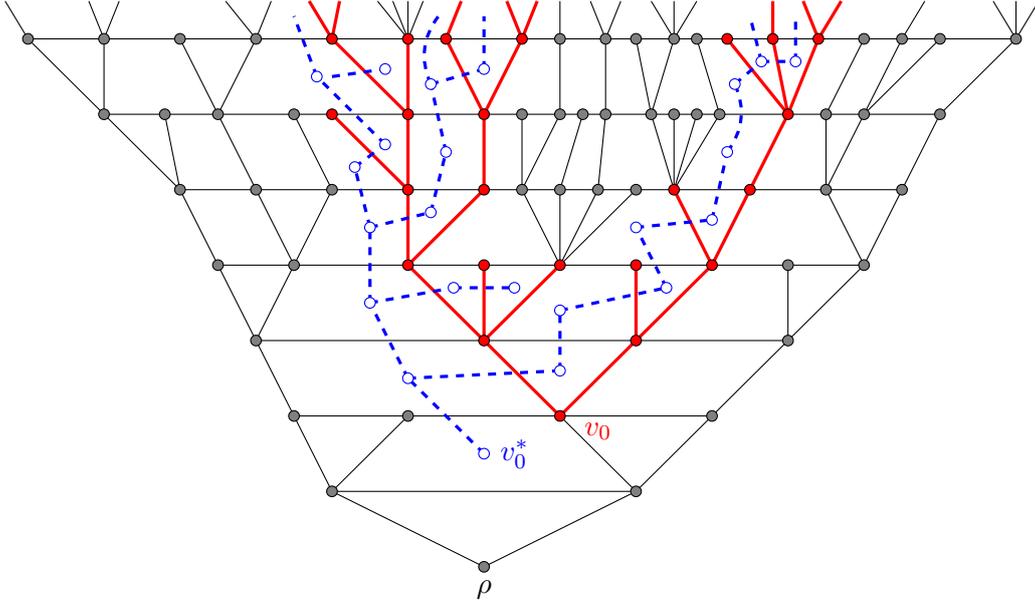


Figure 11: The construction of the dual tree T^* (in blue) from the tree $T[v_0]_{\text{bdd}}$ (in red). Here, we have taken $c_{\max} = 3$. The vertices of T^* are the faces adjacent to the vertices of $T[v_0]_{\text{bdd}}$ at their bottom-left corners.

of high degree become obstacles: if a vertex v of T has degree d in T , we need d dual edges to "move around" v in \mathcal{S}^* . Therefore, it becomes difficult to control the ratio between resistances in T^* and in T . To circumvent this problem, we will use the fact that T contains a supercritical Galton–Watson tree with bounded degrees.

More precisely, we fix a constant c_{\max} large enough to have

$$\sum_{i=0}^{c_{\max}} i\mu(i) > 1.$$

If t is a (finite or infinite) tree, for every vertex v with more than c_{\max} children, we remove all the edges between v and its children, and we call t_{bdd} the connected component of the root. If T' is a Galton–Watson tree with distribution μ , then T'_{bdd} is a Galton–Watson tree with offspring distribution μ_{bdd} given by

$$\mu_{\text{bdd}}(i) = \begin{cases} 0 & \text{if } i > c_{\max}, \\ \mu(i) & \text{if } 0 < i \leq c_{\max}, \\ \mu(0) + \sum_{j>c_{\max}} \mu(j) & \text{if } i = 0. \end{cases}$$

In particular, we have $\sum_i i\mu_{\text{bdd}}(i) = \sum_{i=0}^{c_{\max}} i\mu(i) > 1$, so T'_{bdd} is supercritical, and it survives with positive probability. But T is a Galton–Watson tree conditioned to survive, so it contains infinitely many i.i.d. copies of T (take for examples the trees of descendants of the children of the right boundary). Therefore, there is a.s. a vertex $v_0 \in T$ that is not on the left boundary of \mathcal{S} , such that $T[v_0]_{\text{bdd}}$ is a Galton–Watson tree and survives. In particular, it is transient and, for every $v \in T[v_0]_{\text{bdd}}$, the number of children of v in T is bounded by c_{\max} .

From here, we can build a tree T^* in \mathcal{S}^* whose branches follow the branches of $T[v_0]_{\text{bdd}}$ on their left, and which circumvents branching points of $T[v_0]_{\text{bdd}}$ by the top. See Figure 11 for the

construction of this tree. The tree T^* is then a subgraph of \mathcal{S}^* . Therefore, it is enough to prove that T^* is transient. But since the vertex degrees in $T[v_0]_{\text{bdd}}$ are bounded, it is easy to see that $T[v_0]_{\text{bdd}}$ and T^* are quasi-isometric, so T^* is also transient (by e.g. Section 2.4.4 of [17]). \square

3.3 Consequences on the Poisson boundary

We recall that (X_n) is the simple random walk on \mathcal{C} started from ρ . By a result of Hutchcroft and Peres (Theorem 1.3 of [19]), the first point of Theorem 2 implies the second.

Proof of Theorem 2. We first show the almost sure convergence of (X_n) . By compactness (Lemma 5), it is enough to prove that (X_n) a.s. has a unique subsequential limit in $\widehat{\partial T}$. Note that if $\widehat{\gamma}_1 \neq \widehat{\gamma}_2$ are two distinct points of $\widehat{\partial T}$, then there are two vertices $v, v' \in \mathbf{B}(T)$ such that the slices $\mathcal{S}[v]$ and $\mathcal{S}[v']$ separate γ_1 from γ_2 . Therefore, if γ_1 and γ_2 are subsequential limits of (X_n) , by transience of $\mathcal{C}(T)$, the walk (X_n) crosses infinitely many times either $\mathcal{S}[v]$ or $\mathcal{S}[v']$ horizontally. Therefore, it is enough to prove that for every $v_0 \in \mathbf{B}(T)$, the walk (X_n) cannot cross infinitely many times $\mathcal{S}[v_0]$ horizontally.

For every $v \in \mathcal{S}[v_0]$, let $f(v) = P_{\mathcal{S}[v_0], v}(\tau_{\partial\mathcal{S}[v_0]} < +\infty)$. The function f is harmonic on $\mathcal{S}[v_0] \setminus \partial\mathcal{S}[v_0]$. Moreover, by Proposition 6, there is a vertex $v_1 \in \mathcal{S}[v_0]$ such that $f(v_1) < 1$. Let (Y_n) be a simple random walk started from v_1 and killed when it hits $\partial\mathcal{S}[v_0]$. Then $f(Y_n)$ is a martingale and, by the martingale convergence theorem, it converges a.s. to $\mathbb{1}_{\tau_{\partial\mathcal{S}[v_0]} < +\infty}$. In particular, it has limit zero with positive probability, so there is an infinite path (w_k) going to infinity in $\mathcal{S}[v_0]$, such that $f(w_k) \rightarrow 0$.

We fix $k_0 > 0$. Everytime the walk (X_n) crosses $\mathcal{S}[v_0]$ horizontally at a large enough height, it must cross the path $(w_k)_{k \geq 0}$. Since \mathcal{C} is transient, if X crosses $\mathcal{S}[v_0]$ infinitely many times, it must cross $(w_k)_{k \geq k_0}$, and then hit $\partial\mathcal{S}[v_0]$. If this happens, let K be such that w_K is the first of the points $(w_k)_{k \geq k_0}$ to be hit by X (if none of these points is hit, we take $K = +\infty$). We have

$$\mathbb{P}(X \text{ hits } (w_k)_{k \geq k_0} \text{ and then } \partial\mathcal{S}[v_0]) = \mathbb{E}[\mathbb{1}_{K < +\infty} f(w_K)] \leq \mathbb{E}\left[\sup_{k \geq k_0} f(w_k)\right].$$

Since $f(w_k) \rightarrow 0$, by dominated convergence, this goes to 0, which proves that X cannot cross $\mathcal{S}[v_0]$ infinitely many times. This implies the almost sure convergence of X to a point X_∞ of $\widehat{\partial T}$.

The proof that X_∞ has full support is quite easy. Let $v_0 \in \mathbf{B}(T)$. Then (X_n) has a positive probability to visit the slice $\mathcal{S}[v_0]$ and, by Proposition 6, it a.s. has a positive probability to stay there ever after. But if $X_n \in \mathcal{S}[v_0]$ for n large enough, then X_∞ must correspond to a ray of descendants of v_0 , so the distribution of X_∞ gives a positive mass to rays that are descendants of v_0 . This is almost surely true for any $v_0 \in \mathbf{B}(T)$, so the distribution of X_∞ has a.s. full support.

Finally, to prove the almost sure nonatomicity, it is enough to prove that if X and Y are two independent simple random walks on \mathcal{C} , then $X_\infty \neq Y_\infty$ almost surely. The idea of the proof is that everytime X and Y reach a new height for the first time, by Proposition 6, they have a positive probability to get "swallowed" in two different slices of the form $\mathcal{S}[x]$ and $\mathcal{S}[y]$, so this will almost surely happen at some height.

More precisely, **in this proof and in this proof only**, until the end of Section 3.3, we assume that T is a **non-conditioned** Galton-Watson tree with offspring distribution μ . We

recall that Z_h is the number of vertices of T at height h . For every $h \geq 0$, let

$$\tau_h^X = \min\{n \geq 0 | h(X_n) = h\} \quad \text{and} \quad \tau_h^Y = \min\{n \geq 0 | h(Y_n) = h\}.$$

Note that if T survives, then $\tau_h^X, \tau_h^Y < +\infty$ for every h . Let also \mathcal{F}_h be the σ -algebra generated by $B_h(\mathcal{C}(T))$, $(X_n)_{0 \leq n \leq \tau_h^X}$ and $(Y_n)_{0 \leq n \leq \tau_h^Y}$, and let \mathcal{F}_∞ be the σ -algebra generated by $\bigcup_{h \geq 0} \mathcal{F}_h$. We note right now that $(\mathcal{F}_h)_{h \geq 0}$ is nondecreasing and that \mathcal{F}_∞ is the σ -algebra generated by (\mathcal{C}, X, Y) . Finally, for every $h > 0$, let A_h be the event

- $\left\{ \begin{array}{l} \text{There are four distinct vertices } (x_i)_{1 \leq i \leq 4} \text{ of } T \text{ at height } h \text{ such that:} \\ \bullet \text{ the vertices } x_1, x_2, x_3 \text{ and } x_4 \text{ lie in this cyclic order,} \\ \bullet \text{ the trees } T[x_i] \text{ all survive,} \\ \bullet \text{ for every } n \geq \tau_h^X, \text{ we have } X_n \in T[x_1] \\ \bullet \text{ for every } n \geq \tau_h^Y + 2, \text{ we have } Y_n \in T[x_3]. \end{array} \right\}$

Lemma 11. There is a constant $\delta > 0$ such that for every h , if $Z_h \geq 4$, then

$$\mathbb{P}(A_h | \mathcal{F}_h) \geq \delta.$$

Once this lemma is known, the end of the proof is quite easy: let $A = \bigcup_{h \geq 0} A_h$. If T survives, then $Z_h \geq 4$ for h large enough, so

$$\delta \leq \mathbb{P}(A_h | \mathcal{F}_h) \leq \mathbb{P}(A | \mathcal{F}_h) \xrightarrow[h \rightarrow +\infty]{a.s.} \mathbb{P}(A | \mathcal{F}_\infty) = \mathbb{1}_A,$$

by the martingale convergence theorem and the fact that (\mathcal{C}, X, Y) is \mathcal{F}_∞ -measurable. Therefore, almost surely, if T survives, there is an h such that A_h occurs. But if it does, the slices $\mathcal{S}[x_2]$ and $\mathcal{S}[x_4]$ separate X and Y eventually, so they separate X_∞ from Y_∞ , so $X_\infty \neq Y_\infty$, which ends the proof of Theorem 2. \square

Remark 12. It is easy to show by using Proposition 6 that $\mathbb{P}(A_h | \mathcal{F}_h) > 0$ a.s.. However, this is not sufficient to prove Lemma 11. Indeed, this is the one point in our proof of Theorem 2 at which our argument fails to hold in a more general setting. More precisely, in the setting of a tree \mathbf{T} filled with i.i.d. strips, the lower degree of $X_{\tau_h^X}$ (i.e. the number of edges joining this vertex to a lower vertex) is not constant but depends on \mathcal{F}_h , and we might imagine that it goes to $+\infty$ as $h \rightarrow +\infty$. In this case, we might have $\mathbb{P}(A_h | \mathcal{F}_h) \rightarrow 0$ (with high probability, X goes back down right after τ_h^X). This problem does not occur in $\mathcal{C}(T)$, where the lower degree is always equal to 1, but it explains why the proof of Lemma 11 needs to be treated with some care.

Proof of Lemma 11. The proof will be split into three cases: the case where $X_{\tau_h^X} = Y_{\tau_h^Y}$, the case where $X_{\tau_h^X}$ and $Y_{\tau_h^Y}$ are distinct but neighbours, and the case where they are not neighbours. We treat carefully the first one, which is slightly more complicated than the others.

In the first case, we write $x_1 = X_{\tau_h^X} = Y_{\tau_h^Y}$. We also denote by x_2 and x_3 the two vertices at height h on the right of x_1 , and by x_4 the left neighbour of x_1 . Let also A'_h be the following event:

{The trees $T[x_i]$ for $1 \leq i \leq 4$ survive. Moreover, we have $X_n \in T[x_1]$ for every $n \geq \tau_h^X$ and $Y_{\tau_h^Y+1} = x_2$, $Y_{\tau_h^Y+2} = x_3$ and $Y_n \in T[x_3]$ for every $n \geq \tau_h^Y + 2$.}

If A'_h occurs, then so does A_h . Moreover, we claim that the probability for A'_h to occur is independent of h and \mathcal{F}_h . The reason why this is true is that for every vertex v in one of these trees (say $T[x_1]$), the number of neighbours of v in \mathcal{C} that are not in $\mathcal{S}[x_1]$ is fixed: there are 3 such neighbours if $v = x_1$, there is 1 such neighbour if $v \neq x_1$ is on the boundary of $\mathcal{S}[x_1]$ and 0 if it is not. Therefore, the probability given \mathcal{C} that X stays in $\mathcal{S}[x_1]$ after time τ_h^X only depends on $T[x_1]$. Similarly, the probability for Y to perform the right first two steps after time τ_h^Y only depends on the numbers of children of x_1 and x_2 , and the probability to stay in $T[x_3]$ ever after only depends on $T[x_3]$. All this is independent of h and \mathcal{F}_h , which proves our claim. If we write $\delta_1 = \mathbb{P}(A'_h|\mathcal{F}_h)$, we have $\delta_1 > 0$ by Proposition 6 and $\mathbb{P}(A_h|\mathcal{F}_h) \geq \delta_1$ for every h in this first case.

The other two cases can be treated similarly with minor adaptations in the choices of the vertices x_i , and the first two steps of Y after τ_h^Y . While we needed to control exactly the first two steps of Y in the first case, we only need one step for the second case and zero step for the third one. The other two cases yield two constants δ_2 and δ_3 , which proves the lemma by taking $\delta = \min(\delta_1, \delta_2, \delta_3)$. □

3.4 Robustness of Proposition 6

The goal of this subsection is to explain why Proposition 6 still holds in a quite general setting and to deduce the following result. We recall that a graph G is *Liouville* if its Poisson boundary is trivial, i.e. if every bounded harmonic function on G is constant.

Theorem 4. Let \mathbf{T} be a supercritical Galton–Watson tree with offspring distribution μ such that $\mu(0) = 0$, and let $(S_i)_{i \geq 0}$ be an i.i.d. sequence of random strips. We assume that (S_i) is also independent from \mathbf{T} .

1. Proposition 6 holds if we replace \mathcal{S} by $\mathcal{S}(\mathbf{T}, (S_i)_{i \geq 0})$.
2. The map $\mathcal{M}(\mathbf{T}, (S_i))$ is a.s. non-Liouville.
3. If furthermore the strips (S_i) are a.s. recurrent and bounded-degree, then $\widehat{\partial}\mathbf{T}$ is the Poisson boundary of $\mathcal{M}(\mathbf{T}, (S_i))$.

Note that the assumption that the S_i are recurrent is necessary. For example, if some strips S_i have a non-trivial Poisson boundary, then the Poisson boundary of $\mathcal{M}(\mathbf{T}, (S_i))$ is larger than $\widehat{\partial}\mathbf{T}$. See Section 5 for a more developed discussion.

Proof of the first point. Most of the proof works exactly along the same lines as the proof of Proposition 6, with \mathbf{T} playing the same role as the backbone tree. In particular, we choose for (x_n) a nonbacktracking random walk on \mathbf{T} , and the sets A_k are built in the same way as in the original proof, but from \mathbf{T} instead of $\mathbf{B}(T)$. The proof of Proposition 6 from Lemma 7 is very similar, as well as points (i) and (ii) of Lemma 7. The only difference is that the proof of point (ii) of Lemma 7 is easier in our new framework, because of the independence of the strips, and we do not need anymore to restrict ourselves to even values of k .

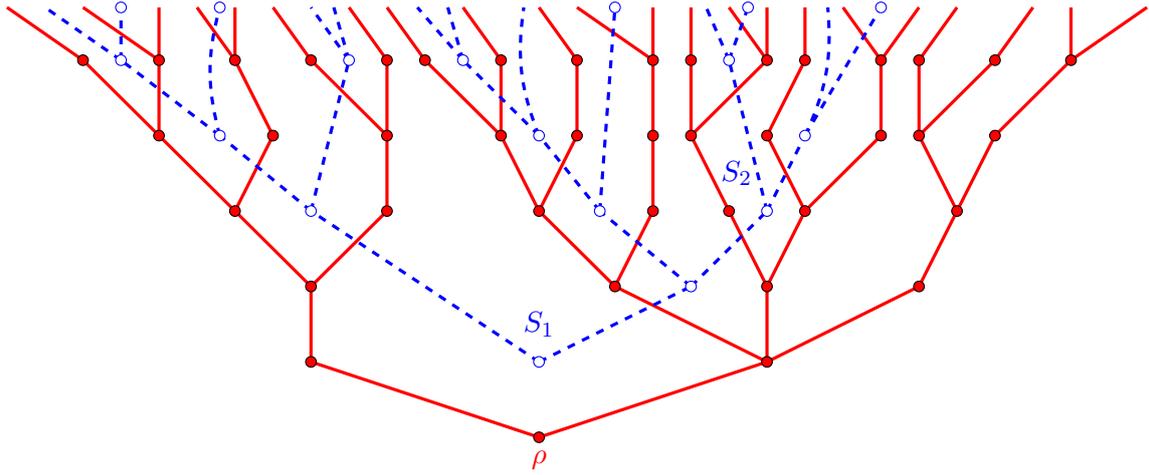


Figure 12: The tree \mathbf{T} (in red), and the tree T_{str} of descendants of the strip S_1 (in blue). Note that the strip S_2 has one parent on each side, but the right parent is not a descendant of S_1 .

Exactly as in the first proof, by using the "self-similarity" property of $\mathcal{S}(\mathbf{T}, (S_i)_{i \geq 0})$, the proof of point (iii) of Lemma 7 can be reduced to the proof that the dual map of $\mathcal{S}(\mathbf{T}, (S_i)_{i \geq 0})$ is transient (it is also important that the sets A_k do not touch each other, which is why we have required that the boundaries of the strips are simple). The adaptation of the proof of Lemma 10 (transience of the dual slice), however, is not obvious.

More precisely, let \mathcal{S}^* be the graph whose vertices are the finite faces of $\mathcal{S}(\mathbf{T}, (S_i)_{i \geq 0})$ and where, for every edge e of $\mathcal{S}(\mathbf{T}, (S_i)_{i \geq 0})$ that is adjacent to two finite faces, we draw an edge e^* between these two faces. We note right now that, since all the strips have only finite faces, the graph \mathcal{S}^* is a connected graph, and we need to prove that it is transient. The idea of the proof is the following: we will build a genealogy on the set of strips, which contains a complete binary tree. As in the proof of Lemma 10, we will then kill the strips whose root is "too far" from its children, in order to preserve some quasi-isometry.

We first build a genealogy on the set of strips. We recall that the *root* of a strip S_i is the lowest vertex of its boundary, and is denoted by ρ_i . The *height* of S_i is the height of ρ_i . We call two strips *adjacent* if their respective boundaries share at least one edge. If S_i is a strip, we consider the first vertex on the left boundary of S_i (apart from ρ_i) that is also a branching point of \mathbf{T} . This vertex is also the root of some strips, exactly one of which is adjacent to S_i . We call this strip the *left child* of S_i (cf. Figure 12). We can similarly define its *right child*. Note that almost surely, every branch of \mathbf{T} branches eventually, so these childs always exist. We now fix a strip S_1 . We claim that the restriction of this genealogy to the set of descendants of S_1 is encoded by a complete binary tree, which we denote by T_{str} . Indeed, all the descendants of the left child of a strip S lie on the left of S , whereas all the descendants of its right child lie on its right. Therefore, it is not possible to obtain the same strip by two different genealogical lines from S_1 (see Figure 12).

We now kill some of the strips. We fix a constant $\ell_{\max} > 0$. For every strip S_i , let e_i^ℓ (resp. e_i^r) be the first edge on the left (resp. right) boundary of S_i that is also adjacent to its left (resp. right) child. We call a strip S_i *good* if, for every face f of S_i that is adjacent to ρ_i , the dual of S_i contains a path of length at most ℓ_{\max} from f to e_i^ℓ , and similarly for e_i^r .

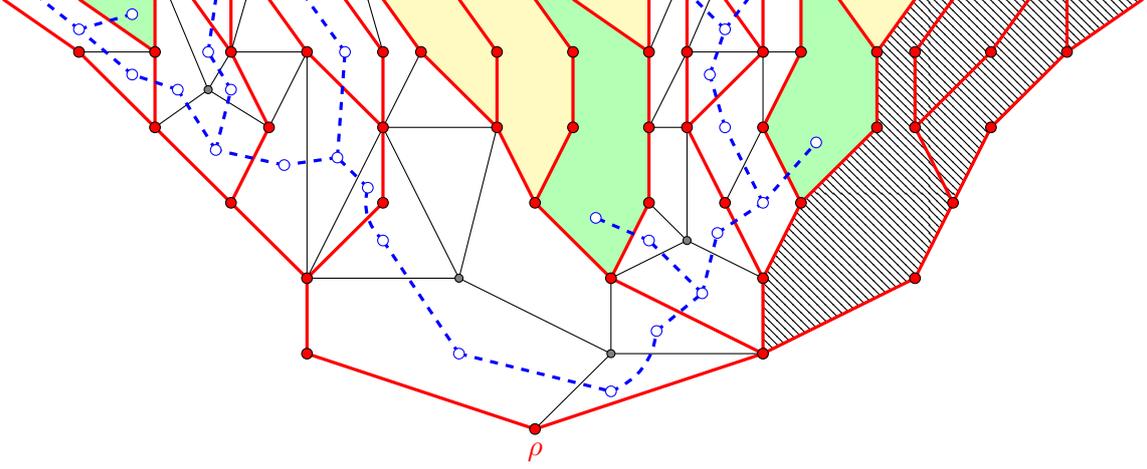


Figure 13: Construction of $\mathcal{S}_{\text{bdd}}^*$ (in blue). The tree \mathbf{T} is in red. The hatched strips do not belong to T_{str} . The green strips are the strips that are not good, and the yellow ones are the descendants of the green ones, so they are not in T'_{str} . For the sake of clarity, we have not drawn the interiors of the strips that do not belong to T'_{str} .

Note that the fact that S_i is good or not only depends on the internal geometry of S_i , and on the numbers of children of the vertices on the part of ∂S_i lying between e_i^ℓ and e_i^r . These parts for different values of $i \in T_{\text{str}}$ are disjoint. Hence, since the strips are i.i.d. and \mathbf{T} is a Galton–Watson tree, the events

$$\{S_i \text{ is good}\}$$

for $i \in T_{\text{str}}$ are independent, and have the same probability. Therefore, removing from T_{str} all the strips that are not good is equivalent to performing a Bernoulli site percolation on the complete binary tree T_{str} . Moreover, the probability for a strip to be good goes to 1 as ℓ_{max} goes to $+\infty$, so we can find ℓ_{max} such that this percolation is supercritical. We fix such an ℓ_{max} until the end of the proof.

Let T'_{str} be an infinite connected component of T_{str} containing only good strips. Then T'_{str} is a supercritical Galton–Watson tree and survives, so it is transient. We can now define a submap $\mathcal{S}_{\text{bdd}}^*$ of \mathcal{S}^* . For every strip $S_i \in T'_{\text{str}}$, let p_i^ℓ (resp. p_i^r) be a dual path of length at most ℓ_{max} joining the face of S_i that is adjacent to its parent to the face adjacent to e_i^ℓ (resp. e_i^r). Then the edges of $\mathcal{S}_{\text{bdd}}^*$ are the edges of these paths for all i such that $S_i \in T'_{\text{str}}$, as well as the dual edges of the edges e_i^ℓ and e_i^r (cf. Figure 13). The vertices of $\mathcal{S}_{\text{bdd}}^*$ are simply the vertices adjacent to these edges. Since the lengths of the paths p_i^ℓ and p_i^r are bounded by ℓ_{max} , it is easy to see that $\mathcal{S}_{\text{bdd}}^*$ is quasi-isometric to the tree T'_{str} , so it is transient. This implies that \mathcal{S}^* is transient as well, which concludes the proof of the first point. \square

Proof of point 2 of Theorem 4. We mimic the beginning of the proof of Theorem 2, but we use the first point of Theorem 4 instead of Proposition 6. The proof of the first two points is robust, and shows that almost surely, the simple random walk X on $\mathcal{M}(\mathbf{T}, (S_i)_{i \geq 0})$ converges a.s. to a point X_∞ of $\partial \mathbf{T}$, and that the distribution of X_∞ has a.s. full support. In particular, for every $x \in \mathbf{T}$, let $\hat{\partial} \mathbf{T}[y]$ be the set of the classes $\hat{\gamma}$, where γ is a ray passing through y . Let y_1, y_2 be

two vertices such that $\widehat{\partial\mathbf{T}}[y_1] \cap \widehat{\partial\mathbf{T}}[y_2] = \emptyset$. We define the function h on \mathcal{M} by

$$h(x) = P_{\mathcal{M},x}(X_\infty \in \partial\mathbf{T}[y_1]).$$

Then h is harmonic and bounded on \mathcal{M} . Moreover, by the same argument as in the beginning of the proof of Theorem 2, there is a sequence (x_n) of vertices in $\mathcal{S}[y_1]$ such that $h(x_n) \rightarrow 1$. On the other hand, by the first point of Theorem 4, there is a positive probability that X stays in $\mathcal{S}[y_2]$ eventually, so $h(x) < 1$ for every x , so h is non-constant. It follows that \mathcal{M} is non-Liouville. \square

Proof of point 3 of Theorem 4. We know from the proof of point 2 of Theorem 4 that the simple random walk X on $\mathcal{M}(\mathbf{T}, (S_i))$ converges a.s. to a point X_∞ on $\widehat{\partial\mathbf{T}}$. As in the proof of Theorem 2, it is enough to prove that the law of X_∞ is a.s. non-atomic, i.e. that if X and Y are two independent simple random walks on the same instance of $\mathcal{M}(\mathbf{T}, (S_i))$, then $X_\infty \neq Y_\infty$ a.s..

To prove this, we rely on a variant of Lemma 11. For $h \geq 0$, we denote by τ_h^X the smallest n for which X_n is a point of \mathbf{T} of height at least h . We define τ_h^Y similarly. The recurrence of the strips guarantees that the times τ_h^X and τ_h^Y are all a.s. finite. However, the vertices $X_{\tau_h^X}$ and $Y_{\tau_h^Y}$ may have height larger than h , which prevents us to re-use Lemma 11 as such. For $h \geq 1$, let \mathcal{F} be the σ -algebra generated by the finite tree $B_h(\mathbf{T})$, the family of the strips whose root is at height at most $h-1$ and the paths $(X_n)_{0 \leq n \leq \tau_h^X}$ and $(Y_n)_{0 \leq n \leq \tau_h^Y}$. Note that this makes sense since up to time τ_h^X , the walk X may only visit strips whose root as height at most $h-1$.

Let also A'_h be the following event:

$\left\{ \text{There are four distinct vertices } (x_i)_{1 \leq i \leq 4} \text{ of } T \text{ at height } \geq h, \text{ neither of which is an ancestor of another, such that:} \right.$

- the trees $\mathbf{T}[x_1]$, $\mathbf{T}[x_2]$, $\mathbf{T}[x_3]$ and $\mathbf{T}[x_4]$ lie in this cyclic order,
- for every $n \geq \tau_h^X + 2$, we have $X_n \in \mathcal{S}(\mathbf{T}, (S_i))[x_1]$
- for every $n \geq \tau_h^Y + 2$, we have $Y_n \in \mathcal{S}(\mathbf{T}, (S_i))[x_3]$.

As in the proof of Theorem 2, it is enough to prove that there is $\delta > 0$ such that almost surely, for h large enough, we have

$$\mathbb{P}(A'_h | \mathcal{F}_h) \geq \delta. \quad (7)$$

As in the proof of Lemma 11, there are several cases to treat separately according to the relative positions of $X_{\tau_h^X}$ and $Y_{\tau_h^Y}$. Since this is the most different case from the proof of Lemma 11, let us treat in details the case where $Y_{\tau_h^Y}$ is a strict descendant of $X_{\tau_h^X}$.

If this is the case, then right before τ_h^Y the walk Y lies in a strip S^Y intersecting $B_{h-1}(\mathbf{T})$. Without loss of generality, we assume that S^Y is on the right of $X_{\tau_h^X}$. As on Figure 14, the point $Y_{\tau_h^Y}$ must be the rightmost descendant of $X_{\tau_h^X}$ at its height. Note also that the conditioning on \mathcal{F}_h does not give any information about the numbers of children of the vertices of \mathbf{T} between $X_{\tau_h^X}$ and $Y_{\tau_h^Y}$.

Let K be a bound on the degrees in the strips. With conditional probability at least $(1 - \mu(1)) \times \frac{1}{3K}$, the vertex $X_{\tau_h^X}$ is a branching point in \mathbf{T} and $X_{\tau_h^X+1}$ is a child of $X_{\tau_h^X}$ which is not the rightmost one. If this occurs, let $x_1 = X_{\tau_h^X+1}$. Then the slice $\mathcal{S}(\mathbf{T}, (S_i))[x_1]$ has the same law as $\mathcal{S}(\mathbf{T}, (S_i))$, so by the first point of Theorem 4, the walk X has a probability bounded away from 0 to stay in $\mathcal{S}(\mathbf{T}, (S_i))[x_1]$ after $\tau_h^X + 1$. Similarly, independently of the behaviour of X , there is a probability at least $\frac{1-\mu(1)}{3K}$ that $x_3 = Y_{\tau_h^Y+1}$ is a child of $Y_{\tau_h^Y}$ which is not the

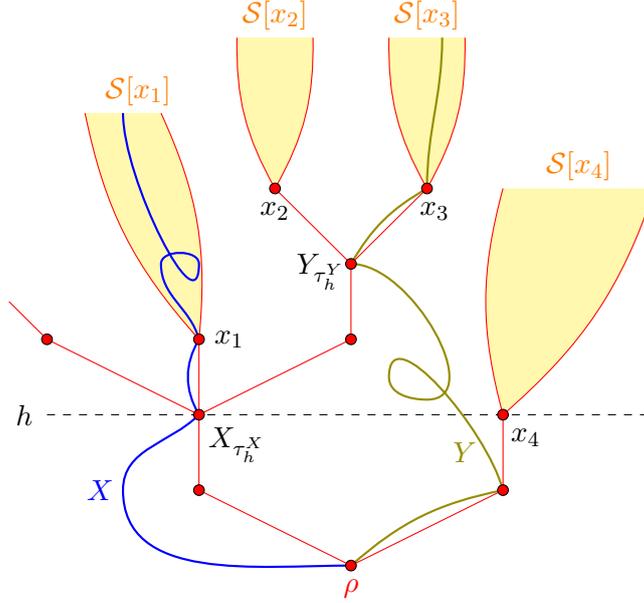


Figure 14: Sketch of the proof of (7) in the case where $Y_{\tau_h^Y}$ is a descendant of $X_{\tau_h^X}$. The tree \mathbf{T} is in red. The trajectories of X and Y are in blue and green respectively. For the sake of clarity, the interiors of the strips have not been drawn.

leftmost one, and that Y stays in $\mathcal{S}(\mathbf{T}, (S_i))[x_3]$ ever after τ_h^Y . To see that A'_h occurs if this is the case, we just need to take as x_2 the leftmost child of $Y_{\tau_h^Y}$, and as x_4 a vertex of \mathbf{T} which is different from $X_{\tau_h^X}$ (which a.s. exists if h is large enough).

This proves (7) in the case where $Y_{\tau_h^Y}$ is a strict descendant of $X_{\tau_h^X}$. As in the proof of Theorem 2, the other cases can be treated in a very similar way. This finishes the proof of Theorem 4. \square

4 Positive speed of the simple random walk

4.1 Sketch of the proof and definition of the half-plane model \mathcal{H}

The goal of this section is to prove Theorem 3. In all this section, we assume $\mu(0) = 0$. We first give a quick sketch of the proof. Our main task will be to prove positive speed in a half-planar model \mathcal{H} , constructed from an infinite forest of supercritical Galton–Watson trees, and where we have more vertical stationarity than in \mathcal{C} . The results of Section 3 will allow us to pass from \mathcal{H} to \mathcal{C} . Note that Theorem 3 is very easy in the case where $\mu(0) = \mu(1) = 0$, since then the height of the simple random walk dominates a random walk on \mathbb{Z} with positive drift. Therefore, we need to make sure that the vertices with only one child do not slow down the walk too much. Our proof in \mathcal{H} relies on two ingredients:

- An exploration method of \mathcal{H} will allow us to prove that the walk cannot be too far away from a point with at least two children. This will guarantee that the walk spends some time at vertices with at least two children. At these vertices, the height of the walk accumulates a positive drift. This will give a "quasi-positive speed" result: the height at

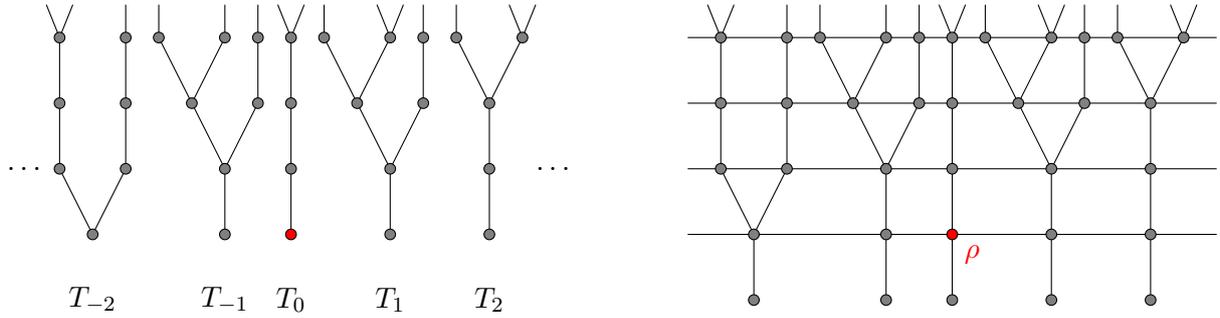


Figure 15: On the left, an infinite forest $(T_i)_{i \in \mathbb{Z}}$. On the right, the half-planar map \mathcal{H} obtained from this forest. The root vertex is in red.

time n is $n^{1-o(1)}$.

- Thanks to the stationarity properties of \mathcal{H} , we can study regeneration times, i.e. times at which the walk reaches some height for the first time, and stays above this height ever after. The estimates obtained in the first point are sufficient to prove that the first regeneration time has finite expectation. As in many other models (like random walks in random environments, see [28]), this is enough to ensure positive speed.

We now define our half-plane model. Let $(T_i)_{i \in \mathbb{Z}}$ be a family of i.i.d. Galton–Watson trees with offspring distribution μ . We draw the trees T_i with their roots on a horizontal line, and for every $h \geq 0$, we add horizontal connections between successive vertices of height h . Finally, for every vertex v of height 0, we add a parent of v at height -1 , which is linked only to v . We root the obtained map at the root of T_0 (which has height 0), and denote it by \mathcal{H} (see Figure 15). As for \mathcal{C} , if v is a vertex of \mathcal{H} , the height $h(v)$ of v in \mathcal{H} is defined as its height in the forest (T_i) . We denote by $\partial\mathcal{H}$ the set of vertices of \mathcal{H} at height -1 .

We first explain why it is enough to prove the result in \mathcal{H} instead of \mathcal{C} . We denote by X^G the simple random walk on a graph G .

Lemma 13. If in \mathcal{H} we have $\frac{1}{n}h(X_n^{\mathcal{H}}) \rightarrow v_\mu > 0$ a.s. as $n \rightarrow +\infty$, then Theorem 3 is true.

Proof. For every $i \in \mathbb{Z}$, let $\mathcal{S}(T_i)$ be the slice associated to the tree T_i . Then almost surely, if the walk (X_n) stays in $\mathcal{S}(T_i)$ eventually, the distance between X_n and the root of T_i is equivalent to $v_\mu n$. Since $\mathcal{S}(T_i)$ has the same distribution as \mathcal{S} , the simple random walk on \mathcal{S} has a.s. speed v_μ on the event that it does not hit $\partial\mathcal{S}$.

Back in \mathcal{C} , by Theorem 2, we know that $X^{\mathcal{C}}$ converges to a point $X_\infty \in \widehat{\partial}T$ and that X_∞ is a.s. not the point of $\widehat{\partial}T$ corresponding to the leftmost and rightmost rays of T . Therefore, almost surely, the walk $X^{\mathcal{C}}$ only hits the boundary of $\mathcal{S}(T)$ finitely many times, so it has a.s. speed $v_\mu > 0$. \square

4.2 An exploration method of \mathcal{H}

Let $k > 0$ and let x be a vertex of \mathcal{H} at height $h \geq 0$. We write $x_0 = x$. Let also x_{-1}, \dots, x_{-k} be the k first neighbours at height h on the left of x , and let x_1, \dots, x_k be the k first neighbours of x on its right.

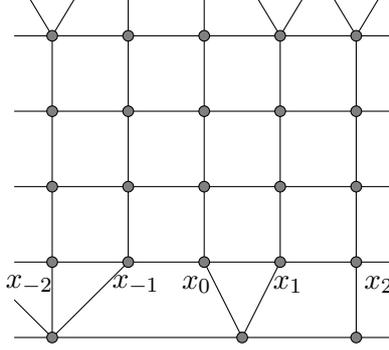


Figure 16: An example for Definition 14: the point $x = x_0$ is 2-bad, but not 3-bad.

Definition 14. Let x be a vertex of \mathcal{H} at height $h \geq 0$. We say that x is k -bad if all the descendants of $x_{-k}, x_{-(k-1)}, \dots, x_k$ at heights $h, h+1, \dots, h+k$ have only one child (cf. Figure 16).

The goal of this section is to show that the probability for the walk to visit a "very bad" point in its first n steps is very small. More precisely, we will prove the following result.

Lemma 15. There is a constant $c > 0$ such that for every $k, n > 0$, we have

$$\mathbb{P}(\text{one of the points } X_0, X_1, \dots, X_n \text{ is } k\text{-bad}) \leq ck(n+1)^2 \mu(1)^{k^2}.$$

The presence of the factor $\mu(1)^{k^2}$ is not surprising. For example, the probability for the root vertex X_0 to be k -bad is exactly $\mu(1)^{(k+1)(2k+1)}$. The idea of the proof is to explore \mathcal{H} at the same time as the random walk moves, in such a way that every time we discover a new vertex, either its k neighbours on the right or its k neighbours on the left have all their descendants undiscovered. If this is the case, the probability for the discovered vertex to be k -bad is at most $\mu(1)^{k^2}$. The factor $ck(n+1)^2$ means that our exploration method needs to explore at most $ck(n+1)^2$ vertices to discover $\{X_0, X_1, \dots, X_n\}$.

The rest of Section 4.2 is devoted to the proof of Lemma 15. We first define our exploration method. We then state precisely the properties that we need this exploration to satisfy (Lemmas 16, 17 and 18), and explain how to conclude the proof from here. Finally, we prove these properties.

Exploration methods. By an *exploration method*, we mean a nondecreasing sequence $(E_i)_{i \geq 0}$ of finite sets of vertices of \mathcal{H} . For every $i \geq 0$, the part of \mathcal{H} discovered at time i is the finite map formed by:

- the vertices of E_i ,
- the edges of \mathcal{H} whose two endpoints belong to E_i ,
- for every edge of \mathcal{H} with exactly one endpoint in E_i , the half-edge adjacent to E_i .

We denote this map by \mathcal{E}_i (see Figure 17). In particular, when we explore a vertex, we know how many children it has, even if these children are yet undiscovered. Moreover, for every i ,

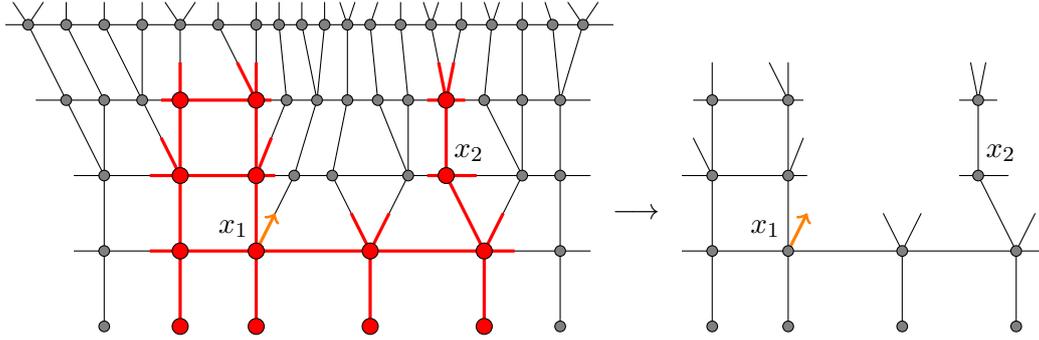


Figure 17: On the left, the map \mathcal{H} and the set of vertices E_i (in red). On the right, the explored map \mathcal{E}_i . The edge in orange means that at the current time, the random walk is leaving the vertex x_1 towards its right child. The vertex x_1 is 2-free on the right, but not 3-free. The vertex x_2 is k -free on the right for every $k \geq 0$. It is also 5-free on the left, but not 6-free. By only looking at \mathcal{E}_i , we can be sure that x_1 is 2-free on the right and x_2 is 4-free on the left, but not that x_2 is 5-free on the left.

the map \mathcal{E}_i will be equipped with an additional marked oriented edge or half-edge e_i describing the current position of the simple random walk. In all the explorations we will consider, we can pass from \mathcal{E}_{i-1} to \mathcal{E}_i by either adding one vertex to the explored set E_{i-1} , or moving the marked edge to a neighbour edge or half-edge. In the first case, we call i an *exploration step* and, in the second, we call i a *walk step*. In particular, the time of the exploration is not the same as the time of the random walk.

We say that a vertex $v \in E_i$ is *k -free on the right* in E_i if none of the k first neighbours on the right of the rightmost child of E_i belongs to E_i . We define similarly a *k -free on the left* vertex. We would like to build an exploration of \mathcal{H} such that at every exploration step i , the unique vertex of $E_{i+1} \setminus E_i$ is either k -free on the left or on the right in E_{i+1} . Note that it is not always possible by looking at \mathcal{E}_i to decide whether a vertex $v \in E_i$ is k -free or not (cf. vertex v_2 on Figure 17). However, as we will see later (proof of Lemma 18), it is sometimes possible to be sure that a vertex is k -free.

Choice of the exploration method. The simplest exploration method coming into mind is to explore a vertex when it is hit for the first time by the walk (X_n) . However, this is not suitable for our purpose. If for example we discover some descendants of a vertex v and explore v afterwards, then we have some partial information about the descendance of v , so we cannot control the probability for v to be k -bad. For this reason, we never want to discover a vertex before discovering all its ancestors. We will therefore require that the sets E_i are *stable*, which means that for every vertex $v \in E_i$, all the ancestors of v lie in E_i as well.

A second natural exploration method is now the following: everytime the walk (X_n) hits a vertex v for the first time, we discover all the ancestors of v that are yet undiscovered (including v), from the lowest to the highest. Although more convenient than the first one, this method is not sufficient either. Indeed, assume that at some point we explore a vertex v such that the k first neighbours on the left and on the right of v have already been discovered, as well as many of their descendants. Then we have accumulated some partial information about the descendants

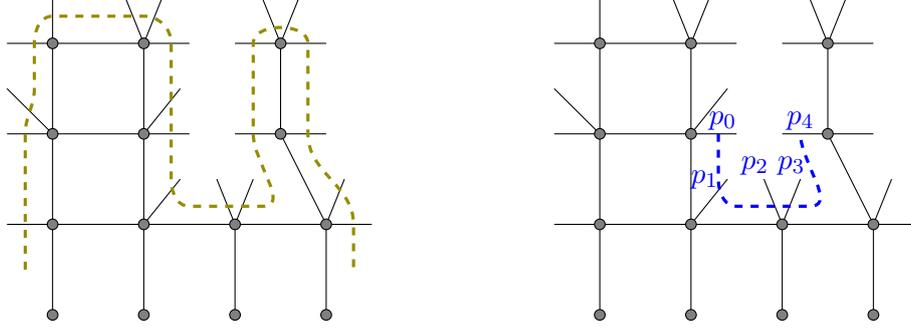


Figure 18: On the left, we move around \mathcal{E}_i by crossing all the half-edges. On the right, the unique pit of \mathcal{E}_i . It has width 3 and height 0. In particular, the map \mathcal{E}_i is 1-flat, but not 2-flat.

of all the neighbours of v , so we cannot control the probability for v to be k -bad.

More generally, this problem occurs if we allow the formation of "narrow pits" in the explored part of \mathcal{H} . Therefore, the idea of our exploration is the following: everytime the walk (X_n) hits a vertex v for the first time, we explore all the ancestors of v that are yet undiscovered (including v), from the lowest to the highest. If by doing so we create a pit of width at most $2k$, then we explore completely the bottom of this pit, until it has width greater than $2k$.

To define this exploration more precisely, we need to define precisely a pit. We assume that the set E_i is stable, which implies that \mathcal{E}_i is simply connected. Then there is a unique way to move around the map \mathcal{E}_i from left to right by crossing all the half-edges exactly once, as on Figure 18. A *pit* is a sequence of consecutive half-edges p_0, p_1, \dots, p_{j+1} such that:

- the half-edges p_1, \dots, p_j point upwards, and start from the same height h ,
- the half-edge p_0 points to the right and lies at height $h + 1$,
- the half-edge p_{j+1} points to the left and lies at height $h + 1$.

Note that all the pits are half-edge-disjoint. We call j and h the *width* and the *height* of the pit. Finally, we say that \mathcal{E}_i is k -flat if it has no pit of width $j \leq 2k$ (see Figure 18).

If e is an oriented edge or half-edge, we will denote by e^- its starting point and e^+ its endpoint. We can now describe our exploration algorithm precisely. We take for E_0 the set formed by the root vertex ρ of \mathcal{H} and its parent at height -1 , and pick e_0 uniformly among all the edges and half-edges started from ρ . For every $i \geq 0$, we recall that e_i is the oriented edge or half-edge of \mathcal{E}_i marking the position of the simple random walk. For every $i \geq 1$, given $(\mathcal{E}_{i-1}, e_{i-1})$, we construct (\mathcal{E}_i, e_i) as follows.

- If the marked edge e_{i-1} is a full edge and \mathcal{E}_{i-1} is k -flat, we perform a *walk step*: we set $E_i = E_{i-1}$ and pick e_i uniformly among all the edges and half-edges whose starting point is e_{i-1}^+ .
- If e_{i-1} is a half-edge, we perform an *exploration step*: we denote by v_i the lowest ancestor of e_{i-1}^+ that does not belong to E_{i-1} . If v_i lies at height -1 , then E_i is the union of E_{i-1} , the vertex v_i and its child at height 0 (this is the only case where we explore two vertices at once, to make sure that \mathcal{E}_i remains connected). If not, then $E_i = E_{i-1} \cup \{v_i\}$. Note that if $v_i = e_{i-1}^+$, then the marked half-edge becomes a full edge.

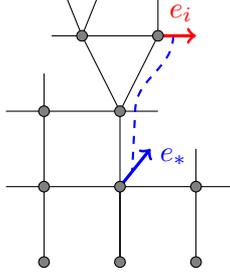


Figure 19: The second case of our exploration algorithm. We move around the boundary of \mathcal{E}_i towards the right: we encounter one horizontal half-edge, and then a vertical half-edge e_* . This e_* has descendants at the same height as e_i , and the leftmost such descendant is e_i^+ .

- (iii) If e_{i-1} is a full edge but \mathcal{E}_{i-1} is not k -flat, we also perform an *exploration step*: let $(p_0, p_1, \dots, p_{j+1})$ be the leftmost pit of width at most $2k$. We then take $E_i = E_{i-1} \cup \{p_1^+\}$. This means that we explore the endpoint of the leftmost vertical half-edge of the pit.

It is easy to check that the sets E_i we just defined are all stable.

The exploration is Markovian. An important feature of our exploration that we need to check is that it is Markovian. More precisely, for every $i \geq 0$, let ∂E_i be the set of vertices consisting of:

- the endpoints of the half-edges of \mathcal{E}_i pointing upwards,
- the vertices of \mathcal{H} of height 0 that do not belong to E_i .

Lemma 16. Conditionally on $(\mathcal{E}_j, e_j)_{0 \leq j \leq i}$, the trees of descendants of the vertices of ∂E_i are independent Galton–Watson trees with offspring distribution μ .

Proof. Given the Markovian structure of supercritical Galton–Watson tree, it is enough to check that at every step, our exploration is independent of the part of \mathcal{H} that has not yet been discovered. This is easy in the case (i): conditionally on \mathcal{E}_i , the choice of e_i is independent of the rest of \mathcal{H} . We claim that in the other two cases, the discovered vertex is either the endpoint of a vertical half-edge which is a deterministic function of (\mathcal{E}_i, e_i) , or the root of the first infinite tree on the left or on the right of \mathcal{E}_i . This claim is obvious in the case (iii).

In the case (ii), the half-edge e_i points either to the top, the left or the right. If it points to the top, then all the ancestors of e_i^+ have been discovered, so the explored vertex is e_i^+ . If e_i does not point to the top, we assume without loss of generality that it points to the right. We then start from the half-edge e_i and move around \mathcal{E}_i towards the right. We first cross horizontal half-edges pointing to the right at decreasing heights, until either we reach height 0, or we cross a vertical half-edge pointing to the top. If we cross a first vertical half-edge e , then e must be an ancestor of e_i^+ (indeed, e has descendants at the same height as e_i^+ , and there is no other half-edge pointing to the top between e_i^- and e_*^- , see Figure 19). Therefore, the explored vertex must be e_*^+ . Finally, if we reach the bottom boundary, then the explored vertex is the root of the first tree on the right of \mathcal{E}_i (and its parent at height -1).

□

We denote by $\varphi(n)$ the n -th walk step of our exploration, with $\varphi(0) = 0$. Note that the marked edge or half-edge from the time $\varphi(n)$ to the time $\varphi(n+1) - 1$ corresponds to the edge in \mathcal{H} from X_n to X_{n+1} . Therefore, at time $\varphi(n)$, the explored part covers $\{X_0, X_1, \dots, X_n\}$. As explained in the beginning of this subsection, it is important to control φ . We can now state the two important properties that our exploration satisfies.

Lemma 17. There is a deterministic constant c such that for every $n \geq 0$, we have

$$\varphi(n+1) - \varphi(n) \leq ck(n+1).$$

Lemma 18. For every exploration step i , the unique vertex v_i of $E_i \setminus E_{i-1}$ that does not lie at height -1 is either k -free on the left or k -free on the right.

Note that our exploration method was precisely designed to satisfy Lemma 18. Both these lemmas are completely deterministic: they hold if we replace \mathcal{H} by any infinite causal map with no leaf, and (X_n) by any infinite path. Before proving them, we explain how to conclude the proof of Lemma 15 given these two results.

Proof of Lemma 15 given Lemmas 17 and 18. For every $i \geq 0$, let \mathcal{F}_i be the σ -algebra generated by $(\mathcal{E}_j, e_j)_{0 \leq j \leq i}$. We first show that, for every exploration step $i \geq 0$, we have

$$\mathbb{P}(v_i \text{ is } k\text{-bad} | \mathcal{F}_i) \leq \mu(1)^{k^2}. \quad (8)$$

Let $i \geq 0$ be an exploration step. By Lemma 18, without loss of generality, we may assume that v_i is k -free on the right. Let v_i^1, \dots, v_i^k be the k first neighbours on the right of the rightmost child of v_i . Since v_i is k -free on the right, these vertices do not belong to E_i . By Lemma 16, conditionally on \mathcal{F}_i , their trees of descendants are i.i.d. Galton–Watson trees with offspring distribution μ . If the vertex v_i is k -bad, each of the vertices v_i^1, \dots, v_i^k has only one descendant at height $h+k+1$, so we have

$$\begin{aligned} \mathbb{P}(v_i \text{ is } k\text{-bad} | \mathcal{F}_i) &\leq \mathbb{P}(v_i^1, \dots, v_i^k \text{ have one descendant each at height } h+k+1 | \mathcal{F}_i) \\ &= \prod_{j=1}^k \mathbb{P}(v_i^j \text{ has exactly one descendant at height } h+k+1 | \mathcal{F}_i) \\ &= \left(\mu(1)^k\right)^k, \end{aligned}$$

and we obtain (8), which implies

$$\mathbb{P}(i \text{ is an exploration step and } v_i \text{ is } k\text{-bad}) \leq \mu(1)^{k^2}$$

for every $i \geq 0$. By summing Lemma 17, we obtain $\varphi(n+1) \leq ck(n+1)^2$. Therefore, the vertices X_0, X_1, \dots, X_n all lie in $\mathcal{E}_{ck(n+1)^2}$. If one of these points is k -bad, it cannot lie at height -1 by definition of a bad point, so it is equal to v_i for some $0 \leq i \leq ck(n+1)^2$. Therefore, we have

$$\begin{aligned} \mathbb{P}(\text{one of the points } X_0, X_1, \dots, X_n \text{ is } k\text{-bad}) &\leq \sum_{i=0}^{ck(n+1)^2} \mathbb{P}(v_i \text{ is } k\text{-bad}) \\ &\leq ck(n+1)^2 \mu(1)^{k^2}, \end{aligned}$$

which ends the proof. □

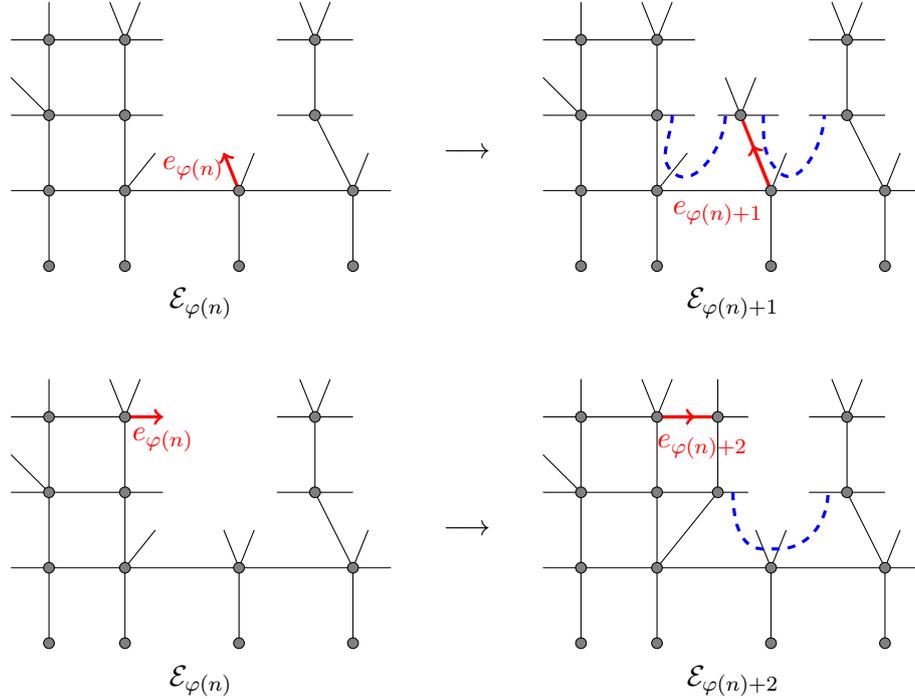


Figure 20: The two ways our exploration can create new pits between two walk steps. On the top, a pit of width 3 is split into two pits of width 1. On the bottom, the width of a pit decreases from 3 to 2. The new pits of width at most 2 are indicated in blue.

Proof of Lemma 17. To bound $\varphi(n+1) - \varphi(n)$, we describe precisely what happens between the times $\varphi(n)$ and $\varphi(n+1)$. Note that $\mathcal{E}_{\varphi(n)} = \mathcal{E}_{\varphi(n)-1}$ is k -flat (if it was not, $\varphi(n)$ would have to be an exploration step). Hence, if $e_{\varphi(n)}$ is a full edge, then $\varphi(n)+1$ is a walk step and we have $\varphi(n+1) = \varphi(n)+1$, so it is only necessary to treat the case where $e_{\varphi(n)}$ is a half-edge.

In this case, the first thing our algorithm does is to explore the vertex $e_{\varphi(n)}^+$ and all its undiscovered ancestors. We know that $e_{\varphi(n)}^+ = X_{n+1}$, so in particular its height is at most $n+1$. Therefore, exploring all its ancestors takes at most $n+2$ steps.

We can now perform the $(n+1)$ -th walk step, except if exploring $e_{\varphi(n)}^+$ and its ancestors has created a new pit of width at most $2k$. This can happen in two different ways, as on Figure 20:

- if $e_{\varphi(n)}$ is vertical, exploring $e_{\varphi(n)}^+$ may split an existing pit in two,
- if $e_{\varphi(n)}$ is horizontal (say it points to the right), exploring its ancestors may decrease the width of an existing pit on the right of $e_{\varphi(n)}$.

Note that in the first case, we can create at most two new pits, whereas in the second, we can shrink only one (cf. Figure 20). Hence, we will create at most 2 narrow pits, at the same height h .

In the first case (top of Figure 20), our algorithm will then fill the pit on the left if it has width at most $2k$, and then the pit on the right. The number of steps this takes is at most $2 \times 2k = 4k$. By doing so, we may create a new pit at height $h+1$. However, since the trees we work with have no leaf, this pit is at least as wide as the pit of $\mathcal{E}_{\varphi(n)}$ in which $e_{\varphi(n)}$ lies. Hence,

the new pit at height $h + 1$ has width greater than $2k$, and does not need to be filled. Therefore, in the first case, the number of exploration steps needed to fill all the narrow pits is at most $4k$.

In the second case (bottom of Figure 20), if the pit has width at most $2k$, our algorithm will explore all the vertical half-edges of this pit, from left to right. This takes at most $2k$ steps. Once again, this creates a new pit at height $h + 1$, but this time this pit may have width $2k$ or less. If this is the case, our algorithm will explore all its half-edges and perhaps create a pit at height $h + 2$, and so on. Note that the maximal height of $\mathcal{E}_{\varphi(n)}$ is at most n . Indeed, the only times at which this maximal height increases is when the random walk reaches some height for the first time, so the maximal height of $\mathcal{E}_{\varphi(n)}$ cannot be larger than the maximal height of the random walk during its first n steps. Therefore, we will need to fill at most n pits (at heights between 0 and $n - 1$), each of which taking at most $2k$ steps. Hence, filling the narrow pits takes at most $2kn$ steps.

Therefore, in both cases, the number of steps needed to obtain a k -flat map and perform a new walk step is bounded by $\max(4k, 2kn)$. If we add the number of steps needed to explore the ancestors of $e_{\varphi(n)}^+$ and the walk step $\varphi(n + 1)$, we obtain

$$\varphi(n + 1) - \varphi(n) \leq \max(4k, 2kn) + (n + 2) + 1 \leq ck(n + 1),$$

with e.g. $c = 7$. □

We finally prove Lemma 18. The proof will make use of the "step by step" description of our exploration that we also used in the last proof. We recall that for every exploration step i , we call v_i the unique vertex of nonnegative height in $E_i \setminus E_{i-1}$, and e_i the oriented edge or half-edge marking the current position of the random walk.

Proof of Lemma 18. We fix an exploration step $i \geq 0$. Note that the vertex v_i is always the endpoint of some half-edge of \mathcal{E}_{i-1} , that we denote by e_* .

Before moving on to the details of the proof, we explain how it is possible, by only looking at the map \mathcal{E}_{i-1} , to be sure that the vertex v_i is k -free on the right in E_i . We move along the boundary of \mathcal{E}_{i-1} from e_* towards the right, and stop when we encounter a vertex of height $h(v_i) + 1$. If this never occurs, it means that in \mathcal{E}_{i-1} and \mathcal{E}_i , there is no vertex at height $h(v_i) + 1$ on the right of v_i , so v_i is k -free on the right. If this occurs, assume that by moving so, we cross at least k vertical half-edges. Since the trees we consider have no leaf, all these vertical half-edges have descendants at height $h(v_i) + 1$, which lie on the right of all the children of v_i . Moreover, none of these descendants belongs to \mathcal{E}_i . Therefore, v_i must be k -free on the right in \mathcal{E}_i . Of course, this is also true for k -free on the left vertices (see the end of the caption of Figure 17 for an example). This remark will be implicitly used in all the cases below.

Let n be the integer such that $\varphi(n) < i < \varphi(n + 1)$. We distinguish two cases, corresponding to the two "phases" of exploration between $\varphi(n)$ and $\varphi(n + 1)$ that we described in the proof of Lemma 17. Both of these cases will be separated in a few subcases.

- We first treat the case where e_{i-1} is a half-edge, so the explored vertex v_i is an ancestor of e_{i-1}^+ .
 - We start with the subcase where e_{i-1} is vertical, and lies in a pit p . Since e_{i-1} is vertical, exploring the ancestors of e_{i-1}^+ takes only one step, so $i = \varphi(n) + 1$ and $v_i = e_{i-1}^+$. Since $i - 1$ is a walk step, the pit p has width at least $2k + 1$. Without

loss of generality, we may assume that at least k of the vertical half-edges of p are on the right of e_{i-1} , so v_i is k -free on the right.

- If e_{i-1} is vertical but is not in a pit, as in the previous case, we have $i = \varphi(n) + 1$ and v_i is the endpoint of e_{i-1} . Moreover, there is a direction (left or right) such that when we start from e_{i-1} and move along the boundary of \mathcal{E}_{i-1} in this direction, the height decreases before increasing for the first time (if not, e_{i-1} would be in a pit). Without loss of generality, this direction is the right. Let p be the first pit that we encounter on the right of e_{i-1} . If p does not exist, it means that the height never increases again, so there is no vertex on the right of e_{i-1} in \mathcal{E}_{i-1} that is higher than e_{i-1}^- . Therefore, the vertex $v_i = e_{i-1}^+$ is k -free on the right in \mathcal{E}_i . If p exists, it has width at least $2k + 1$, so v_i is $(2k + 1)$ -free on the right, and in particular k -free.
 - If e_{i-1} is horizontal, without loss of generality it points to the right. As explained earlier (see Figure 19), the edge e_* is in this case the first vertical half-edge we meet when we move around \mathcal{E}_{i-1} from e_{i-1} towards the right (except if no such vertical half-edge exists, in which case v_i is the root of a new tree on the right of \mathcal{E}_{i-1} , and v_i is obviously k -free on the right). By the same argument as in the previous case, if there is no pit on the right of e_* , then v_i is k -free on the right (and even ∞ -free). If there is one and p is the first such pit, note that between the times $\varphi(n)$ and i , the pit p has either been untouched, or has been shrunk by 1. Therefore, at time i it has width at least $2k$, so v_i is k -free in E_i .
- We now consider the second "phase", i.e. the case where \mathcal{E}_{i-1} has a pit of width at most $2k$, and the goal of the exploration step i is to fill it.
 - If e_{i-1} points to the top, then the pit has been created at time $\varphi(n) + 1$ as in the top part of Figure 20. Hence, the half-edge e_* belonged at time $\varphi(n)$ to a pit $(p_0, p_1, \dots, p_{j+1})$ of height h and width $j > 2k$. Therefore, either k of the vertical half-edges p_1, \dots, p_j lie on the left of e_* , or k of them lie on its right (the two cases are not symmetric since the pit is filled from left to right). If k of these half-edges lie on the left of e_* , then their k endpoints (of height $h + 1$) have been explored before v_i , but none of the descendants of these endpoints has been discovered. Therefore, the map \mathcal{E}_i contains at least k vertical half-edges at height $h + 1$ on the left of v_i , so v_i is k -free on the left. This case is the reason why, in the definition of a k -free vertex, we asked the neighbours of the *children* of v to be undiscovered, and not simply the neighbours of v . If k of the half-edges of p lie on the right of e_* , the argument is similar (it is actually simpler since the half-edges on the right of e_* have not yet been explored).
 - If e_{i-1} points to the right, then we are in the bottom case of Figure 20: a pit p of width $2k + 1$ has been shrunk to width $2k$ during the first phase, resulting in a pit $(p_0, p_1, \dots, p_{2k+1})$ of width $2k$ at some height h . Let also $h' \geq h$ be the height of e_* . Since the pit is filled layer by layer from the bottom, the half-edge e_* must be a descendant of a half-edge p_{ℓ_0} with $1 \leq \ell_0 \leq 2k$. Moreover, our algorithm fills the layers from left to right. Therefore, at time i , for every $1 \leq \ell \leq 2k$, we have already explored the descendants of e_ℓ up to height $h' + 1$ if $\ell \leq \ell_0$ and up to height h' if $\ell > \ell_0$. But the vertex v_i lies at height $h' + 1$. Therefore, if $\ell_0 \geq k + 1$, then v_i has k

vertical half-edges on its left so it is k -free on the left. On the other hand, if $\ell_0 \leq k$, then v_i has k vertical half-edges on its right at height h , so it is k -free on the right. This concludes the proof. □

4.3 Quasi-positive speed in \mathcal{H}

The goal of this subsection is to use Lemma 15 to prove that the walk has a speed $n^{1-o(1)}$, which is slightly weaker than positive speed. We will need to "bootstrap" this result in Section 4.4 to obtain positive speed. We denote by H_n the height of X_n . We also write

$$D_n = \max\{H_k - H_\ell \mid 0 \leq k < \ell \leq n\}$$

for the greatest "descent" of X before time n .

Proposition 19. Let $0 < \delta < 1$ and $\beta > 0$. Then we have

$$\mathbb{P}(H_n \leq n^{1-\delta}) = o(n^{-\beta}) \quad \text{and} \quad \mathbb{P}(D_n \geq n^\delta) = o(n^{-\beta})$$

as $n \rightarrow +\infty$.

Proof. We start with the proof of the first estimate, the proof of the second will follow the same lines. We call a vertex of \mathcal{H} *good* if it has height -1 or if it has at least two children. The idea of the proof is the following: by Lemma 15, with high probability, the walk does not visit any k -bad point before time n for some k . Hence, it is never too far from a good point. Therefore, the walk always has a reasonable probability to reach a good point in a near future. It follows that X will visit many good points, so (H_n) will accumulate a large positive drift.

More precisely, let $a > 0$ (we will take a large later). We define two events formalizing the ideas we just explained:

$$\begin{aligned} A_1 &= \{\text{none of the vertices } X_0, X_1, \dots, X_n \text{ is } a\sqrt{\log n}\text{-bad}\}, \\ A_2 &= \{\text{for every } 0 \leq m \leq n - n^{\delta/2}, \text{ one of the points } X_m, X_{m+1}, \dots, X_{m+n^{\delta/2}} \text{ is good}\}. \end{aligned}$$

Then we have

$$\mathbb{P}(H_n \leq n^{1-\delta}) \leq \mathbb{P}(A_1^c) + \mathbb{P}(A_1 \setminus A_2) + \mathbb{P}(A_2 \cap \{H_n \leq n^{1-\delta}\}). \quad (9)$$

We start with the first term. By Lemma 15, we have

$$\mathbb{P}(A_1^c) \leq ca\sqrt{\log n} (n+1)^2 \mu(1)^{a^2 \log n}.$$

Hence, if we choose a large enough (i.e. $a^2 > \frac{\beta+2}{-\log \mu(1)}$), we have $\mathbb{P}(A_1^c) = o(n^{-\beta})$.

We now bound the second term of (9). For every $0 \leq m \leq n$, let \mathcal{F}_m be the σ -algebra generated by \mathcal{H} and (X_0, X_1, \dots, X_m) . If X_m is not $a\sqrt{\log n}$ -bad (which is an \mathcal{F}_m -measurable event), let Y be the closest good vertex from X_m (we may have $Y = X_m$). We have $d_{\mathcal{H}}(X_m, Y) \leq 2a\sqrt{\log n}$, so there is a path from X_m to Y of length at most $2a\sqrt{\log n}$, and visiting only vertices of degree 4 (except of course Y). Therefore, we have

$$\mathbb{P}(X \text{ visits the vertex } Y \text{ between time } m \text{ and time } m + 2a\sqrt{\log n} \mid \mathcal{F}_m) \geq \left(\frac{1}{4}\right)^{2a\sqrt{\log n}}$$

if X_m is not $a\sqrt{\log n}$ -bad. By induction on i , we easily obtain, for every $i \geq 0$,

$$\begin{aligned} \mathbb{P}\left(X_m, X_{m+1}, \dots, X_{m+2ia\sqrt{\log n}} \text{ are neither good nor } a\sqrt{\log n}\text{-bad}\right) &\leq \left(1 - \frac{1}{4^{2a\sqrt{\log n}}}\right)^i \\ &\leq \exp\left(-\frac{i}{4^{2a\sqrt{\log n}}}\right). \end{aligned}$$

In particular, by taking $i = \frac{n^{\delta/2}}{2a\sqrt{\log n}}$, we obtain, for every m :

$$\mathbb{P}\left(X_m, X_{m+1}, \dots, X_{m+n^{\delta/2}} \text{ are neither good nor } a\sqrt{\log n}\text{-bad}\right) \leq \exp\left(-\frac{n^{\delta/2}}{2a\sqrt{\log n} 4^{2a\sqrt{\log n}}}\right).$$

If the event $A_1 \setminus A_2$ occurs, then there is an m with $0 \leq m \leq n - n^{\delta/2}$ such that the above event occurs. Therefore, by summing the last equation over $0 \leq m \leq n - n^{\delta/2}$, we obtain

$$\mathbb{P}(A_1 \setminus A_2) \leq n \exp\left(-\frac{n^{\delta/2}}{2a\sqrt{\log n} 4^{2a\sqrt{\log n}}}\right) = o(n^{-\beta}).$$

Finally, we bound the third term of (9) by the Azuma inequality. For every $n \geq 0$, let

$$M_n = H_n - \sum_{i=0}^{n-1} E_{\mathcal{H}}[H_{i+1} - H_i | X_0, X_1, \dots, X_i].$$

It is clear that M is a martingale with $|M_{n+1} - M_n| \leq 2$ for every n , and $M_0 = 0$. Moreover, we have

$$E_{\mathcal{H}}[H_{i+1} - H_i | X_0, X_1, \dots, X_i] = \frac{c(X_i) - 1}{c(X_i) + 3} \mathbb{1}_{X_i \notin \partial\mathcal{H}} + \mathbb{1}_{X_i \in \partial\mathcal{H}},$$

where we recall that $c(v)$ is the number of children of a vertex v . In particular, we have $E_{\mathcal{H}}[H_{i+1} - H_i | X_0, X_1, \dots, X_i] \geq 0$, and $E_{\mathcal{H}}[H_{i+1} - H_i | X_0, X_1, \dots, X_i] \geq \frac{1}{5}$ if X_i is a good vertex. If A_2 occurs, the walk X must visit at least $n^{1-\delta/2}$ good vertices before time n , so we have

$$\sum_{i=0}^{n-1} E_{\mathcal{H}}[H_{i+1} - H_i | X_0, X_1, \dots, X_i] \geq \frac{1}{5} n^{1-\delta/2}.$$

Therefore, if the event in the third term of (9) occurs, we have

$$M_n \leq n^{1-\delta} - \frac{1}{5} n^{1-\delta/2} < 0.$$

On the other hand, the Azuma inequality applied to M gives

$$P_{\mathcal{H},\rho}\left(M_n \leq n^{1-\delta} - \frac{1}{5} n^{1-\delta/2}\right) \leq \exp\left(-\frac{1}{8n} \left(\frac{1}{5} n^{1-\delta/2} - n^{1-\delta}\right)^2\right),$$

so

$$\mathbb{P}\left(M_n \leq n^{1-\delta} - \frac{1}{5} n^{1-\delta/2}\right) = o(n^{-\beta}),$$

which bounds the third term of (9), and proves the first part of Proposition 19.

To prove the second part, we decompose the event $\{D_n \geq n^\delta\}$ in the same way as in (9). By the definition of D_n , it is enough to show

$$\max_{0 \leq k \leq \ell \leq n} \mathbb{P} \left(A_2 \cap \{H_\ell - H_k \leq -n^\delta\} \right) = o(n^{-(\beta+2)}), \quad (10)$$

and then to sum over k and ℓ . To prove (10), note that if $\ell < k + n^\delta$, then $H_\ell - H_k > -n^\delta$ deterministically. If $\ell \geq k + n^\delta$, we use the same argument based on the Azuma inequality as for the first part. Let $0 \leq k < k + n^\delta \leq \ell \leq n$. If A_2 occurs, then X visits at least $\frac{\ell-k}{n^{\delta/2}}$ good vertices between times k and ℓ , so

$$\sum_{i=k}^{\ell-1} E_{\mathcal{H}}[H_{i+1} - H_i | X_0, \dots, X_i] \geq \frac{1}{5} \frac{\ell - k}{n^{\delta/2}}.$$

Hence, if the event of (10) occurs for k and ℓ , we have

$$M_\ell - M_k \leq H_\ell - H_k - \frac{1}{5} \frac{\ell - k}{n^{\delta/2}} \leq -n^\delta - \frac{1}{5} \frac{\ell - k}{n^{\delta/2}} \leq -\frac{2}{\sqrt{5}} (\ell - k)^{1/2} n^{\delta/4}.$$

But the Azuma inequality gives

$$\begin{aligned} \mathbb{P} \left(M_\ell - M_k \leq -\frac{2}{\sqrt{5}} (\ell - k)^{1/2} n^{\delta/4} \right) &\leq \exp \left(-\frac{(2(\ell - k)^{1/2} n^{\delta/4})^2}{8 \times 5(\ell - k)} \right) \\ &= \exp \left(-\frac{n^{\delta/2}}{10} \right) \\ &= o(n^{-(\beta+2)}), \end{aligned}$$

which proves (10) and the second point of Proposition 19. \square

4.4 Positive speed in \mathcal{H} via regeneration times

For every $0 \leq h < h' \leq +\infty$, we denote by $\mathcal{B}_{h,h'}$ the map formed by the vertices of \mathcal{H} with height in $\{h, h+1, \dots, h'\}$, in which for every vertex v at height h , we have added a vertex below v that is linked only to v . We root $\mathcal{B}_{h,h'}$ at the vertex ρ_h corresponding to the leftmost descendant of ρ at generation h . The *height* of a vertex in $\mathcal{B}_{h,h'}$ is its height in \mathcal{H} , minus h , and the height of the additional vertices is -1 . We denote by $\partial\mathcal{B}_{h,h'}$ the set of these additional vertices. Note that for any $h \geq 0$, the rooted map $(\mathcal{B}_{h,\infty}, \rho_h)$ is independent of $\mathcal{B}_{0,h}$ and has the same distribution as (\mathcal{H}, ρ) . Since this distribution is invariant by horizontal root translation, this is still true for any choice of the root vertex of $\mathcal{B}_{h,h'}$ at height 0, as long as the choice of the root is independent of $\mathcal{B}_{h,\infty}$.

Definition 20. We say that $n > 0$ is a *regeneration time* if $H_i < H_n$ for every $i < n$, and $H_i \geq H_n$ for every $i \geq n$. We denote by $\tau^1 < \tau^2 < \dots$ the list of regeneration times in increasing order.

We also denote by T_∂ the first time at which the simple random walk X on \mathcal{H} hits $\partial\mathcal{H}$. The key of the proof of Theorem 3 will be to combine the two following results.

Proposition 21. We have $\mathbb{E}[\tau^1] < +\infty$. In particular, $\tau^1 < +\infty$ a.s..

Lemma 22. 1. Almost surely, $\tau^j < +\infty$ for every $j \geq 1$.

2. The path-decorated maps $\left(\mathcal{B}_{H_{\tau^j}, H_{\tau^{j+1}}}, (X_{\tau^j+i})_{0 \leq i \leq \tau^{j+1} - \tau^j}\right)$ for $j \geq 1$ are i.i.d. and have the same distribution as $\left(\mathcal{B}_{0, H_{\tau^1}}, (X_i)_{0 \leq i \leq \tau^1}\right)$ conditioned on $\{T_\partial = +\infty\}$.

3. In particular, the pairs $(\tau^{j+1} - \tau^j, H_{\tau^{j+1}} - H_{\tau^j})$ for $j \geq 1$ are i.i.d. and have the same distribution as (τ^1, H_{τ^1}) conditioned on $\{T_\partial = +\infty\}$.

Proposition 21 will be deduced from the results of Sections 4.2 and 4.3. On the other hand, Lemma 22 is the reason why regeneration times have been used to prove positive speed for many other models. The same property has already been observed and used in various contexts such as random walks in random environments [28], or biased random walks on Galton–Watson trees [25]. Although the proof is basically the same for our model, we write it formally in Appendix A.

Finally, we note that the finiteness of the times τ^i could be deduced directly from the results of Section 3, even in the case $\mu(0) > 0$. However, this is not sufficient to ensure positive speed.

We now explain how to conclude the proof of Theorem 3 from the last two results.

Proof of Theorem 3 given Proposition 21 and Lemma 22. By Lemma 13, it is enough to prove the result on \mathcal{H} . By item 3 of Lemma 22 and Proposition 21, we have

$$\mathbb{E}[\tau^2 - \tau^1] = \frac{\mathbb{E}[\tau^1 \mathbb{1}_{\forall n \geq 0, H_n \geq 0}]}{\mathbb{P}(\forall n \geq 0, H_n \geq 0)} < +\infty.$$

Moreover, $H_{\tau^2} - H_{\tau^1} \leq \tau^2 - \tau^1$, so $\mathbb{E}[H_{\tau^2} - H_{\tau^1}] < +\infty$ as well. By Lemma 22 and the law of large numbers, we have

$$\frac{\tau^j}{j} \xrightarrow{j \rightarrow +\infty} \mathbb{E}[\tau^2 - \tau^1] \quad \text{and} \quad \frac{H_{\tau^j}}{j} \xrightarrow{j \rightarrow +\infty} \mathbb{E}[H_{\tau^2} - H_{\tau^1}].$$

For every $n > \tau^1$, let $j(n)$ be the index such that $\tau^{j(n)} \leq n < \tau^{j(n)+1}$. Then we have $\frac{j(n)}{n} \rightarrow \mathbb{E}[\tau^2 - \tau^1]^{-1}$ a.s.. Moreover, we have $H_{\tau^{j(n)}} \leq H_n \leq H_{\tau^{j(n)+1}}$ by the definition of regeneration times, so $\frac{H_n}{j(n)} \rightarrow \mathbb{E}[H_{\tau^2} - H_{\tau^1}]$ a.s.. The result follows, with

$$v_\mu = \frac{\mathbb{E}[H_{\tau^2} - H_{\tau^1}]}{\mathbb{E}[\tau^2 - \tau^1]} > 0.$$

□

Proof of Proposition 21. We will actually show that τ^1 has a subpolynomial tail, i.e. for every $\beta > 0$, we have

$$\mathbb{P}(\tau^1 > n) = o(n^{-\beta}).$$

We first need to introduce a few notation. We define by induction stopping times τ_j and τ'_j for every $j \geq 1$:

- $\tau_1 = \inf\{n | H_n > 0\}$,
- $\tau'_j = \inf\{n \geq \tau_j | H_n < H_{\tau_j}\}$ for every $j \geq 1$,
- $\tau_{j+1} = \inf\left\{n > \tau'_j | H_n > \max\left(H_0, H_1, \dots, H_{\tau'_j}\right)\right\}$.

Let also J be the largest index such that $\tau_J < +\infty$. We claim that J is a geometric variable. Indeed, on the one hand, we know that $H_n \rightarrow +\infty$ when $n \rightarrow +\infty$, so almost surely, if $\tau'_j < +\infty$, then $\tau_{j+1} < +\infty$. On the other hand, if $\tau_j < +\infty$, let \mathcal{F}_{τ_j} be the σ -algebra generated by $\mathcal{B}_{0, H_{\tau_j}}$ and $(X_0, X_1, \dots, X_{\tau_j})$. Then the variable

$$\left(\mathcal{B}_{H_{\tau_j}, \infty}, (X_{\tau_j+i})_{0 \leq i \leq \tau'_j - \tau_j} \right)$$

is independent of \mathcal{F}_{τ_j} and has the same distribution as $\left(\mathcal{B}_{1, \infty}, (X_{\tau_1+i})_{0 \leq i \leq \tau'_1 - \tau_1} \right)$. In particular, if $\tau_j < +\infty$, we have

$$\mathbb{P}(\tau'_j < +\infty | \mathcal{F}_{\tau_j}) = \mathbb{P}(\tau'_1 < +\infty) = \mathbb{P}(T_\partial = +\infty),$$

so $\mathbb{P}(\tau_{j+1} < +\infty | \tau_j < +\infty)$ does not depend on j . This shows that J is a.s. finite and geometric. Note that $\tau^1 = \tau_J$. For any $n > 0$, we also denote by J_n the largest index j such that $\tau_j \leq n$. Finally, we recall that D_n is the greatest "descent" of X before time n .

In order to estimate the tail of τ^1 , we partition the event $\{\tau^1 > n\}$ into several "bad" events. Let $\delta > 0$ be small (we will actually only need $\delta < 1/3$). We have

$$\begin{aligned} \mathbb{P}(\tau^1 > n) &= \mathbb{P}(\tau^1 \neq \tau_{J_n}) \\ &\leq \mathbb{P}(J_n \geq n^\delta) + \mathbb{P}(D_n \geq n^\delta) + \mathbb{P}(J_n < n^\delta, D_n < n^\delta, \tau^1 \neq \tau_{J_n}). \end{aligned} \quad (11)$$

We now bound these terms one by one. First, we know that $J_n \leq J$, which is a geometric variable. Hence, the first term is at most $\exp(-cn^\delta)$ for some constant c , so it is $o(n^{-\beta})$ for any $\beta > 0$. Moreover, the second part of Proposition 19 shows that the second term is $o(n^{-\beta})$ as well.

Finally, we study the third term of (11). We first show that if $D_n < n^\delta$ and $J_n < n^\delta$, then $H_{\tau_{J_n}} < n^{2\delta}$ (this is a deterministic statement). If $D_n < n^\delta$, let $1 \leq j < J_n$. We have $\tau_{j+1} \leq n$, so $\tau'_j \leq n$ and $H_{\tau'_j} = H_{\tau_j} - 1$ by the definition of τ'_j . By the definitions of τ_{j+1} and of D_n , we have

$$H_{\tau_{j+1}} - H_{\tau_j} = 1 + \max_{[0, \tau'_j]} H - (H_{\tau'_j} + 1) \leq D_n < n^\delta.$$

By summing over j (and remembering $H_{\tau_1} = 1$), we obtain

$$H_{\tau_{J_n}} \leq 1 + n^\delta(J_n - 1) < 1 + n^\delta(n^\delta - 1) < n^{2\delta}.$$

Therefore, if the event in the third term of (11) occurs, we have $H_{\tau_{J_n}} < n^{2\delta}$ but $\tau'_{J_n} < +\infty$, so there is $k > n$ such that $H_k < n^{2\delta}$. On the other hand, if $\delta < 1/3$, we have

$$\mathbb{P}(\exists k > n, H_k \leq n^{2\delta}) \leq \mathbb{P}(\exists k > n, H_k \leq k^{1-\delta}) \leq \sum_{k>n} \mathbb{P}(H_k \leq k^{1-\delta}) = \sum_{k>n} o(k^{-(\beta+1)})$$

by the first point of Proposition 19. This proves that the third term of (11) decays superpolynomially, which concludes the proof. \square

5 Counterexamples and open questions

We finally discuss the necessity of the various assumptions made in the results of this paper, and we state a few conjectures. See Figure 3 for a quick summary.

Liouville property. We first note that if we do not require the strips (S_i) to be i.i.d., then Theorem 4 fails. Indeed, we start from $\mathcal{C}(\mathbf{T})$ and choose a ray γ_0 of \mathbf{T} . We then duplicate many times the horizontal edges to add a very strong lateral drift towards γ_0 . If we also duplicate the edges of γ_0 enough times, we can make sure that the simple random walk eventually stays on the path γ_0 . This yields a map of the form $\mathcal{M}(\mathbf{T}, (S_i))$ which has the intersection property, so it is Liouville.

Poisson boundary. The description of the Poisson boundary given by Theorem 2 cannot be true for any map of the form $\mathcal{M}(\mathbf{T}, (S_i))$, even if the strips (S_i) are i.i.d.. Indeed, it is possible to choose S_i such that the walk (X_n) has a positive probability to stay in S_i forever, and such that S_i itself has a non-trivial Poisson boundary. In this case, the Poisson boundary of $\mathcal{M}(\mathbf{T}, (S_i))$ is larger than $\widehat{\partial}\mathbf{T}$, and nonatomicity in Theorem 2 is false. On the other hand, we conjecture that if we furthermore assume that all the slices S_i are recurrent graphs (and i.i.d.), then $\widehat{\partial}\mathbf{T}$ is a realization of the Poisson boundary. As explained in Remark 12, our arguments cannot handle this general setting.

Positive speed. The positive speed is also false in general maps of the form $\mathcal{M}(\mathbf{T}, (S_i))$ if the strips S_i are too large and do not add vertical drift. For example, if they are equal to the half-planar regular triangular lattice, then the random walk will spend long periods in the same strip, where it has speed zero. On the other hand, we conjecture that the assumption $\mu(0) = 0$ is not necessary in Theorem 3.

As for Galton–Watson trees, another process of interest on the maps $\mathcal{C}(T)$ is the λ -biased random walk X^λ . If a vertex x has $c(x)$ children and $X_n^\lambda = x$, then X_{n+1}^λ is equal to y with probability $\frac{1}{c(x)+3\lambda}$ for every child y of x , and to z with probability $\frac{\lambda}{c(x)+3\lambda}$ if z is the parent or one of the two neighbours of x .

If $\lambda > 1$, we expect that, whether $\mu(0) = 0$ or not, the process behaves in the same way as on trees [25]: the walk is recurrent for $\lambda > \lambda_c$ (as easily shown by the Nash–Williams criterion) and should have positive speed for $\lambda < \lambda_c$, where $\lambda_c = \sum i\mu(i)$. If $\lambda < 1$ and $\mu(0) = 0$, it is easy to see that the speed is positive on $\mathcal{C}(T)$ since the drift at every vertex is positive. For $\lambda < 1$ and $\mu(0) > 0$, the λ -biased walk on T has speed zero for λ small enough ($\lambda \leq f'(q)$, where q is the extinction probability of T and f the generating function of μ). We believe that this regime disappears on causal maps, and that the λ -biased walk on $\mathcal{C}(T)$ has positive speed for every $\lambda < 1$.

Other properties of the simple random walk (for $\mu(0) = 0$). As shown by Theorem 2, the harmonic measure of $\mathcal{C}(T)$ on $\widehat{\partial}T$ is a.s. nonatomic and has full support. It would be interesting to investigate finer properties of this measure, as it has been done for Galton–Watson trees [24, 22]. We believe that as for Galton–Watson trees, the harmonic measure is not absolutely continuous with respect to the mass measure, and should satisfy a dimension drop.

Another quantity of interest related to the simple random walk is the heat kernel decay, i.e. the probability of returning to the root at time n . Perhaps surprisingly, the annealed and quenched heat kernels might have different behaviours: if $\mu(1) > 0$, the possibility that T does not branch during the first $n^{1/3}$ steps gives an annealed lower bound of order $e^{-n^{1/3}}$. On the other hand, the worst possible traps after the first branching points seem to be large portions of square lattice, which yield a quenched lower bound of order $e^{-n^{1/2}}$. Our argument for quasi-positive

speed could be adapted to prove that the heat kernel decays quicker than any polynomial, which seems far from optimal. On the other hand, a natural first step to show that the lower bounds are tight would be to prove anchored expansion for $\mathcal{C}(T)$. However, this property does not seem well suited to the study of causal maps since connected subsets of $\mathcal{C}(T)$ can be quite nasty.

Other random processes. Finally, other random processes such as percolation on $\mathcal{C}(T)$ might be investigated. We expect that we should have $p_c < p_u$, i.e. there is a regime where infinitely many infinite components coexist, as it is generally conjectured for graphs with a hyperbolic behaviour (like for example nonamenable transitive graphs [10]). We note that oriented percolation is studied in a work in progress of David Marchand.

More generally, for unimodular, planar graphs, other notions of hyperbolicity (including $p_c < p_u$) have been studied in [4] and proved to be equivalent to each other. It might be interesting to study the relation with our setting: if it is true that any hyperbolic (in the sense of [4]) unimodular map contains a supercritical Galton–Watson tree, then our results of Section 2 apply. On the other hand, it is clear that every unimodular planar map containing a supercritical Galton–Watson tree is hyperbolic in the sense of [4].

A The regeneration structure

The goal of this appendix is to prove Lemma 22. We recall that $\partial\mathcal{H}$ is the set of vertices at height -1 , and that T_∂ is the first time at which X hits $\partial\mathcal{H}$. We will first prove the following intermediate result.

Lemma 23. 1. We have $\tau^1 < +\infty$ a.s..

2. The path-decorated map $(\mathcal{B}_{H_{\tau^1}, \infty}, (X_{\tau^1+i})_{i \geq 0})$ is independent of $(\mathcal{B}_{0, H_{\tau^1}}, (X_i)_{0 \leq i \leq \tau^1})$ and has the same distribution as $(\mathcal{H}, (X_i)_{i \geq 0})$ conditioned on the event $\{T_\partial = +\infty\}$.

Note that the first point follows from Proposition 21, so we only need to focus on the second point.

Proof of Lemma 23. We first note that, by Proposition 19, we have $H_n \rightarrow +\infty$ a.s., so the conditioning on $\{T_\partial = +\infty\}$ is non-degenerate.

For every $h \geq 0$, let $T_h = \min\{n \geq 0 | H_n = h\}$, and let $T'_h = \min\{n \geq T_h | H_n < h\}$. By Proposition 19, we have $T_h < +\infty$ a.s.. We also know that the rooted map $(\mathcal{B}_{h, \infty}, X_{T_h})$ is independent of $(\mathcal{B}_{0, h}, (X_i)_{0 \leq i \leq T_h})$ and has the same distribution as \mathcal{H} . Therefore, the path-decorated map $(\mathcal{B}_{h, \infty}, (X_{T_h+i})_{0 \leq i \leq T'_h - T_h})$ is independent of $(\mathcal{B}_{0, h}, (X_i)_{0 \leq i \leq T_h})$ and has the same distribution as $(\mathcal{H}, (X_i)_{0 \leq i \leq T_\partial})$.

It follows that, for any two measurable sets A and B of path-decorated maps, we have

$$\begin{aligned} & \mathbb{P} \left((\mathcal{B}_{0, H_{\tau^1}}, (X_i)_{0 \leq i \leq \tau^1}) \in A \text{ and } (\mathcal{B}_{H_{\tau^1}, \infty}, (X_{\tau^1+i})_{i \geq 0}) \in B \right) \\ &= \sum_{h \geq 0} \mathbb{P} \left((\mathcal{B}_{0, h}, (X_i)_{0 \leq i \leq T_h}) \in A \text{ and } (\mathcal{B}_{h, \infty}, (X_{T_h+i})_{i \geq 0}) \in B \text{ and } \tau^1 = h \right) \\ &= \sum_{h \geq 0} \mathbb{P} \left((\mathcal{B}_{0, h}, (X_i)_{0 \leq i \leq T_h}) \in A \text{ and } \forall i < h, T'_i \leq T_h \right. \\ & \qquad \qquad \qquad \left. \text{and } (\mathcal{B}_{h, \infty}, (X_{T_h+i})_{i \geq 0}) \in B \text{ and } T'_h = +\infty \right), \end{aligned}$$

by noting that H_{τ^1} is the smallest height i such that $T'_i = +\infty$. Note that the event $\{\forall i < h, T'_i \leq T_h\}$ is a measurable function of $(\mathcal{B}_{0,h}, (X_i)_{0 \leq i \leq T_h})$, and the event $\{T'_h = +\infty\}$ is a measurable function of $(\mathcal{B}_{h,\infty}, (X_{T_h+i})_{i \geq 0})$. Hence, by the independence and the distribution of $(\mathcal{B}_{h,\infty}, (X_{T_h+i})_{i \geq 0})$ found above, we have

$$\begin{aligned} & \mathbb{P} \left((\mathcal{B}_{0,H_{\tau^1}}, (X_i)_{0 \leq i \leq \tau^1}) \in A \text{ and } (\mathcal{B}_{H_{\tau^1},\infty}, (X_{\tau^1+i})_{i \geq 0}) \in B \right) \\ &= \sum_{h \geq 0} \mathbb{P} \left((\mathcal{B}_{0,h}, (X_i)_{0 \leq i \leq T_h}) \in A \text{ and } \forall i < h, T'_i \leq T_h \right) \mathbb{P} \left((\mathcal{B}_{h,\infty}, (X_i)_{i \geq 0}) \in B \text{ and } T_\partial = +\infty \right) \\ &= \mathbb{P} \left((\mathcal{B}_{0,\infty}, (X_i)_{i \geq 0}) \in B \mid T_\partial = +\infty \right) f(A), \end{aligned}$$

where $f(A)$ is a function of A . Therefore, the path-decorated maps $(\mathcal{B}_{0,H_{\tau^1}}, (X_i)_{0 \leq i \leq \tau^1})$ and $(\mathcal{B}_{H_{\tau^1},\infty}, (X_{\tau^1+i})_{i \geq 0})$ are independent and, by taking $A = \Omega$, we obtain that the distribution of the second is a multiple of the distribution of $(\mathcal{H}, (X_i)_{i \geq 0})$ conditioned on the event $\{T_\partial = +\infty\}$. Since both are probability measures, they coincide. \square

Proof of Lemma 22. We define the shift operator θ as follows:

$$(\mathcal{H}, (X_i)_{i \geq 0}) \circ \theta = (\mathcal{B}_{H_{\tau^1},\infty}(\mathcal{H}), (X_{\tau^1+i})_{i \geq 0}).$$

We first notice that Lemma 23 remains true if we consider (\mathcal{H}, X) under the measure $\mathbb{P}(\cdot \mid T_\partial = +\infty)$ instead of \mathbb{P} . Indeed, conditioning on an event of positive probability does not change the fact that $\tau^1 < +\infty$ a.s.. Moreover, the event $\{T_\partial = +\infty\}$ only depends on $(\mathcal{B}_{0,H_{\tau^1}}, (X_i)_{0 \leq i \leq \tau^1})$ and not on $(\mathcal{B}_{H_{\tau^1},\infty}, (X_{\tau^1+i})_{i \geq 0})$, so conditioning on this event affects neither the independence of these two path-decorated maps, nor the distribution of the second.

But by Lemma 23, the map $(\mathcal{H}, (X_i)_{i \geq 0}) \circ \theta$ has the same distribution as $(\mathcal{H}, (X_i)_{i \geq 0})$ under $\mathbb{P}(\cdot \mid T_\partial = +\infty)$, so Lemma 23 applies after composition by θ . In particular, we have $\tau^1 \circ \theta < +\infty$ a.s., i.e. $\tau^2 < +\infty$ a.s.. Moreover, the two following path-decorated maps are independent:

- $(\mathcal{B}_{0,H_{\tau^1}}, (X_i)_{0 \leq i \leq \tau^1}) \circ \theta = (\mathcal{B}_{H_{\tau^1},H_{\tau^2}}, (X_{\tau^1+i})_{0 \leq i \leq \tau^2 - \tau^1}),$
- $(\mathcal{B}_{H_{\tau^1},\infty}, (X_{\tau^1+i})_{i \geq 0}) \circ \theta = (\mathcal{B}_{H_{\tau^2},\infty}, (X_{\tau^2+i})_{i \geq 0}),$

and the second one has the same distribution as $(\mathcal{H}, (X_i)_{i \geq 0})$ under $\mathbb{P}(\cdot \mid T_\partial = +\infty)$. From here, an easy induction on j shows that for any $j \geq 1$, we have $\tau^j < +\infty$ and the path-decorated map $(\mathcal{B}_{H_{\tau^j},H_{\tau^{j+1}}}, (X_{\tau^j+i})_{0 \leq i \leq \tau^{j+1} - \tau^j})$ has indeed the right distribution and is independent of $(\mathcal{B}_{H_{\tau^{j+1}},\infty}, (X_{\tau^{j+1}+i})_{i \geq 0})$. This proves Lemma 22 (the third item is a direct consequence of the first two). \square

References

- [1] E. Aïdékon. Speed of the biased random walk on a Galton–Watson tree. *Probability Theory and Related Fields*, 159(3):597–617, Aug 2014.
- [2] J. Ambjørn and R. Loll. Non-perturbative Lorentzian quantum gravity, causality and topology change. *Nuclear Physics B*, 536(1):407 – 434, 1998.
- [3] O. Angel, T. Hutchcroft, A. Nachmias, and G. Ray. Unimodular hyperbolic triangulations: circle packing and random walk. *Inventiones mathematicae*, 206(1):229–268, Oct 2016.

- [4] O. Angel, T. Hutchcroft, A. Nachmias, and G. Ray. Hyperbolic and parabolic unimodular random maps. *Geometric and Functional Analysis*, Jun 2018.
- [5] O. Angel, A. Nachmias, and G. Ray. Random walks on stochastic hyperbolic half planar triangulations. *Random Structures and Algorithms*, 49(2):213–234, 2016.
- [6] O. Angel and O. Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.*, 241(2-3):191–213, 2003.
- [7] I. Benjamini. Gaussian free field on hyperbolic lattices. *Lecture Notes in Mathematics*, 2116:39–45, 08 2014.
- [8] I. Benjamini and N. Curien. Ergodic theory on stationary random graphs. *Electron. J. Probab.*, 17(93):1–20, 2012.
- [9] I. Benjamini, E. Paquette, and J. Pfeffer. Anchored expansion, speed and the Poisson–Voronoi tessellation in symmetric spaces. *Ann. Probab.*, 46(4):1917–1956, 07 2018.
- [10] I. Benjamini and O. Schramm. Percolation beyond Z^d , many questions and a few answers. *Electron. Commun. Probab.*, 1:71–82, 1996.
- [11] I. Benjamini and R. Tessera. First passage percolation on a hyperbolic graph admits bi-infinite geodesics. *Electron. Commun. Probab.*, 22:8 pp., 2017.
- [12] T. Budzinski. Infinite geodesics in hyperbolic random triangulations. *arXiv:1804.07711*, 2018.
- [13] N. Curien. Planar stochastic hyperbolic triangulations. *Probability Theory and Related Fields*, 165(3):509–540, 2016.
- [14] N. Curien. Scaling limits of stable causal maps. *to appear*, 2018+.
- [15] N. Curien, T. Hutchcroft, and A. Nachmias. Geometric and spectral properties of causal maps. *arXiv:1710.03137*, 2017.
- [16] N. Curien and J.-F. Le Gall. First-passage percolation and local modifications of distances in random triangulations. *Annales Scientifiques de l'École Normale Supérieure*, to appear.
- [17] P. G. Doyle and J. L. Snell. *Random walks and electric networks*, volume 22 of *Carus Mathematical Monographs*. Mathematical Association of America, Washington, DC, 1984.
- [18] B. Durhuus, T. Jonsson, and J. Wheeler. On the spectral dimension of causal triangulations. *Journal of Statistical Physics*, 139:859–881, 2010.
- [19] T. Hutchcroft and Y. Peres. Boundaries of planar graphs: a unified approach. *Electron. J. Probab.*, 22, 2017.
- [20] V. A. Kaimanovich and A. M. Vershik. Random walks on discrete groups: boundary and entropy. *Ann. Probab.*, 11(3):457–490, 1983.
- [21] M. Krikun. A uniformly distributed infinite planar triangulation and a related branching process. *Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI)*, 307(Teor. Predst. Din. Sist. Komb. i Algoritm. Metody. 10):141–174, 282–283, 2004.

- [22] S. Lin. Harmonic measure for biased random walk in a supercritical Galton-Watson tree. *arXiv:1707.01811*, July 2017.
- [23] R. Lyons. Random walks, capacity and percolation on trees. *The Annals of Probability*, 20(4):2043–2088, 1992.
- [24] R. Lyons, R. Pemantle, and Y. Peres. Ergodic theory on Galton-Watson trees: speed of random walk and dimension of harmonic measure. *Ergodic Theory Dynam. Systems*, 15(3):593–619, 1995.
- [25] R. Lyons, R. Pemantle, and Y. Peres. Biased random walks on Galton-Watson trees. *Probab. Theory Related Fields*, 106(2):249–264, 1996.
- [26] R. Lyons and Y. Peres. *Probability on Trees and Networks*. Cambridge University Press, 2017.
- [27] R. Lyons and G. White. A Stationary Planar Random Graph with Singular Stationary Dual: Dyadic Lattice Graphs. *arXiv:1807.06165*, Jul 2018.
- [28] A.-S. Sznitman and M. Zerner. A law of large numbers for random walks in random environment. *Ann. Probab.*, 27(4):1851–1869, 10 1999.
- [29] B. Virág. Anchored expansion and random walk. *Geom. Funct. Anal.*, 10(6):1588–1605, 2000.