Annales de l'Institut Henri Poincaré - Probabilités et Statistiques
2016, Vol. 52, No. 1, 162–172
DOI: 10.1214/14-AIHP636

ANNALES DE L'INSTITUT HENRI POINCARÉ
PROBABILITÉS ET STATISTIQUES
www.imstat.org/aihp

The Brownian web is a two-dimensional black noise

Tom Ellis[1] and Ohad N. Feldheim[2]

School of Mathematics, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel.
E-mail: ohad_f@netvision.net.il

Received 4 June 2013; revised 2 June 2014; accepted 25 July 2014

Abstract. The Brownian web is a random variable consisting of a Brownian motion starting from each space–time point on the plane. These are independent until they hit each other, at which point they coalesce. Tsirelson mentions this model in (*Scaling Limit, Noise, Stability* (2004) Springer), along with planar percolation, in suggesting the existence of a two-dimensional black noise. A two-dimensional noise is, roughly speaking, a random object on the plane whose distribution is translation invariant and whose behavior on disjoint subsets is independent. Black means sensitive to the resampling of sets of arbitrarily small total area.

Tsirelson implicitly asks: "Is the Brownian web a two-dimensional black noise?." We give a positive answer to this question, providing the second known example of such after the scaling limit of critical planar percolation.

Résumé. La toile brownienne est une collection de mouvements browniens issus de tout point du plan, indépendants jusqu' à l'instant où ils se rencontrent, à la suite de quoi ils coalescent. Tsirelson mentionne ce modèle dans (*Scaling Limit, Noise, Stability* (2004) Springer), avec la percolation planaire, comme modèles possibles de bruit noir bidimensionnel. Un bruit bidimensionnel est, en gros, un objet aléatoire dans le plan dont la distribution est invariante par translation et dont le comportement sur des ensembles disjoints est indépendant. Noir signifie sensible au re-échantillonage d'ensembles dont l'aire totale est arbitrairement petite.

Tsirelson demande implicitement : « La toile brownienne est-elle un bruit noir bidimensionnel ? ». Nous répondons positivement à cette question, en donnant le deuxième exemple connu d'un tel bruit après la limite d'échelle de la percolation critique planaire.

MSC: 60G20

Keywords: Brownian web; Black noise

1. Introduction

In this paper we study a stochastic object called the *Brownian web*. We research this object in the context of the theory of classical and non-classical noises, developed by Boris Tsirelson (see [13] for a survey). Our main result is that, in the terminology of this theory, the Brownian web is a two-dimensional noise, and we then demonstrate that this noise is black. Roughly speaking, the Brownian web is a random variable which assigns to every space–time point in $\mathbb{R} \times \mathbb{R}$ a standard Brownian motion starting at that point. The motions in each finite subcollection are independent until the first time that one hits another and from thereon those two coalesce, continuing together. This object was originally studied more then twenty-five years ago by Arratia [1], motivated by a study of the asymptotics of one-dimensional voter models, and later by Tóth and Werner [9], motivated by the problem of constructing continuum "self-repelling motions," by Fontes, Isopi, Newman and Ravishankar [3], motivated by its relevance to "aging" in the statistical physics of one-dimensional coarsening, and by Norris and Turner [6] regarding a scaling limit of a two-dimensional aggregation process. A rigorous notion of the Brownian web in our context can be found in [11] for the

[1]This work was supported in part by a Wyndham Deedes Memorial Travel Scholarship from The Anglo-Israel Association.
[2]Research supported by an ERC advanced grant.

case of coalescing Brownian motions on a circle. The above also provide their own constructions of the Brownian web.

The Brownian web functions as an important example in the theory of classical and non-classical noises. In this theory a noise is a probability space equipped with a collection of sub-σ-fields indexed by the open rectangles (possibly infinite) of $\mathbb{R}^d$. (Throughout, all σ-fields will be assumed to contain all sets of probability zero.) The sub-σ-field associated to a rectangle is intended to represent the behavior of a stochastic object within that rectangle. The σ-fields must satisfy the following three properties:

- the σ-fields associated to disjoint rectangles of $\mathbb{R}^d$ are independent,
- translations on $\mathbb{R}^d$ act in a way that preserves the probability measure,
- the σ-field associated to a rectangle is generated by the two σ-fields associated to two smaller rectangles which partition it.

Two natural examples of noises are the Gaussian white noise and the Poisson noise. These noises are called classical, or white, meaning that resampling of a small portion of $\mathbb{R}^d$ doesn't change the observables of the process very much.

The foundational result of Tsirelson and Vershik [10] showed that there exist non-classical noises. Indeed they showed the existence of non-classical noises that are as far from white as could be, and these are called black. The defining property of a black noise is that all its observables are sensitive, i.e. for any particular observable, resampling a small scattered portion of the noise renders that observable nearly independent of its original value. (For a thorough discussion of black and white, classical and non-classical noises see [13].)

In [12], Tsirelson gives several examples of one-dimensional noises, and shows (Theorem 7c2) that the Brownian web, when considered as a time-indexed random process, is a (one-dimensional) black noise. He suggests the Brownian web and critical planar percolation as possible sources for the discovery of two-dimensional black noises.

Following a result by Benjamini, Kalai and Schramm about the sensitivity of local events in [2], and using a result in [8], Tsirelson shows that if a scaling limit for critical planar percolation were a noise, then this noise would necessarily be black. Such a limit was later demonstrated by Schramm and Smirnov [7], giving the first example of a two-dimensional black noise. Here we give only the second known example of a two-dimensional black noise through the following theorem:

Theorem 1.1. *The Brownian web is a two-dimensional black noise.*

Of the three properties required for a process to be a noise, the first two, i.e. translation invariance and independence on disjoint domains, hold trivially for the Brownian web. Furthermore, once we have shown that the Brownian web is a two-dimensional noise it will follow that it is a two-dimensional *black* noise, through a general argument.

The main difficulty in proving Theorem 1.1 is to show that the σ-field associated to any rectangle is generated by the two σ-fields associated to any two rectangles that partition it. The key step towards this result is to show the special case when the large rectangle is the whole plane, and the smaller rectangles that partition it are the upper and lower half-planes.

Theorem 1.2. *In the Brownian web, the σ-field associated to the whole plane is generated by that associated to the upper half-plane and that associated to the lower half-plane.*

The rest of the paper goes as follows: in Section 2 we define the Brownian web formally; in Section 3 we restate Theorem 1.2 as Theorem 3.1; we then reduce this theorem to a convergence result for an auxiliary process which we prove in Section 4. In Section 5 we extend Theorem 1.2 to hold for the σ-fields associated to horizontal strips as well; in Section 6 we extend further to all rectangles. In addition we define noises properly and conclude by proving Theorem 1.1. Section 7 is devoted to remarks and open problems.

2. Definition of the Brownian web

The Brownian web is the continuum scaling limit of a system of independent-coalescing random walks (see [11]). Constructing the continuum version raises several technical difficulties addressed in [3,6,9]. Different authors have introduced different state spaces in order to the give the Brownian web a unique distribution on each of those spaces (see Theorem 2.1 of [3], and Theorem 5.1 of [6]). Nonetheless, all the definitions share the following defining property.

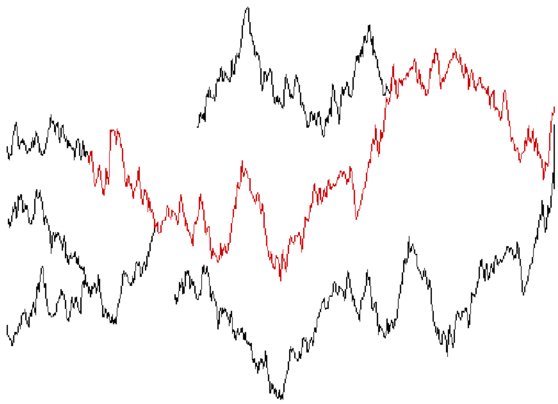

Fig. 1. Some trajectories of the Brownian web. A particular trajectory is marked.

Definition 2.1. *Denote $\mathcal{S} = \{(s,t) \in \mathbb{R}^2 : s \le t\}$.*

- *A Brownian web on a probability space Ω induces a (jointly) measurable mapping $\phi : \Omega \times \mathcal{S} \times \mathbb{R} \to \mathbb{R}$, $(\omega, (s,t), x) \mapsto \phi_{s,t}(x)$ (suppressing ω in the notation).*
- *On every finite collection of starting points $(s_1, x_1), (s_2, x_2), \dots, (s_n, x_n)$, the collection of processes $\phi_{s_1,\cdot}(x_1)$, $\phi_{s_2,\cdot}(x_2), \dots, \phi_{s_n,\cdot}(x_n)$ forms a system of n independent-coalescing Brownian motions.*
- *There exists a countable subset $S \subseteq \mathbb{R}^2$ such that the collection $\{\phi_{s,t}(x) : s, t \in S\}$ generates the entire σ-field of the Brownian web.*

A system of n independent-coalescing Brownian motions is a finite collection of stochastic processes $(X^1, X^2, \dots, X^n)$ such that each X^i starts at some point x_i at some time s_i, and $(X^1, X^2, \dots, X^n)$ are independent until the first time T at which $X^i(T) = X^j(T)$ for some $i \ne j$. From this time onwards $X^i(T)$ and $X^j(T)$ coalesce and continue with the rest of the X^k (for $k \ne i, j$) as a system of $n-1$ independent-coalescing Brownian motions.

To observe that the last property is indeed a property of the Brownian web, see for example equation (2.6) in [9]. In this equation we can see that Λ (which is the same as our ϕ) can be recovered from the countable collection $F_n = \Lambda_{(s_n,t_n)}$. There $\{(s_n,t_n) : n \in \mathbb{N}\}$ is some ordering of $D \times D$ where $D \subseteq \mathbb{R}^2$ is the collection of diadic rationals.

For the equivalent construction in [3], item (2) of Theorem 2.1 guarantees that the third property in our definition holds.

Several trajectories of a Brownian web can be seen in Figure 1.

3. Recovering the web from half-planes

In this section we introduce three σ-fields generated by the Brownian web and use them to restate Theorem 1.2 as Theorem 3.1.

Write $X_{(s,x)}$ for the random trajectory $t \mapsto \phi_{s,t}(x)$. The value of this process at time t is $X_{(s,x),t}$ and we will simplify the notation to X_t when s and x are clear from the context. Write $\mathcal{X}$ for the collection of all trajectories, that is $\{X_{(s,x)} : (s,x) \in \mathbb{R}^2\}$. Let $\mathcal{F}$ be the σ-field generated by $\mathcal{X}$. By the third property in Definition 2.1, $\mathcal{F}$ is also generated by $\{X_{(s,x)} : (s,x) \in S\}$ where $S \subseteq \mathbb{R}^2$ is some countable subset.

We further introduce $\mathcal{F}_+$ and $\mathcal{F}_-$, the σ-fields generated by the web on the upper and lower half-planes respectively. Formally,

Definition. *For any path f we write $\mathcal{R}_+(f)$ for f stopped at the first time it is outside the upper half-plane. $\mathcal{R}_+(X_{(s,x)})$ is therefore the trajectory of the web ϕ started at the point (s,x) and stopped at the first time it is outside the upper half-plane (if (s,x) is outside the upper half-plane, f is stopped immediately). We define $\mathcal{F}_+$ to be the σ-field generated by the collection of paths $\{\mathcal{R}_+(X_{(s,x)}) : (s,x) \in \mathbb{R}^2\}$, and define $\mathcal{F}_-$ analogously.*

Note that $\mathcal{F}_+$ and $\mathcal{F}_-$ are independent by the definition of a system of n independent-coalescing Brownian motions.

For σ-fields $\mathcal{F}_a$, $\mathcal{F}_b$, $\mathcal{F}_c$ we write $\mathcal{F}_a = \mathcal{F}_b \otimes \mathcal{F}_c$ when $\mathcal{F}_a$ is generated by $\mathcal{F}_b$ and $\mathcal{F}_c$ (up to sets of measure 0), and $\mathcal{F}_b$ and $\mathcal{F}_c$ are independent. We are now ready to restate Theorem 1.2.

Theorem 3.1. $\mathcal{F} = \mathcal{F}_- \otimes \mathcal{F}_+$.

The difficulty is to show that $\mathcal{F} \subseteq \mathcal{F}_- \otimes \mathcal{F}_+$, because the reverse inclusion follows directly from the definitions. To avoid having to state results as "for all $X \in \mathcal{X}$ ...," here and for the rest of the paper we consider the trajectory starting at an arbitrary start point (s_0, x_0), and write $X_t = X_{(s_0,x_0),t}$.

$\mathcal{F}$ is generated by such processes, so using the third property in Definition 2.1 our theorem reduces to the following lemma.

Lemma 3.2. *X is $\mathcal{F}_- \otimes \mathcal{F}_+$-measurable.*

To prove this we must show that X can be recovered using trajectories starting in the upper half-plane which stop when they hit 0, and trajectories starting on the lower half-plane which stop when they hit 0.

To do so, we might have liked to recover X (starting on, say, the upper half-plane) by following it until it hits 0, then by following its continuation within the lower half-plane. However, this is impossible since when trajectories of Brownian motion cross 0 they do so infinitely often within a finite period of time.

We introduce in Section 3.1 a process X^ε, for each $\varepsilon > 0$, which is an approximation of X in the sense that X^ε converges to X (in probability) as $\varepsilon \to 0$. We use X^ε to show that X is is $F_- \otimes F_+$-measurable.

Observe that here and in the rest of the paper we treat s_0 and x_0 as fixed and so rates of convergence may very well depend on their value.

3.1. *The perturbed process*

From our arbitrary X we now construct X^ε, a "perturbed" version of X, which depends also on ψ_t, some Brownian motion independent of ϕ (measurable with respect to $\mathcal{G}$, say, where $\mathcal{G}$ is independent of $\mathcal{F}$).

In the definition of X^ε we give the word "follows" two distinct meanings. We say X^ε follows ϕ on $[s,u]$ if $X^\varepsilon_t = \phi_{s,t}(X^\varepsilon_s)$ for $t \in [s,u]$. That is if the trajectory of X^ε follows the trajectory of the web starting from point X^ε_s at time s and up to time u. We say X^ε follows ψ on $[s,u]$ if $X^\varepsilon_t = X^\varepsilon_s + \psi_t - \psi_s$ for $t \in [s,u]$.

Definition 3.3. *The perturbed process X^ε starts at the same time and position as X, i.e. at (s_0, x_0), and alternates between two states. In state S_ϕ it follows ϕ while in state S_ψ it follows ψ. The process starts in state S_ϕ and the transition from state S_ϕ to state S_ψ occurs when X^ε hits 0, while the transition from state S_ψ to state S_ϕ occurs when X^ε hits $\pm\varepsilon$. See Figure 2 for an illustration of sample paths of X and X^ε.*

The following lemma specifies in what sense the perturbed process is an approximation of the trajectory of the web. Here and in the rest of the paper the convergence is uniform on compacts in probability.

Lemma 3.4. *$X^\varepsilon \xrightarrow{\mathbb{P}} X$ as $\varepsilon \to 0$.*

That is to say, for all $\delta > 0$, $u > s_0$

$$\mathbb{P}\big(\big|X(t) - X^\varepsilon(t)\big| > \delta \text{ for some } t \in [s_0, u]\big) \to 0 \quad \text{as } \varepsilon \to 0.$$

Lemma 3.2 reduces to Lemma 3.4.

Proof of reduction. X^ε is $\mathcal{F}_- \otimes \mathcal{F}_+ \otimes \mathcal{G}$-measurable, so we use Lemma 3.4 to conclude that X is also $\mathcal{F}_- \otimes \mathcal{F}_+ \otimes \mathcal{G}$-measurable. However, since X is actually independent of $\mathcal{G}$ we can use a basic result on tensor products of Hilbert spaces (for example Equation (2c8) of [14]) to conclude that X is in fact $\mathcal{F}_- \otimes \mathcal{F}_+$-measurable. □

We devote the following section to the proof of Lemma 3.4.

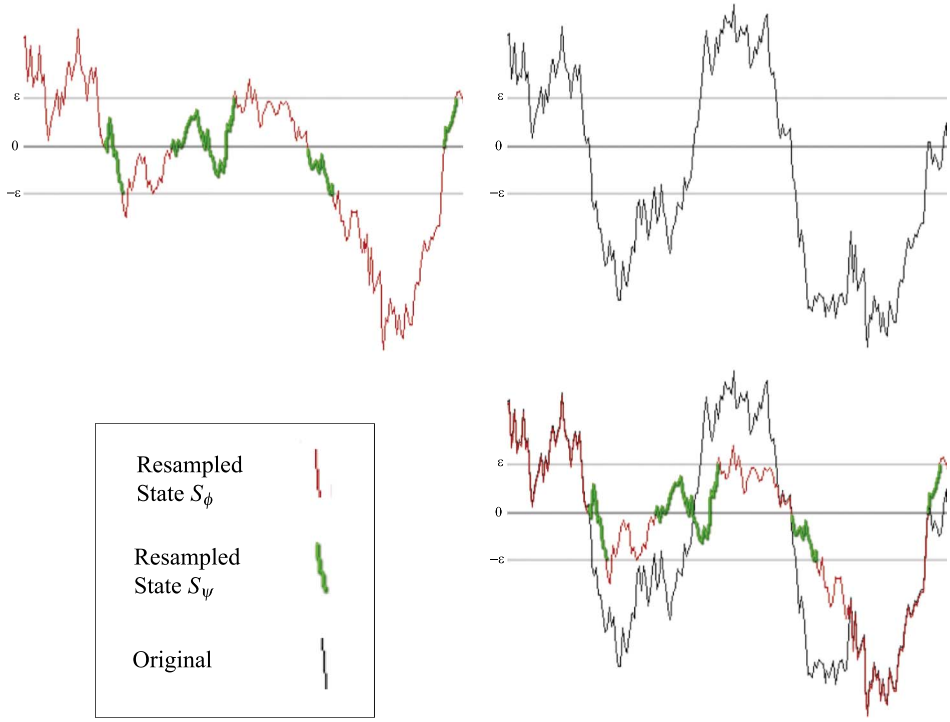

Fig. 2. A sample of a web trajectory and the corresponding perturbed processes. The top left image depicts the perturbed process, with state S_ψ in bold. The top right image depicts the original web trajectory. The bottom right image illustrates both processes together, showing the segments where they coalesce.

4. Convergence of the perturbed process

In this section we prove that $X^\varepsilon \xrightarrow{\mathbb{P}} X$ as $\varepsilon \to 0$. This statement depends only on the joint distribution of X and X^ε. We therefore define $Y(\varepsilon) = Y = (X^\varepsilon, X)$ (generally suppressing the ε dependence in the notation). Let us describe the distribution of Y as a two-dimensional random process.

We classify the behavior of the process into three states according to the behavior of X^ε with respect to X.

- If X^ε is in state S_ϕ and is coalesced with X we say Y is in state S_ϕ^{1D}.
- If X^ε is in state S_ϕ and *is not* coalesced with X we say Y is in state S_ϕ^{2D}.
- If X^ε is in state S_ψ we say Y is in state S_ψ^{2D}.

Y starts at point (x_0, x_0) at time s_0 in state S_ϕ^{1D}. From S_ϕ^{1D}, Y can only transition to S_ψ^{2D}. This transition occurs when Y hits the origin, as the coalesced X and X^ε will continue together until they leave their current half-plane. From S_ψ^{2D}, Y can only transition to S_ϕ^{2D}. This transition occurs when X^ε leaves the $(-\varepsilon, \varepsilon)$ interval (i.e. Y hits either of the $x = \pm\varepsilon$ lines). From S_ϕ^{2D}, Y can either transition to S_ϕ^{1D} if X and X^ε coalesce (i.e. Y hits the line $x = y$) or transition to S_ψ^{2D} if X^ε hits 0 (i.e. Y hits the $x = 0$ line). States and possible transitions of Y are summarized in Figure 3.

The form of the labels given to the states is justified by the following.

Observation. *In S_ϕ^{1D}, Y follows the law of a (time scaled) one-dimensional Brownian motion on the line $x = y$. In S_ϕ^{2D} and S_ψ^{2D}, Y follows the law of a two-dimensional Brownian motion.*

Additionally observe that by the scale-invariance of Brownian motion, the distribution of the sample paths of $Y(\varepsilon)/\varepsilon$ is independent of ε (modulo time scaling).

	State	Illustration	Law	Next	Trans. cond.
X^ε	S_ϕ^{1D}	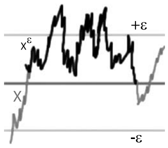	Equal	S_ψ^{2D}	Hits 0
	S_ϕ^{2D}	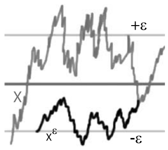	Indep.	S_ψ^{2D} S_ϕ^{1D}	Hits 0 Hits X
	S_ψ^{2D}	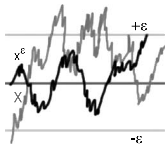	Indep.	S_ϕ^{2D}	Hits $\pm\varepsilon$
$Y = (x, y)$ $x = X^\varepsilon$ $y = X$	S_ϕ^{1D}	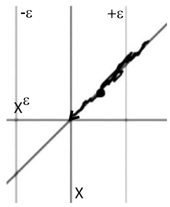	Equal	S_ψ^{2D}	$x = y = 0$
	S_ϕ^{2D}	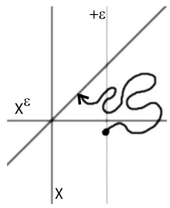	Indep.	S_ψ^{2D} S_ϕ^{1D}	$x = 0$ $x = y$
	S_ψ^{2D}	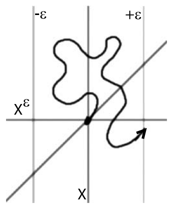	Indep.	S_ϕ^{2D}	$x = \pm\varepsilon$

Fig. 3. Illustrated states and transitions of Y.

Define $A_\delta = \{(x, y) : |x - y| = \delta\}$. To prove Lemma 3.4 we use the following property of Y:

Lemma 4.1. *For given $P > 0$, $\delta > 0$ the probability that Y hits A_δ before it hits (P, P) is* $o(1)$ *as* $\varepsilon \to 0$.

Recall that Y starts in state S_ϕ^{1D}, with X equal to X^ε. (If $x_0 = 0$ then Y immediately transitions to state S_ψ^{2D}.) We delay the proof of Lemma 4.1 to Section 4.1.

Proof of Lemma 3.4. Lemma 3.4 is equivalent to: for all $\delta > 0$, $u > s_0$

$$\mathbb{P}\big(Y_t \in \{(x, y) : |x - y| < \delta\} \text{ for all } t \in [s_0, u]\big) \to 1 \quad \text{as } \varepsilon \to 0.$$

The above statement can be rephrased as

For all $\delta > 0$, $u > s_0$, the probability that before time u, Y has hit A_δ is $o(1)$ as $\varepsilon \to 0$.

We prove this as follows. For any $\eta > 0$, choose P so that the probability that standard Brownian motion travels from 0 to P in time less than $u - s_0$ is less than η. Apply Lemma 4.1 to choose ε_0 such that, for all $\varepsilon < \varepsilon_0$, the probability that $Y(\varepsilon)$ hits A_δ before (P, P) is less than η. Then the probability that $Y(\varepsilon)$ hits A_δ before (P, P) or takes less time than $u - s_0$ to reach (P, P) is less than 2η. Thus the probability that $Y(\varepsilon)$ hits A_δ before time u is less than 2η. □

4.1. *Excursions of Y*

In this section we prove the following, which is slightly stronger than Lemma 4.1.

Lemma 4.2. *For given $P > 0$, $\delta > 0$ the probability that Y hits A_δ before it hits (P, P) is* $O(\frac{1}{\log 1/\varepsilon})$.

The random set $\{t \geq s_0 : Y_t = (0, 0)\}$ is almost surely infinite and discrete. We can therefore order this set, writing T_0 for its smallest element, T_1 for the next, and so on.

Let $U \subset \mathbb{R}^2$ and write H_i for the event $\{Y_t \in U \text{ for some } t \in [T_i, T_{i+1}]\}$. Observe that since our process starts and ends every time interval of the form $[T_i, T_{i+1}]$ at $(0, 0)$ (and in state S_ψ^{2D}) the events $\{H_i : i \geq 0\}$ are independent and have the same probability. We call this probability "the probability that an excursion hits a set U."

Our approach to proving Lemma 4.2 is, to demonstrate that,

$$\mathbb{P}\,\big(\text{an excursion hits } (P, P)\big) \gg \mathbb{P}\,(\text{an excursion hits } A_\delta) \quad \text{as } \varepsilon \to 0.$$

This is realized through the next pair of lemmas whose proofs we delay until Section 4.2.

Lemma 4.3. *For given $\delta > 0$, $\mathbb{P}$ (an excursion hits A_δ) is* $O(\varepsilon)$.

Lemma 4.4. *For given $P > 0$, $\mathbb{P}$ (an excursion hits (P, P)) is $\Omega(\varepsilon \log 1/\varepsilon)$.*

Proof of Lemma 4.2. Before time T_0, Y is in state S_ϕ^{1D} and so the probability that Y hits A_δ before it hits (P, P) is 0.

After time T_0, Y consists of a sequence of excursions, each of which satisfies exactly one of the following conditions

- the excursion hits A_δ (with probability $O(\varepsilon)$),
- the excursion does not hit A_δ but does hit (P, P) (with probability $\Omega(\varepsilon \log 1/\varepsilon) - O(\varepsilon)$, which is itself $\Omega(\varepsilon \log 1/\varepsilon)$),
- the excursion does not hit A_δ or (P, P).

The excursions are independent, so the probability that Y hits A_δ before (P, P) is

$$\frac{O(\varepsilon)}{\Omega(\varepsilon \log 1/\varepsilon) + O(\varepsilon)} = O\left(\frac{1}{\log 1/\varepsilon}\right).$$ □

4.2. *Proofs of the excursion lemmas*

In this section we give the proof of Lemmas 4.3 and 4.4. For convenience we rotate (and scale) $Y = (X^\varepsilon, X)$, defining

$$Z(\varepsilon) = Z = (Z^1, Z^2) = \frac{1}{2}(X^\varepsilon + X, X^\varepsilon - X).$$

Observe that when Y is in S_{ϕ}^{1D}, Z follows the law of a standard one-dimensional Brownian motion on the x-axis. The condition that Y hits A_δ is equivalent to Z^2 hitting $\pm\delta/2$. Like Y, Z has the following "scale invariance" property: the distribution of sample paths of $Z(\varepsilon)/\varepsilon$ is independent of ε (modulo time scaling).

Proof of Lemma 4.3. Consider the process Y between times T_0 and T_1. Once Y arrives at S_{ϕ}^{1D} it can never hit A_δ before hitting $(0,0)$. Thus our goal is to show that with probability at least $1-\mathrm{O}(\varepsilon)$, Y arrives in S_{ϕ}^{1D} before hitting A_δ. Next, we observe the following two auxiliary claims:

Claim 1. *Whenever $Z^2=0$ the probability that subsequently Y arrives at S_{ϕ}^{1D} before Z^2 hits $\pm\varepsilon/2$ is at least a constant* (*independent of* ε).

Claim 2. *Whenever $Z^2=\pm\varepsilon/2$ then there is probability equal to ε/δ of Z^2 hitting $\pm\delta/2$ before it hits* 0.

Claim 1 follows from scale invariance, while Claim 2 is a standard martingale result on Brownian motion (observing that on the relevant time interval Z^2 is a standard Brownian motion).

The reduction of Lemma 4.3 to those two claims is similar to the proof of Lemma 4.2. Claims 1 and 2 imply that between a time when $Z^2=0$ and the next time that $Z^2=0$ after having hit $\pm\varepsilon/2$,

$$\frac{\mathbb{P}(Y \text{ hits } A_\delta)}{\mathbb{P}(Y \text{ arrives at } S_{\phi}^{1D})} \leq \frac{(1-C)(\varepsilon/\delta)}{C} = \mathrm{O}(\varepsilon),$$

where C is the constant of Claim 1. As the behavior of Y is independent on those intervals, we deduce Lemma 4.3. □

Proof of Lemma 4.4. We bound below the probability that an excursion hits (P,P), i.e. Z hits $(P,0)$ before returning to $(0,0)$. We do this by considering the probability that the excursion takes the following form: Z travels from $(0,0)$ to the line segment $Q=[0,\varepsilon]\times\{\varepsilon\}$, then hits the horizontal axis for the first time in $[\varepsilon,1]\times\{0\}$, then travels to $(P,0)$, before returning to $(0,0)$.

After a stopping time at which $Y=Z=(0,0)$ there is a positive probability K that Z hits Q before returning to $(0,0)$. By scale invariance K is independent of ε.

Consider Z after hitting some point in Q. We now bound the hitting density of this process on the horizontal axis. Regardless of the point in Q, this density for points on $[\varepsilon,1]\times\{0\}$ is at least

$$\frac{1}{\pi\varepsilon}\frac{1}{1+(x/\varepsilon)^2}\,\mathrm{d}x.$$

This follows directly from the classical result that the hitting density on a line of the two-dimensional Brownian motion is a Cauchy distribution (see for example Theorem 2.37 of [5]).

On hitting a point $(x,0)$ for $x\in[\varepsilon,1]$ the process transitions from state S_{ϕ}^{2D} to state S_{ϕ}^{1D}. When in state S_{ϕ}^{1D}, Z behaves as a one-dimensional Brownian motion on the horizontal axis until it hits $(0,0)$. By the same martingale argument which justifies Claim 2, the probability of subsequently hitting $(P,0)$ before $(0,0)$ is x/P. Integrating this against the hitting density we get that the probability that Z started from some point in Q hits the horizontal axis in $[\varepsilon,1]\times\{0\}$ and then travels to $(P,0)$ before returning to $(0,0)$ is at least

$$\frac{1}{\pi P}\int_{\varepsilon}^{1}\frac{x/\varepsilon}{1+(x/\varepsilon)^2}\,\mathrm{d}x = \frac{\varepsilon}{2\pi P}\log\left(\frac{1+(1/\varepsilon)^2}{2}\right), \quad \text{which is } \Omega(\varepsilon\log 1/\varepsilon).$$ □

5. Recovering the web from strips

The next step towards proving Theorem 1.1 is to show that the σ-field associated to a horizontal strip is generated by those σ-fields associated to any two substrips which partition the larger strip. To do so, we follow closely the of Section 3.

In the same way we defined $\mathcal{F}_+$ and $\mathcal{F}_-$, we introduce a σ-field $\mathcal{F}_{y,z}$ for each $y < z \in [-\infty, \infty]$, generated by the web in the horizontal strip $(-\infty, \infty) \times (y, z)$. Formally,

Definition 5.1. *For any path f we write $\mathcal{R}_{y,z}(f)$ for f stopped at the first time it is outside $(-\infty,\infty) \times (y,z)$, i.e. $\mathcal{R}_{y,z}(f)(t) = f(t \wedge T_f)$ where $T_f = \inf\{t : f(t) \notin (y,z)\}$ (as in Section 3, if f starts outside $(-\infty,\infty) \times (y,z)$, it is stopped immediately). We define $\mathcal{F}_{y,z}$ to be the σ-field generated by the collection of paths $\{\mathcal{R}_{y,z}(X) : X \in \mathcal{X}\}$. The association of strips $(-\infty,\infty) \times (y,z)$ to σ-fields $\mathcal{F}_{y,z}$ we call the* horizontal factorization of the Brownian web.

With these definitions, $\mathcal{F}$ of Section 3 is $\mathcal{F}_{-\infty,\infty}$, $\mathcal{F}_+$ is $\mathcal{F}_{0,\infty}$ and $\mathcal{F}_-$ is $\mathcal{F}_{-\infty,0}$. Note that $\mathcal{F}_{w,x}$ and $\mathcal{F}_{y,z}$ are independent if the intervals (w,x) and (y,z) are disjoint.

Theorem 5.2. *$\mathcal{F}_{x,z} = \mathcal{F}_{x,y} \otimes \mathcal{F}_{y,z}$ for all $x < y < z$.*

The Brownian web is translation invariant. Thus without loss of generality we can limit ourselves to $x = a$, $y = 0$, $z = b$ in the above theorem, for arbitrary fixed a and b. This reduces the theorem to

$$\mathcal{F}_{a,b} = \mathcal{F}_{a,0} \otimes \mathcal{F}_{0,b}.$$

In the rest of this section we therefore write $\mathcal{R}(\cdot)$ for $\mathcal{R}_{a,b}(\cdot)$.

Given the definition of $\mathcal{F}_{a,b}$ and since X is arbitrary, Theorem 5.2 is a consequence of the following.

Lemma 5.3. *$\mathcal{R}(X)$ is $\mathcal{F}_{a,0} \otimes \mathcal{F}_{0,b}$-measurable.*

Similarly to Section 3, $\mathcal{R}(X^\varepsilon)$ is constructed from trajectories that are $\mathcal{F}_{a,0}$, $\mathcal{F}_{0,b}$ and $\mathcal{G}$-measurable only. Formally, $\mathcal{R}(X^\varepsilon)$ is $\mathcal{F}_{a,0} \otimes \mathcal{F}_{0,b} \otimes \mathcal{G}$-measurable. Thus as in the reduction of Lemma 3.2 to Lemma 3.4, Lemma 5.3 follows from

Lemma 5.4. *$\mathcal{R}(X^\varepsilon) \xrightarrow{\mathbb{P}} \mathcal{R}(X)$ as $\varepsilon \to 0$.*

We could show this convergence result directly by an extension of the argument we used for the half-planes in Section 3. However, knowing that $X^\varepsilon \xrightarrow{\mathbb{P}} X$ is nearly enough, and all that is required in addition is that this convergence is preserved by $\mathcal{R}(\cdot)$. For this we use the following straightforward result in classical analysis.

Lemma 5.5. *If T_f is not a local extremum of the path f, then the map $f \mapsto \mathcal{R}(f)$ is continuous at f in the topology of uniform convergence on compacts.*

Proof of Lemma 5.4. We know from Lemma 3.4 that $X^\varepsilon \xrightarrow{\mathbb{P}} X$. In addition X is a Brownian motion so almost surely satisfies the condition of Lemma 5.5, i.e. "T_X is not a local extremum of the path X." We conclude that $\mathcal{R}(X^\varepsilon) \xrightarrow{\mathbb{P}} \mathcal{R}(X)$ by the continuous mapping theorem (see for example [15], Theorem 2.3, p. 7). □

6. Conclusions about the noise

We conclude by supplying a formal framework for the statement of Theorem 1.1 followed by its proof. The following definition of noise is a straightforward extension of that due to Tsirelson (Definition 3d1 of [13]) to multiple dimensions.

A *d-dimensional noise* consists of a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, sub-$\sigma$-fields $\mathcal{F}_R \subset \mathcal{F}$ given for all open d-dimensional rectangles $R \subset \mathbb{R}^d$, and a measurable action $(T_h)_h$ of the additive group of $\mathbb{R}^d$ on Ω, having the following properties:

(a) $\mathcal{F}_R \otimes \mathcal{F}_{R'} = \mathcal{F}_{R''}$ whenever R and R' partition R'', in the sense that $R \cap R' = \varnothing$ and the closure of $R \cup R'$ is the closure of R'',

(b) T_h sends $\mathcal{F}_R$ to $\mathcal{F}_{R+h}$ for each $h \in \mathbb{R}^d$,
(c) $\mathcal{F}$ is generated by the union of all $\mathcal{F}_R$.

Recall that for σ-fields $\mathcal{F}_a$, $\mathcal{F}_b$, $\mathcal{F}_c$ we write $\mathcal{F}_a = \mathcal{F}_b \otimes \mathcal{F}_c$ when $\mathcal{F}_a$ is generated by $\mathcal{F}_b$ and $\mathcal{F}_c$ (up to sets of measure 0), and $\mathcal{F}_b$ and $\mathcal{F}_c$ are independent. When $d = 1$ our definition coincides with that of Tsirelson. In that case, R ranges over all open intervals and condition (a) translates to $\mathcal{F}_{(s,t)} \otimes \mathcal{F}_{(t,u)} = \mathcal{F}_{(s,u)}$ whenever $s < t < u$.

As conditions (b) and (c) are immediate for the horizontal factorization of the Brownian web, Theorem 5.2 immediately implies the following:

Proposition. *The horizontal factorization of the Brownian web is a* (*one-dimensional*) *noise*.

Recall that the horizontal factorization of the Brownian web is an association of a σ-field to any horizontal strip (see Definition 5.1). Observe that the association arises from considering trajectories of the Brownian web stopped at the first time they are outside a particular strip. Similarly, we can associate a σ-field to any vertical strip, or indeed to any rectangle. The former association is the *vertical factorization of the Brownian web* and the latter the *two-dimensional factorization*.

We can extend existing results to derive the following:

Proposition. *The Brownian web factorized on two-dimensional rectangles is a two-dimensional noise*.

That is, when a rectangle is partitioned horizontally or vertically into two smaller rectangles, the σ-field of the larger is generated by the σ-fields of the two smaller. To see this holds for a rectangle partitioned by a horizontal split observe that this is a consequence of our result restricted to a finite time interval. The see it holds for a vertical split, observe that this is a straightforward modification of the earlier result that the vertical factorization of the Brownian web is a noise (see, for example, [12]).

Furthermore, by a general result of Tsirelson (see [14], Theorem 1e2), a two-dimensional noise is black when one of its one-dimensional factorizations is black. As the vertical factorization of the Brownian web is black (see [13]), we deduce Theorem 1.1.

7. Remarks and open problems

In this paper we present the second known example of a two-dimensional black noise, after the scaling limit of critical planar percolation, as proved by Schramm and Smirnov [7].

Let us illustrate some similarities and differences between our proof and the proof in [7]. Here as in [7] the challenge is to show that a certain object is a noise, while its blackness follows from an observation of Tsirelson. In both proofs the noise property is obtained by showing that the perturbation introduced when resampling a neighbourhood of the interface between the two domains, can be made arbitrarily small.

However, a major difference between the setups is that while all observables in our proof are described in terms of coalescing Brownian motions, in [7] the observables are intangible, being defined only through scaling limits of discrete models. Since our proof involves no scaling limit argument we are able to take advantage of standard techniques for dealing with Brownian motion, allowing a more straightforward presentation.

They also remark that σ-fields can be associated to a larger class of domains than just rectangles in a way that still allows the σ-field of a larger domain to be recovered from the σ-fields of two smaller domains that partition it. Here we call this procedure "extending the noise." In particular, in Remark 1.8 the authors refer to the work of Garban, Pete and Schramm (Remark 8.5 of [4]) in mentioning that this can be done for the scaling limit of site percolation on the triangular lattice, as long as border between those domains has Hausdorff dimension less than $5/4$, and cannot be done if the border has Hausdorff dimension greater then than $5/4$. Consequently, the noise of the scaling limit of site percolation on the triangular lattice can be extended to the class of all domains whose boundaries have Hausdorff dimension less than $5/4$. This raises the following question:

Open Problem 1. *To what class of two-dimensional domains can the Brownian web noise be extended?*

We expect the answer to be more sophisticated than for percolation, since the Brownian web is not rotationally invariant. This suggests that Hausdorff dimension is not a sufficient measurement to determine from which subdomains the Brownian web can be reconstructed. In some sense it is easier to reconstruct the Brownian web from vertical strips than it is from horizontal strips.

Given that we now have a second example of a two-dimensional black noise, one may also ask the following:

Open Problem 2. *Are there black noises in three dimensions and higher?*

Readers may wish to note that in [14] Tsirelson extends the concept of a noise to a much more abstract and general setting. This allows results on noises to be formulated and proved without having an explicit underlying geometric base. However, our methods here which are concrete and geometric in nature are more conveniently described in terms of the earlier formulation.

Acknowledgements

The authors would like to thank Boris Tsirelson for his guidance in the field of noises, for introducing the question to us, and for helpful discussions. In addition we would like to thank Ron Peled and Naomi Feldheim for simplifying some of our arguments and their useful comments on preliminary drafts.

Finally, the authors would like to express our wholehearted gratitude to an anonymous referee whose devoted work in reviewing the paper helped clarifying several vague and misleading points in our original formulation.

References

[1] R. A. Arratia. Coalescing Brownian motions on the line. Ph.D. thesis, Univ. Wisconsin, Madison, 1979. MR2630231
[2] I. Benjamini, G. Kalai and O. Schramm. Noise sensitivity of Boolean functions and applications to percolation. *Inst. Hautes Études Sci. Publ. Math.* **90** (1999) 5–43. MR1813223
[3] L. R. G. Fontes, M. Isopi, C. M. Newman and K. Ravishankar. The Brownian web: Characterization and convergence. *Ann. Probab.* **32** (4) (2004) 2857–2883. MR2094432
[4] C. Garban, G. Pete and O. Schramm. The Fourier spectrum of critical percolation. *Acta Math.* **205** (1) (2010) 19–104. MR2736153
[5] P. Mörters and Y. Peres. *Brownian Motion*. Cambridge Univ. Press, Cambridge, 2010. MR2604525
[6] J. Norris and A. Turner. Planar aggregation and the coalescing Brownian flow. Preprint, 2008. Available at arXiv:0810.0211v1 [math.PR].
[7] O. Schramm and S. Smirnov. On the scaling limits of planar percolation. *Ann. Probab.* **39** (5) (2011) 1768–1814. MR2884873
[8] S. Smirnov and W. Werner. Critical exponents for two-dimensional percolation. *Math. Res. Lett.* **8** (2001) 729–744. MR1879816
[9] B. Toth and W. Werner. The true self-repelling motion. *Probab. Theory Related Fields* **111** (3) (1998) 375–452. MR1640799
[10] B. Tsirelson and A. Vershik. Examples of nonlinear continuous tensor products of measure spaces and non-Fock factorizations. *Rev. Math. Phys.* **10** (1) (1998) 81–145. MR1606855
[11] B. Tsirelson. White noises, black noises and other scaling limits. Lecture course, Tel Aviv Univ., 2002. Available at http://www.tau.ac.il/~tsirel/Courses/Noises/syllabus.html.
[12] B. Tsirelson Scaling limit, noise, stability. In *Lectures on Probability Theory and Statistics* 1–106. *Lecture Notes in Mathematics* **1840**. Springer, Berlin, 2004. MR2079671
[13] B. Tsirelson. Nonclassical stochastic flows and continuous products. *Probab. Surv.* **1** (2004) 173–298. MR2068474
[14] B. Tsirelson. Noise as a Boolean algebra of sigma-fields. *Ann. Probab.* **42** (1) (2014) 311–353. MR3161487
[15] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge Univ. Press, Cambridge, 1998. MR1652247