6

The Brownian Web, the Brownian Net, and their Universality

EMMANUEL SCHERTZER, RONGFENG SUN
AND JAN M. SWART

6.1 Introduction

The Brownian web originated from the work of Arratia's Ph.D. thesis [1], where he studied diffusive scaling limits of coalescing random walk paths starting from everywhere on $\mathbb{Z}$, which can be seen as the spatial genealogies of the population in the dual voter model on $\mathbb{Z}$. Arratia showed that the collection of coalescing random walks converge to a collection of coalescing Brownian motions on $\mathbb{R}$, starting from every point on $\mathbb{R}$ at time 0. Subsequently, Arratia [2] attempted to generalize his result by constructing a system of coalescing Brownian motions starting from everywhere in the space-time plane $\mathbb{R}^2$, which would be the scaling limit of coalescing random walk paths starting from everywhere on $\mathbb{Z}$ at every time $t \in \mathbb{R}$. However, the manuscript [2] was never completed, even though fundamental ideas have been laid down. This topic remained dormant until Tóth and Werner [99] discovered a surprising connection between the one-dimensional space-time coalescing Brownian motions that Arratia tried to construct, and an unusual process called the *true self-repelling motion*, which is repelled by its own local time profile. Building on ideas from [2], Tóth and Werner [99] gave a construction of the system of space-time coalescing Brownian motions, and then used it to construct the true self-repelling motion.

On the other hand, Fontes, Isopi, Newman and Stein [37] discovered that this system of space-time coalescing Brownian motions also arises in the study of aging and scaling limits of one-dimensional spin systems. To establish weak convergence of discrete models to the system of coalescing Brownian motions, Fontes et al. [38, 40] introduced a topology that the system of coalescing Brownian motions starting from every space-time point can be realized as a random variable taking values in a Polish space, and they named this random variable *the Brownian web*. An extension to the Brownian web was later introduced by the authors in [86], and independently by Newman, Ravishankar and Schertzer in [73]. This object was named *the Brownian net* in [86], where the coalescing paths in the Brownian web

are also allowed to branch. To counter the effect of instantaneous coalescence, the branching occurs at an effectively "infinite" rate.

The Brownian web and net have very interesting properties. Their construction is nontrivial due to the uncountable number of starting points in space-time. Coalescence allows one to reduce the system to a countable number of starting points. In fact, the collection of coalescing paths starting from every point on $\mathbb{R}$ at time 0 immediately becomes locally finite when time becomes positive, similar to the phenomenon of *coming down from infinity* in Kingman's coalescent (see e.g., [10]). In fact, the Brownian web can be regarded as the spatial analogue of Kingman's coalescent, with the former arising as the limit of genealogies of the voter model on $\mathbb{Z}$, and the latter arising as the limit of genealogies of the voter model on the complete graph. The key tool in the analysis of the Brownian web, as well as the Brownian net, is its self-duality, similar to the self-duality of critical bond percolation on $\mathbb{Z}^2$. Duality allows one to show that there exist random space-time points where multiple paths originate, and one can give a complete classification of these points. The Brownian web and net also admit a coupling, where the web can be constructed by sampling paths in the net, and conversely, the net can be constructed from the web by Poisson marking a set of "pivotal" points in the web and turning these into points where paths can branch. The latter construction is similar to the construction of scaling limits of near-critical planar percolation from that of critical percolation [19, 48, 49].

The Brownian web and net give rise to a new universality class. In particular, they are expected to arise as the universal scaling limits of one-dimensional interacting particle systems with coalescence, respectively branching-coalescence. One such class of models are population genetic models with resampling and selection, whose spatial genealogies undergo branching and coalescence. Establishing weak convergence to the Brownian web or net can also help in the study of the discrete particle systems themselves. Related models which have been shown to converge to the Brownian web include coalescing random walks [72], succession lines in Poisson trees [36, 23, 46] and drainage network type models [17, 22, 80]. Interesting connections with the Brownian web and net have also emerged from many unexpected sources, including supercritical oriented percolation on $\mathbb{Z}^{1+1}$ [6], planar aggregation models [76, 77], true self-avoiding random walks on $\mathbb{Z}$ [94, 99], random matrix theory [101, 100], and also one-dimensional random walks in i.i.d. space-time random environments [92]. There are also close parallels between the Brownian web and the scaling limit of critical planar percolation, which are the only known examples of two-dimensional *black noise* [95, 96, 87, 33].

The goal of this article is to give an introduction to the Brownian web and net, their basic properties, and how they arise in the scaling limits of one-dimensional interacting particle systems with branching and coalescence. We will focus on the

key ideas, while referring many details to the literature. Our emphasis is naturally biased toward our own research. However, we will also briefly survey related work, including the many interesting connections mentioned above. We have left out many other closely related studies, including diffusion-limited reactions [28, 8] where a dynamic phase transition is observed for branching-coalescing random walks, the propagation of cracks in a sheet [27], rill erosion [31] and the directed Abelian Sandpile Model [25], quantum spin chains [60], etc., which all lie within the general framework of non-equilibrium critical phenomena discussed in the physics surveys [79, 53].

The rest of this article is organized as follows. In Section 6.2, we will construct and give a characterization of the Brownian web and study its properties. In Section 6.3, we do the same for the Brownian net. In Section 6.4, we introduce a coupling between the Brownian web and net and show how one can be constructed from the other. In Section 6.5, we will explain how the Brownian web and net can be used to construct the scaling limits of one-dimensional random walks in i.i.d. space-time random environments. In Section 6.6, we formulate convergence criteria for the Brownian web, which are then applied to coalescing random walks. We will also discuss strategies for proving convergence to the Brownian net. In Section 6.7, we survey other interesting models and results connected to the Brownian web and net. Lastly, in Section 6.8, we conclude with some interesting open questions.

6.2 The Brownian Web

The Brownian web is best motivated by its discrete analogue, the collection of discrete time coalescing simple symmetric random walks on $\mathbb{Z}$, with one walker starting from every site in the space-time lattice $\mathbb{Z}^2_{\text{even}} := \{(x,n) \in \mathbb{Z}^2 : x+n \text{ is even}\}$. The restriction to the sublattice $\mathbb{Z}^2_{\text{even}}$ is necessary due to parity. Figure 6.1 illustrates a graphical construction, where from each $(x,n) \in \mathbb{Z}^2_{\text{even}}$ an independent arrow is drawn from (x,n) to either $(x-1,n+1)$ or $(x+1,n+1)$ with probability $1/2$ each, determining whether the walk starting at x at time n should move to $x-1$ or $x+1$ at time $n+1$. The objects of interest for us are the collection of upward random walk paths (obtained by following the arrows) starting from every space-time lattice point. The question is:

Q.1 What is the diffusive scaling limit of this collection of coalescing random walk paths if space and time are scaled by $1/\sqrt{n}$ and $1/n$ respectively?

Intuitively, it is not difficult to see that the limit should be a collection of coalescing Brownian motions, starting from everywhere in the space-time plane $\mathbb{R}^2$. This is what we will call the *Brownian web*. However, a conceptual difficulty arises, namely that we need to construct the joint realization of *uncountably* many

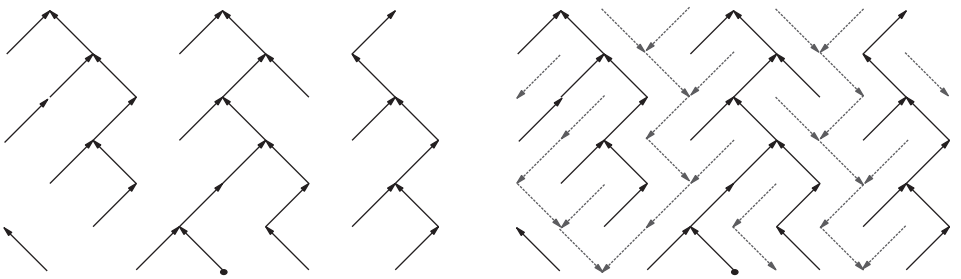

Figure 6.1. Discrete space-time coalescing random walks on $\mathbb{Z}^2_{\text{even}}$, and its dual on $\mathbb{Z}^2_{\text{odd}}$.

Brownian motions. Fortunately it turns out that coalescence allows us to reduce the construction to only a countable collection of coalescing Brownian motions.

Note that in Figure 6.1, we have also drawn a collection of downward arrows connecting points in the odd space-time lattice $\mathbb{Z}^2_{\text{odd}} := \{(x,n) \in \mathbb{Z}^2 : x+n \text{ is odd}\}$, which are dual to the upward arrows by the constraint that the upward and backward arrows do not cross each other. This is the same duality as that for planar bond percolation, and the collection of upward arrows uniquely determines the downward arrows, and vice versa. The collection of downward arrows determines a collection of coalescing random walk paths running backward in time, with one walker starting from each site in $\mathbb{Z}^2_{\text{odd}}$. We may thus strengthen **Q.1** to the following:

Q.2 What is the diffusive scaling limit of the joint realization of the collection of forward and backward coalescing random walk paths?

Observe that the collection of backward coalescing random walk paths has the same distribution as the forward collection, except for a rotation in space-time by 180^o and a lattice shift. Therefore, the natural answer to **Q.2** is that the limit consists of two collections of coalescing Brownian motions starting from everywhere in space-time – one running forward in time and the other backward – and the two collections are equally distributed except for a time-reversal. This is what we will call the (forward) *Brownian web* and the *dual* (backward) *Brownian web*.

In the discrete system, we observe that the collection of forward and the collection of backward coalescing random walk paths uniquely determine each other by the constraint that forward and backward paths cannot cross. It is natural to expect the same for their continuum limits, namely that the Brownian web and the dual Brownian web almost surely uniquely determine each other by the constraint that their paths cannot cross.

The heuristic considerations above, based on discrete approximations, outline the key properties that we expect the Brownian web to satisfy and provide a guide for our analysis.

Before proceeding to a proper construction of the Brownian web and establishing its basic properties, we first define a suitable Polish space in which the Brownian web takes its value. This will be essential to prove weak convergence to the Brownian web.

6.2.1 The Space of Compact Sets of Paths

Following Fontes et al. [40], we regard the collection of colaescing Brownian motions as a *set* of space-time paths, which can be shown to be almost surely relatively compact if space and time are suitably compactified. It is known that given a Polish space E (the space of paths in our case), the space of compact subsets of E, equipped with the induced Hausdorff topology, is a Polish space itself. Therefore a natural space for the Brownian web is the space of *compact sets of paths* (after compactifying space and time), with the Brownian web taken to be the almost sure closure of the set of colaescing Brownian motions. This *paths topology* was inspired by a similar topology proposed by Aizenman [3] to study two-dimensional percolation configurations as closed sets of curves, called the *percolation web*, which was then studied rigorously by Aizenman and Burchard in [4]. We now give the details.

We first compactify $\mathbb{R}^2$. Let R_c^2 denote the completion of the space-time plane $\mathbb{R}^2$ w.r.t. the metric

$$\rho((x_1,t_1),(x_2,t_2)) = |\tanh(t_1) - \tanh(t_2)| \vee \left| \frac{\tanh(x_1)}{1+|t_1|} - \frac{\tanh(x_2)}{1+|t_2|} \right|. \tag{6.1}$$

Note that R_c^2 can be identified with the continuous image of $[-\infty,\infty]^2$ under a map that identifies the line $[-\infty,\infty] \times \{\infty\}$ with a single point $(*,\infty)$, and the line $[-\infty,\infty] \times \{-\infty\}$ with the point $(*,-\infty)$ (see Figure 6.2).

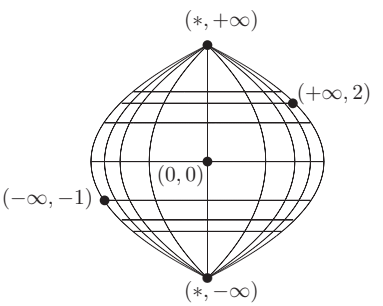

Figure 6.2. The compactification R_c^2 of $\mathbb{R}^2$.

A path π in R_c^2, whose starting time we denote by $\sigma_\pi \in [-\infty,\infty]$, is a mapping $\pi : [\sigma_\pi,\infty] \to [-\infty,\infty] \cup \{*\}$ such that $\pi(\infty) = *$, $\pi(\sigma_\pi) = *$ if $\sigma_\pi = -\infty$, and $t \to (\pi(t),t)$ is a continuous map from $[\sigma_\pi,\infty]$ to (R_c^2,ρ). We then define Π to be the space of all paths in R_c^2 with all possible starting times in $[-\infty,\infty]$. Endowed with the metric

$$d(\pi_1,\pi_2) = \left|\tanh(\sigma_{\pi_1}) - \tanh(\sigma_{\pi_2})\right| \vee \sup_{t \geq \sigma_{\pi_1} \wedge \sigma_{\pi_2}} \left| \frac{\tanh(\pi_1(t \vee \sigma_{\pi_1}))}{1+|t|} - \frac{\tanh(\pi_2(t \vee \sigma_{\pi_2}))}{1+|t|} \right|, \tag{6.2}$$

(Π,d) is a complete separable metric space. Note that convergence in the metric d can be described as locally uniform convergence of paths plus convergence of starting times. (The metric d differs slightly from the original choice in [40], which is somewhat less natural, as explained in the appendix of [86].)

Let $\mathcal{H}$ denote the *space of compact subsets of* (Π,d), equipped with the Hausdorff metric

$$d_{\mathcal{H}}(K_1,K_2) = \sup_{\pi_1 \in K_1} \inf_{\pi_2 \in K_2} d(\pi_1,\pi_2) \vee \sup_{\pi_2 \in K_2} \inf_{\pi_1 \in K_1} d(\pi_1,\pi_2), \tag{6.3}$$

and let $\mathcal{B}_{\mathcal{H}}$ be the Borel σ-algebra associated with $d_{\mathcal{H}}$.

Exercise 6.2.1 *Show that $(\mathcal{H},d_{\mathcal{H}})$ is a complete separable metric space.*

Exercise 6.2.2 *Let $K \subset \Pi$ be compact. Show that $\mathcal{K} := \{\bar{A} : A \subset K\}$ is a compact subset of $\mathcal{H}$.*

For further properties of $(\mathcal{H},d_{\mathcal{H}})$, such as textcolorreda criterion for the convergence of a sequence of elements in $\mathcal{H}$, or necessary and sufficient conditions for the precompactness of a subset of $\mathcal{H}$, see e.g., [92, Appendix B].

We will construct the Brownian web $\mathcal{W}$ as an $(\mathcal{H},\mathcal{B}_{\mathcal{H}})$-valued random variable. The following notational convention will be adopted in the rest of this article:

- For $K \in \mathcal{H}$ and $A \subset R_c^2$, $K(A)$ will denote the set of paths in K with starting points in A.
- When $A = \{z\}$ for $z \in R_c^2$, we also write $K(z)$ instead of $K(\{z\})$.

6.2.2 Construction and Characterization of the Brownian Web

The basic ideas in constructing the Brownian web are the following. First we can construct coalescing Brownian motions starting from a deterministic countable dense subset $\mathcal{D}$ of the space-time plane $\mathbb{R}^2$. It is easily seen that coalescence forces paths started at typical points outside $\mathcal{D}$ to be squeezed between coalescing paths

started from $\mathcal{D}$. Therefore, to construct paths starting from outside $\mathcal{D}$, we only need to take the closure of the set of paths starting from $\mathcal{D}$. Lastly one shows that the law of the random set of paths constructed does not depend on the choice of $\mathcal{D}$. This construction procedure is effectively contained in the following result from [40, Theorem 2.1], which gives a characterization of the Brownian web $\mathcal{W}$ as an $(\mathcal{H}, \mathcal{B}_{\mathcal{H}})$-valued random variable, i.e., a random compact set of paths.

Theorem 6.2.3 (Characterization of the Brownian web) *There exists an* $(\mathcal{H}, \mathcal{B}_{\mathcal{H}})$*-valued random variable* $\mathcal{W}$*, called the standard Brownian web, whose distribution is uniquely determined by the following properties:*

(a) *For each deterministic* $z \in \mathbb{R}^2$*, almost surely there is a unique path* $\pi_z \in \mathcal{W}(z)$*.*
(b) *For any finite deterministic set of points* $z_1, \ldots, z_k \in \mathbb{R}^2$*, the collection* $(\pi_{z_1}, \ldots, \pi_{z_k})$ *is distributed as coalescing Brownian motions.*
(c) *For any deterministic countable dense subset* $\mathcal{D} \subset \mathbb{R}^2$*, almost surely,* $\mathcal{W}$ *is the closure of* $\{\pi_z : z \in \mathcal{D}\}$ *in* (Π, d)*.*

Proof Sketch. We will sketch the main ideas and ingredients and refer the details to [39, 40]. The main steps are:

(1) Let $\mathcal{D} = \{(x,t) : x,t \in \mathbb{Q}\}$ and construct the collection of coalescing Brownian motions $\mathcal{W}(\mathcal{D}) := \{\pi_z\}_{z \in \mathcal{D}}$, where π_z is the Brownian motion starting at z.
(2) Show that $\mathcal{W}(\mathcal{D})$ is almost surely a precompact set in the space of paths (Π, d), and hence $\mathcal{W} := \overline{\mathcal{W}(\mathcal{D})}$ defines a random compact set, i.e., an $(\mathcal{H}, \mathcal{B}_{\mathcal{H}})$-valued random variable.
(3) Show that properties (a) and (b) hold for $\mathcal{W}$, which can be easily seen to imply that property (c) also holds for $\mathcal{W}$.

The above steps construct a random variable $\mathcal{W}$ satisfying properties (a)–(c). Its law is uniquely determined, since if $\widetilde{\mathcal{W}}$ is another random variable satisfying the same properties, then both $\mathcal{W}(\mathcal{D})$ and $\widetilde{\mathcal{W}}(\mathcal{D})$ are coalescing Brownian motions starting from $\mathcal{D}$, and hence can be coupled to equal almost surely. Property (c) then implies that $\mathcal{W} = \widetilde{\mathcal{W}}$ almost surely under this coupling.

Step (1) Fix an order for points in $\mathcal{D}$, so that $\mathcal{D} = \{z_k\}_{k \in \mathbb{N}}$. Coalescing Brownian motions $(\pi_k)_{k \in \mathbb{N}}$ starting respectively from $(z_k)_{k \in \mathbb{N}}$ can be constructed inductively from independent Brownian motions $(\tilde{\pi}_k)_{k \in \mathbb{N}}$ starting from $(z_k)_{k \in \mathbb{N}}$. First let $\pi_1 := \tilde{\pi}_1$. Assuming that $\pi_1, \ldots, \pi_k$ have already been constructed from $\tilde{\pi}_1, \ldots, \tilde{\pi}_k$, then we define the path π_{k+1} to coincide with the independent Brownian motion $\tilde{\pi}_{k+1}$ until the first time τ when it meets one of the already constructed coalescing paths, say π_j, for some $1 \le j \le k$. From time τ onward, we just set π_{k+1} to coincide with π_j. It is not difficult to see that for any $k \in \mathbb{N}$, $(\pi_i)_{1 \le i \le k}$ is a collection of coalescing Brownian motions characterized by the property that different paths

evolve as independent Brownian motions when they are apart, and evolve as the same Brownian motion from the time when they first meet. Furthermore, any subset of a collection of coalescing Brownian motions is also a collection of coalescing Brownian motions.

Step (2) The main idea is the following. The compactification of space-time as shown in Figure 6.2 allows us to approximate $\mathbb{R}^2$ by a large space-time box $\Lambda_{L,T} := [-L,L] \times [-T,T]$, and proving precompactness of $\mathcal{W}(\mathcal{D})$ can be reduced to proving the equicontinuity of paths in $\mathcal{W}(\mathcal{D})$ restricted to $\Lambda_{L,T}$ (for further details, see [40, Appendix B]). More precisely, it suffices to show that for any $\varepsilon > 0$, almost surely we can choose $\delta > 0$ such that the modulus of continuity

$$\psi_{\mathcal{W}(\mathcal{D}),L,T}(\delta) := \sup\{\, |\pi_z(t) - \pi_z(s)| \,:\, z \in \mathcal{D}, (\pi_z(s), s) \in \Lambda_{L,T}, t \in [s, s+\delta]\} \;\le\; \varepsilon. \tag{6.4}$$

Assuming w.l.o.g. that $\varepsilon, \delta \in \mathbb{Q}$, we will control $\psi_{\mathcal{W}(\mathcal{D}),L,T}(\delta) > \varepsilon$ in terms of the modulus of continuity of coalescing Brownian motions starting from the grid

$$G_{\varepsilon,\delta} := \{(m\varepsilon/4, n\delta) : m, n \in \mathbb{N}\} \cap [-L-\varepsilon, L+\varepsilon] \times [-T-\delta, T] \;\subset \mathcal{D}.$$

Indeed, $\psi_{\mathcal{W}(\mathcal{D}),L,T}(\delta) > \varepsilon$ means that $|\pi_z(t) - \pi_z(s)| > \varepsilon$ for some $z \in \mathcal{D}$ with $(\pi_z(s), s) \in \Lambda_{L,T}$ and $t \in [s, s+\delta]$. Then there exists a point in the grid $\tilde{z} = (\tilde{x}, \tilde{t}) \in G_{\varepsilon,\delta}$ with $s \in [\tilde{t}, \tilde{t} + \delta)$ and $\tilde{x} \in (\pi(s) \wedge \pi(t) + \varepsilon/4, \pi(s) \wedge \pi(t) - \varepsilon/4)$. Since π_z and $\pi_{\tilde{z}}$ are coalescing Brownian motions, either $\pi_{\tilde{z}}$ coalesces with π_z before time t, or $\pi_{\tilde{z}}$ avoids π_z up to time t. Either way, we must have

$$\sup_{h \in [0, 2\delta]} |\pi_{\tilde{z}}(\tilde{t} + h) - \pi_{\tilde{z}}(\tilde{t})| \ge \varepsilon/4.$$

Denote this event by $E_z^{\varepsilon,\delta}$. Then

$$\begin{aligned}\mathbb{P}(\psi_{\mathcal{W}(\mathcal{D}),L,T}(\delta) > \varepsilon) \le \mathbb{P}\Big(\bigcup_{z \in G_{\varepsilon,\delta}} E_z^{\varepsilon,\delta} \Big) &\le \sum_{z \in G_{\varepsilon,\delta}} \mathbb{P}(E_z^{\varepsilon,\delta}) = |G_{\varepsilon,\delta}| \mathbb{P}\Big(\sup_{h \in [0,2\delta]} |B_h| \ge \varepsilon/4 \Big) \\ &\le C_{L,T} \varepsilon^{-1} \delta^{-1} e^{-c\varepsilon^2/\delta},\end{aligned}$$

where $c, C_{L,T} > 0$, B is a standard Brownian motion, and we have used the reflection principle to bound the tail probability for $\sup |B|$. Since $\mathbb{P}(\psi_{\mathcal{W}(\mathcal{D}),L,T}(\delta) > \varepsilon) \to 0$ as $\delta \downarrow 0$, this implies (6.4). Therefore $\mathcal{W}(\mathcal{D})$ is a.s. precompact, and $\mathcal{W} := \overline{\mathcal{W}(\mathcal{D})}$ defines an $(\mathcal{H}, \mathcal{B}_{\mathcal{H}})$-valued random variable.

Step (3) We first show that for each $z = (x,t) \in \mathbb{R}^2$, almost surely $\mathcal{W}(z)$, the paths in $\mathcal{W}$ starting at z, contains a unique path. Let $z_n^- = (x - \varepsilon_n, t - \delta_n) \in \mathcal{D}$, $z_n^+ = (x + \varepsilon_n, t - \delta_n) \in \mathcal{D}$, with $\varepsilon_n, \delta_n \downarrow 0$, and let τ_n be the time when $\pi_{z_n^-}$ and $\pi_{z_n^+}$ coalesce. Note that on the event

$$E_n := \left\{ \pi_{z_n^-}(t) < x < \pi_{z_n^+}(t), \tau_n \le t + \frac{1}{n} \right\},$$

every path in $\mathcal{W}(z)$ must be enclosed between $\pi_{z_n^-}$ and $\pi_{z_n^+}$, and hence is uniquely determined from time $\tau_n \le t + \frac{1}{n}$ onward. It is easy to see that we can choose $\varepsilon_n \downarrow 0$ sufficiently fast, and $\delta_n \downarrow 0$ much faster than ε_n^2, such that $\mathbb{P}(E_n) \to 1$ as $n \to \infty$. In particular, almost surely, E_n occurs infinitely often, which implies that the paths in $\mathcal{W}(z)$ all coincide on (t, ∞) and hence $\mathcal{W}(z)$ contains a unique path.

To show that $\mathcal{W}$ satisfies property (b), let us fix $z_1, \ldots, z_k \in \mathbb{R}^2$. For each $1 \le i \le k$, let $z_{n,i} \in \mathcal{D}$ with $z_{n,i} \to z_i$ as $n \to \infty$. By the a.s. compactness of $\mathcal{W}$, and the fact that $\mathcal{W}(z_i)$ a.s. contains a unique path π_{z_i} by property (a) that we just verified, we must have $\pi_{z_{n,i}} \to \pi_{z_i}$ in (Π, d) for each $1 \le i \le k$. In particular, as a sequence of Π^k-valued random variables, $(\pi_{z_{n,i}})_{1 \le i \le k}$ converges in distribution to $(\pi_{z_i})_{1 \le i \le k}$. On the other hand, as a subset of $\mathcal{W}(\mathcal{D})$, $(\pi_{z_{n,i}})_{1 \le i \le k}$ is a collection of coalescing Brownian motions, and it is easy to show that as their starting points converge, they converge in distribution to a collection of coalescing Brownian motions starting from $(z_i)_{1 \le i \le k}$. Therefore $(\pi_{z_i})_{1 \le i \le k}$ is distributed as a collection of coalescing Brownian motions.

Lastly, to show that $\mathcal{W}$ satisfies property (c), let $\mathcal{D}'$ be another countable dense subset of $\mathbb{R}^2$. Clearly $\overline{\mathcal{W}(D')} \subset \mathcal{W}$. To show the converse, $\mathcal{W} \subset \overline{\mathcal{W}(D')}$, it suffices to show that for each $z \in \mathcal{D}$, $\pi_z \in \overline{\mathcal{W}(D')}$. This can be seen by taking a sequence $z'_n \in \mathcal{D}'$ with $z'_n \to z$, for which we must have $\pi_{z'_n} \to \pi_z \in \overline{\mathcal{W}(D')}$ by the compactness of $\overline{\mathcal{W}(D')} \subset \mathcal{W}$ and the fact that $\mathcal{W}(z) = \{\pi_z\}$. □

6.2.3 The Brownian Web and its Dual

As discussed at the beginning of Section 6.2, similar to the duality between forward and backward coalescing random walks shown in Figure 6.1, we expect the Brownian web $\mathcal{W}$ also to have a dual $\widehat{\mathcal{W}}$. Such a duality provides a powerful tool for analyzing properties of the Brownian web. Since the dual Brownian web $\widehat{\mathcal{W}}$ should be a collection of coalescing paths running backward in time, we first define the space in which $\widehat{\mathcal{W}}$ takes its values.

Given $z = (x, t) \in R_c^2$, which is identified with $[-\infty, \infty]^2$ where $[-\infty, \infty] \times \{\pm\infty\}$ is contracted to a single point $(*, \pm\infty)$, let $-z$ denote $(-x, -t)$. Given a set $A \subset R_c^2$, let $-A$ denote $\{-z : z \in A\}$. Identifying each path $\pi \in \Pi$ with its graph as a subset of R_c^2, $\hat{\pi} := -\pi$ defines a path running backward in time, with starting time $\hat{\sigma}_{\hat{\pi}} = -\sigma_\pi$. Let $\widehat{\Pi} := -\Pi$ denote the set of all such backward paths, equipped with a metric $\hat{d}$ that is inherited from (Π, d) under the mapping $-$. Let $\widehat{\mathcal{H}}$ be the space of compact subsets of $(\widehat{\Pi}, \hat{d})$, equipped with the Hausdorff metric $d_{\widehat{\mathcal{H}}}$ and Borel σ-algebra $\mathcal{B}_{\widehat{\mathcal{H}}}$. For any $K \in \mathcal{H}$, we will let $-K$ denote the set $\{-\pi : \pi \in K\} \in \widehat{\mathcal{H}}$.

The following result characterizes the joint law of the Brownian web $\mathcal{W}$ and its dual $\widehat{\mathcal{W}}$ as a random variable taking values in $\mathcal{H} \times \widehat{\mathcal{H}}$, equipped with the product σ-algebra.

Theorem 6.2.4 (Characterization of the double Brownian web) *There exists an $\mathcal{H} \times \widehat{\mathcal{H}}$-valued random variable $(\mathcal{W}, \widehat{\mathcal{W}})$, called the double Brownian web (with $\widehat{\mathcal{W}}$ called the dual Brownian web), whose distribution is uniquely determined by the following properties:*

(a) *$\mathcal{W}$ and $-\widehat{\mathcal{W}}$ are both distributed as the standard Brownian web.*
(b) *Almost surely, no path $\pi_z \in \mathcal{W}$ crosses any path $\hat{\pi}_{\hat{z}} \in \widehat{\mathcal{W}}$ in the sense that, $z = (x,t)$ and $\hat{z} = (\hat{x}, \hat{t})$ with $t < \hat{t}$, and $(\pi_z(s_1) - \hat{\pi}_{\hat{z}}(s_1))(\pi_z(s_2) - \hat{\pi}_{\hat{z}}(s_2)) < 0$ for some $t < s_1 < s_2 < \hat{t}$.*

Furthermore, for each $z \in \mathbb{R}^2$, $\widehat{\mathcal{W}}(z)$ a.s. consists of a single path $\hat{\pi}_z$ which is the unique path in $\widehat{\Pi}$ that does not cross any path in $\mathcal{W}$, and thus $\widehat{\mathcal{W}}$ is a.s. determined by $\mathcal{W}$ and vice versa.

Proof Sketch. The existence of a double Brownian web $(\mathcal{W}, \widehat{\mathcal{W}})$ satisfying properties (a)–(b) is most easily derived as scaling limits of forward and backward coalescing random walks. We defer this to Section 6.6, after we introduce general criteria for convergence to the Brownian web.

Let us first prove that if $(\mathcal{W}, \widehat{\mathcal{W}})$ satisfies properties (a)–(b), then almost surely $\mathcal{W}$ uniquely determines $\widehat{\mathcal{W}}$. Indeed, fix a deterministic $z = (x,t) \in \mathbb{R}^2$. By the characterization of the Brownian web $\mathcal{W}$, $\mathcal{W}(\mathbb{Q}^2)$ is a collection of colaescing Brownian motions, with a.s. one Brownian motion starting from each point in $\mathbb{Q}^2$. Since Brownian motion has zero probability of hitting a deterministic space-time point, there is zero probability that z lies on $\pi_{z'}$ for any $z' \in \mathbb{Q}^2$. Therefore for any $s \in \mathbb{Q}$ with $s < t$, property (b) implies that for any path $\hat{\pi} \in \widehat{\mathcal{W}}(z)$, we must have

$$\hat{\pi}(s) = \sup\{y \in \mathbb{Q} : \pi_{(y,s)}(t) < x\} = \inf\{y \in \mathbb{Q} : \pi_{(y,s)}(t) > x\}.$$

In other words, $\hat{\pi}$ is uniquely determined at rational times and hence at all times, and $\widehat{\mathcal{W}}(z)$ contains a unique path. It follows that $\widehat{\mathcal{W}}(\mathbb{Q}^2)$ is a.s. uniquely determined by $\mathcal{W}$, and hence so is $\widehat{\mathcal{W}} = \overline{\widehat{\mathcal{W}}(\mathbb{Q}^2)}$.

Lastly we show that the distribution of $(\mathcal{W}, \widehat{\mathcal{W}})$ is uniquely determined by properties (a) and (b). Indeed, if $(\mathcal{W}', \widehat{\mathcal{W}}')$ is another double Brownian web, then $\mathcal{W}$ and $\mathcal{W}'$ can be coupled so that they equal a.s. As we have just shown, $\mathcal{W}$ a.s. uniquely determines $\widehat{\mathcal{W}}$, and $\mathcal{W}'$ determines $\widehat{\mathcal{W}}'$. Therefore $\widehat{\mathcal{W}} = \widehat{\mathcal{W}}'$ a.s., and $(\mathcal{W}', \widehat{\mathcal{W}}')$ has the same distribution as $(\mathcal{W}, \widehat{\mathcal{W}})$. □

Remark 6.2.5 One can characterize the joint law of paths in $(\mathcal{W}, \widehat{\mathcal{W}})$ starting from a finite deterministic set of points. Similar to the construction of coalescing

Brownian motions, we can construct one path at a time. To add a new forward path to an existing collection, we follow an independent Brownian motion until it either meets an existing forward Brownian motion, in which case they coalesce, or it meets an existing dual Brownian motion, in which case it is *Skorohod reflected* by the dual Brownian motion (for further details, see [93]). Extending this pathwise construction to a countable dense set of starting points $\mathcal{D}$ and then taking closure, this gives a direct construction of $(\mathcal{W}, \widehat{\mathcal{W}})$, which is formulated in [41, Theorem 3.7].

Remark 6.2.6 In light of Theorem 6.2.4, one may wonder whether $\mathcal{W}$ a.s. consists of *all* paths in Π which do not cross any path in $\widehat{\mathcal{W}}$, and vice versa. The answer is no, and $\mathcal{W}$ is actually the minimal compact set of paths that do not cross any path in $\widehat{\mathcal{W}}$ while still containing paths starting from every point in $\mathbb{R}^2$. More non-crossing paths can be added to $\mathcal{W}$ by extending paths in $\mathcal{W}$ backward in time, following paths in $\widehat{\mathcal{W}}$ (see [44]). Such paths can be excluded if we impose the further restriction that no path can enter from outside any open region that is enclosed by a pair of paths in $\widehat{\mathcal{W}}$. This is called the *wedge characterization* of the Brownian web, to be discussed in more detail in Remark 6.3.10.

6.2.4 The Coalescing Point Set

The coupling between the Brownian web and its dual given in Theorem 6.2.4 allows one to deduce interesting properties for the Brownian web. The first result is on the density of paths in the Brownian web $\mathcal{W}$ started at time 0.

Given the Brownian web $\mathcal{W}$, and a closed set $A \subset \mathbb{R}$, define the *coalescing point set* by

$$\xi_t^A := \{y \in \mathbb{R} : y = \pi(t) \text{ for some } \pi \in \mathcal{W}(A \times \{0\})\}, \qquad t \geq 0. \tag{6.5}$$

In words, ξ_t^A is the set of points in $\mathbb{R}$ that lie on some path in $\mathcal{W}$ that starts from A at time 0. Note that this process is monotone in the sense that if $A \subset B$, then $\xi_t^A \subset \xi_t^B$ a.s. for all $t \geq 0$.

It turns out that even if started from the whole line, $\xi_t^{\mathbb{R}}$ becomes a.s. locally finite as soon as $t > 0$, as the following density result shows. Such a *coming down from infinity* phenomenon also appears in Kingman's coalescent, see e.g., [10].

Proposition 6.2.7 (Density of the coalescing point set) *Let $\xi_\cdot^{\mathbb{R}}$ be the coalescing point set defined from the Brownian web $\mathcal{W}$ as in* (6.5)*. Then for all $t > 0$ and $a < b$,*

$$\mathbb{E}[|\xi_t^{\mathbb{R}} \cap [a, b]|] = \frac{b-a}{\sqrt{\pi t}}. \tag{6.6}$$

Proof. Let $\widehat{\mathcal{W}}$ be the dual Brownian web determined a.s. by $\mathcal{W}$, as in Theorem 6.2.4. Observe that by the non-crossing property between paths in $\mathcal{W}$ and $\widehat{\mathcal{W}}$, $\xi_t^{\mathbb{R}} \cap (a,b) \neq \emptyset$ implies that the paths $\hat{\pi}_a, \hat{\pi}_b \in \widehat{\mathcal{W}}$, starting respectively at (a,t) and (b,t), do not coalesce in the time interval $(0,t)$ (i.e., $\hat{\tau} \leq 0$ if $\hat{\tau}$ denotes the time when $\hat{\pi}_a$ and $\hat{\pi}_b$ coalesce). Conversely, if $\hat{\tau} \leq 0$, then any path in $\mathcal{W}$ started from $[\hat{\pi}_a(0), \hat{\pi}_b(0)]$ at time 0 will hit $[a,b]$ at time t, i.e., $\xi_t^{\mathbb{R}} \cap [a,b] \neq \emptyset$. Thus,

$$\mathbb{P}(\xi_t^{\mathbb{R}} \cap (a,b) \neq \emptyset) \leq \mathbb{P}(\hat{\tau} \leq 0) \leq \mathbb{P}(\xi_t^{\mathbb{R}} \cap [a,b] \neq \emptyset), \tag{6.7}$$

where we observe that

$$\begin{aligned}\mathbb{P}(\hat{\tau} \leq 0) &= \mathbb{P}\Big(\sup_{s\in(0,t)} (B_2(s) - B_1(s)) \leq b-a \Big) = \mathbb{P}\Big(\sup_{s\in(0,t)} B(s) \leq \frac{b-a}{\sqrt{2}} \Big) \\ &= \frac{1}{\sqrt{2\pi t}} \int_{-\frac{b-a}{\sqrt{2}}}^{\frac{b-a}{\sqrt{2}}} e^{-\frac{x^2}{2t}} \, \mathrm{d}x \sim \frac{b-a}{\sqrt{\pi t}} \quad \text{as } b-a \downarrow 0.\end{aligned}$$

By (6.7), this implies that

$$\mathbb{P}(x \in \xi_t^{\mathbb{R}}) = 0 \quad \text{for all } x \in \mathbb{R},$$

and the inequalities in (6.7) are in fact all equalities.

We can then apply the monotone convergence theorem to obtain

$$\begin{aligned}\mathbb{E}[|\xi_t^{\mathbb{R}} \cap [a,b]|] &= \lim_{n\to\infty} \mathbb{E}\Big[\Big| \Big\{ 1 \leq i < (b-a)2^n : \xi_t^{\mathbb{R}} \cap \Big(a + \frac{i-1}{2^n}, a + \frac{i}{2^n} \Big) \neq \emptyset \Big\} \Big| \Big] \\ &= \lim_{n\to\infty} (b-a)2^n \mathbb{P}(\xi_t^{\mathbb{R}} \cap (0, 2^{-n}) \neq \emptyset) = \lim_{n\to\infty} (b-a)2^n \frac{2^{-n}}{\sqrt{\pi t}} \\ &= \frac{b-a}{\sqrt{\pi t}}.\end{aligned}$$

This concludes the proof of the proposition. □

As a corollary of Proposition 6.2.7, we show that when paths in the Brownian web converge, they converge in a strong sense (see e.g., [86, Lemma 3.4]).

Corollary 6.2.8 (Strong convergence of paths in $\mathcal{W}$) *Let $\mathcal{W}$ be the standard Brownian web. Almost surely, for any sequence $\pi_n \in \mathcal{W}$ with $\pi_n \to \pi \in \mathcal{W}$, the time of coalescence τ_n between π_n and π must tend to σ_π as $n \to \infty$.*

Exercise 6.2.9 *Deduce Corollary 6.2.8 from Proposition 6.2.7.*

Remark 6.2.10 Apart from its density, we actually know quite a bit more about the coalescing point set $\xi_t^{\mathbb{R}}$. It has been shown by Tribe et al. [101, 100] that $\xi_t^{\mathbb{R}}$

is in fact a *Pfaffian point process*, whose kernel also appears in the real Ginibre random matrix ensemble. Furthermore, $\xi_t^{\mathbb{R}}$ (and more generally ξ_t^A for any $A \subset \mathbb{R}$) can be shown (see e.g., [52, Appendix C] and [69]) to be negatively associated in the sense that for any $n \in \mathbb{N}$ and any disjoint open intervals $O_1, \cdots, O_n$, we have

$$\mathbb{P}(\cap_{i=1}^n \{\xi_t^{\mathbb{R}} \cap O_i \neq \emptyset\}) \leq \prod_{i=1}^n \mathbb{P}(\xi_t^{\mathbb{R}} \cap O_i \neq \emptyset). \tag{6.8}$$

For any $B \subset \mathbb{R}$ with positive Lebesgue measure, we also have

$$\mathbb{P}(|\xi_t^{\mathbb{R}} \cap B| \geq m+n) \leq \mathbb{P}(|\xi_t^{\mathbb{R}} \cap B| \geq m)\mathbb{P}(|\xi_t^{\mathbb{R}} \cap B| \geq n) \qquad \forall m,n \in \mathbb{N}. \tag{6.9}$$

On a side note, we remark that when $A \subset \mathbb{R}$ is a finite set, determinantal formulas have also been derived for the distribution of ξ_t^A in [103, Proposition. 9].

6.2.5 Special Points of the Brownian Web

We have seen in Theorem 6.2.3 that for each deterministic $z \in \mathbb{R}^2$, almost surely the Brownian web $\mathcal{W}$ contains a unique path starting from z. However, it is easily seen that there must exist random points $z \in \mathbb{R}^2$ where $\mathcal{W}(z)$ contains multiple paths. Indeed, consider paths in $\mathcal{W}$ starting from $\mathbb{R}$ at time 0. Proposition 6.2.7 shows that these paths coalesce into a locally finite set of points $\xi_t^{\mathbb{R}}$ at any time $t > 0$. Each point $x_i \in \xi_t^{\mathbb{R}}$ (with $x_i < x_{i+1}$ for all $i \in \mathbb{Z}$) can be traced back to an interval (u_i, u_{i+1}) at time 0, where all paths starting there pass through the space-time point (x_i, t). At the boundary u_i between two such intervals, we note however that $\mathcal{W}((u_i, 0))$ must contain at least two paths, which are limits of paths in $\mathcal{W}$ starting from (u_{i-1}, u_i), respectively (u_i, u_{i+1}), at time 0. Are there random space-time points where more than two paths originate? It turns out that we can give a complete classification of the type of multiplicity we see almost surely in a Brownian web. The main tool to accomplish this is the self-duality of the Brownian web discussed in Section 6.2.3.

First we give a classification scheme for $z \in \mathbb{R}^2$ according to the multiplicity of paths in $\mathcal{W}$ entering and leaving z. We say a path π enters $z = (x,t)$ if $\sigma_\pi < t$ and $\pi(t) = x$, and π leaves z if $\sigma_\pi \leq t$ and $\pi(t) = x$. Two paths π and π' leaving z are defined to be equivalent, denoted by $\pi \sim_{\text{out}}^z \pi'$, if $\pi = \pi'$ on $[t, \infty)$. Two paths π and π' entering z are defined to be equivalent, denoted by $\pi \sim_{\text{in}}^z \pi'$, if $\pi = \pi'$ on $[t - \varepsilon, \infty)$ for some $\varepsilon > 0$. Note that $\sim_{\text{in}}^z$ and $\sim_{\text{out}}^z$ are equivalence relations.

Let $m_{\text{in}}(z)$, respectively $m_{\text{out}}(z)$, denote the number of equivalence classes of paths in $\mathcal{W}$ entering, respectively leaving, z, and let $\hat{m}_{\text{in}}(z)$ and $\hat{m}_{\text{out}}(z)$ be defined similarly for the dual Brownian web $\hat{\mathcal{W}}$. Given a realization of the Brownian web $\mathcal{W}$, points $z \in \mathbb{R}^2$ are classified according to the value of $(m_{\text{in}}(z), m_{\text{out}}(z))$. We divide points of type (1,2) further into types $(1,2)_{\text{l}}$ and $(1,2)_{\text{r}}$, where the subscript

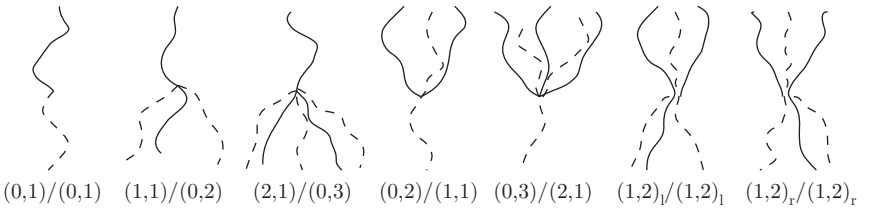

Figure 6.3. Special points of the Brownian web.

l (resp. r) indicates that the left (resp. right) of the two outgoing paths is the continuation of the (up to equivalence) unique incoming path. Points in the dual Brownian web $\widehat{\mathcal{W}}$ are labeled according to their type in the Brownian web obtained by rotating the graph of $\widehat{\mathcal{W}}$ in $\mathbb{R}^2$ by 180°.

We are now ready to state the following classification result (see also [99, Proposition 2.4] and [41, Theorems 3.11–3.14]), illustrated in Figure 6.3.

Theorem 6.2.11 (Special points of the Brownian web) *Let $(\mathcal{W}, \widehat{\mathcal{W}})$ be the standard Brownian web and its dual. Then almost surely, each $z \in \mathbb{R}^2$ satisfies*

$$m_{\text{out}}(z) = \hat{m}_{\text{in}}(z) + 1 \qquad \text{and} \qquad \hat{m}_{\text{out}}(z) = m_{\text{in}}(z) + 1, \tag{6.10}$$

and z is of one of the following seven types according to $(m_{\text{in}}(z), m_{\text{out}}(z))/(\hat{m}_{\text{in}}(z), \hat{m}_{\text{out}}(z))$:

$$(0,1)/(0,1),\ (1,1)/(0,2),\ (0,2)/(1,1),\ (2,1)/(0,3),\ (0,3)/(2,1),\ (1,2)_{\text{l}}/(1,2)_{\text{l}},$$
$$(1,2)_{\text{r}}/(1,2)_{\text{r}}.$$

Almost surely,

(i) *the set of points of type $(0,1)/(0,1)$ has full Lebesgue measure in $\mathbb{R}^2$;*
(ii) *points of type $(1,1)/(0,2)$ are points in the set $\bigcup_{\pi \in \mathcal{W}} \{(\pi(t), t) : t > \sigma_\pi\}$, excluding points of type $(1,2)/(1,2)$ and $(2,1)/(0,3)$;*
(iii) *the set of points of type $(2,1)/(0,3)$ consists of points at which two paths in $\mathcal{W}$ coalesce and is countable;*
(iv) *points of type $(1,2)_{\text{l}}/(1,2)_{\text{l}}$ are points of intersection between some $\pi \in \mathcal{W}$ and $\hat{\pi} \in \widehat{\mathcal{W}}$, with $\sigma_\pi < t < \hat{\sigma}_{\hat{\pi}}$, $\pi(s) \le \hat{\pi}(s)$ for all $s \in [\sigma_\pi, \hat{\sigma}_{\hat{\pi}}]$, and π intersects $\hat{\pi}$ at $(\pi(t), t) = (\hat{\pi}(t), t)$.*

Similar statements hold for the remaining three types by symmetry.

Proof. First we will prove relation (6.10). Let $z = (x,t)$, and assume that $\hat{m}_{\text{in}}(z) = k$ for some $k \in \mathbb{N}_0 := \{0\} \cup \mathbb{N}$. Then there exist $\varepsilon > 0$ and k ordered paths $\hat{\pi}_1, \ldots, \hat{\pi}_k \in \widehat{\mathcal{W}}$ starting at time $t + \varepsilon$, such that these paths are disjoint on $(t, t+\varepsilon]$ and coalesce together at time t at position x. Note that the ordered paths $(\hat{\pi}_i)_{1 \le i \le k}$ divide the

space-time strip $\mathbb{R} \times (t, t+\varepsilon)$ into $k+1$ regions $(I_i)_{1 \le i \le k+1}$, where I_1 is the region to the left of $\hat{\pi}_1$, I_i is the region between $\hat{\pi}_{i-1}$ and $\hat{\pi}_i$ for each $2 \le i \le k$, and I_{k+1} is the region to the right of $\hat{\pi}_k$. From the interior of each region I_i, we can pick a sequence of starting points $(z_i^n)_{n \in \mathbb{N}}$ with $z_i^n \to z$. Since paths in $\mathcal{W}$ and $\widehat{\mathcal{W}}$ do not cross, as formulated in Theorem 6.2.4, each path $\pi_i^n \in \mathcal{W}(z_i^n)$ must stay confined in $\overline{I_i}$ in the time interval $[t, t+\varepsilon]$, and so must any subsequential limit of $(\pi_i^n)_{n \in \mathbb{N}}$. Such subsequential limits must exist by the almost sure compactness of $\mathcal{W}$, and each subsequential limit is a path $\pi_i \in \mathcal{W}(z)$. Therefore $\mathcal{W}(z)$ must contain at least $k+1$ distinct paths, one contained in $\overline{I_i}$ for each $1 \le i \le k+1$. Furthermore, each $\overline{I_i}$ cannot contain more than one path in $\mathcal{W}(z)$. Indeed, if $\overline{I_i}$ contains two distinct paths $\pi, \pi' \in \mathcal{W}(z)$, then any path $\hat{\pi} \in \widehat{\mathcal{W}}$ started strictly between π and π' on the time interval $(t, t+\varepsilon)$ must enter z, and $\hat{\pi}$ is distinct from $(\hat{\pi}_i)_{1 \le i \le k}$, which contradicts the assumption that $\hat{m}_{\text{in}}(z) = k$. Therefore we must have $m_{\text{out}}(z) = k+1 = \hat{m}_{\text{in}}(z) + 1$.

Similar considerations as above show that if z is of type $(1,2)_{\text{l}}$ in $\mathcal{W}$, then it must be of the same type in $\widehat{\mathcal{W}}$. The same holds for type $(1,2)_{\text{r}}$.

We now show that the seven types of points listed are all there is. Note that it suffices to show that almost surely $m_{\text{in}}(z) + m_{\text{out}}(z) \le 3$ for each $z \in \mathbb{R}^2$. There are four possible cases of $m_{\text{in}}(z) + m_{\text{out}}(z) > 3$, which we rule out one by one:

(a) For some $z \in \mathbb{R}^2$, $m_{\text{in}}(z) \ge 3$ and $m_{\text{out}}(z) \ge 1$. Note that Corollary 6.2.8 implies that every path $\pi \in \mathcal{W}$ coincides with some path in $\mathcal{W}(\mathbb{Q}^2)$ on $[\sigma_\pi + \varepsilon, \infty)$, for any given $\varepsilon > 0$. Therefore the event that $m_{\text{in}}(z) \ge 3$ for some $z \in \mathbb{R}^2$ is contained in the event that three distinct Brownian motions among $\mathcal{W}(\mathbb{Q}^2)$ coalesce at the same time. Such an event has probability zero, because there are countably many ways of choosing three Brownian motions from $\mathcal{W}(\mathbb{Q}^2)$, and conditioned on two Brownian motions coalescing at a given space-time point, there is zero probability that a third independent Brownian motion (which evolves independently before coalescing) would visit the same space-time point.

(b) For some $z \in \mathbb{R}^2$, $m_{\text{in}}(z) \ge 2$ and $m_{\text{out}}(z) \ge 2$. In this case, $\hat{m}_{\text{in}}(z) \ge 1$, and, again by Corollary 6.2.8, the event we consider is contained in the event that there exist two paths $\pi_1, \pi_2 \in \mathcal{W}(\mathbb{Q}^2)$ and a path $\hat{\pi} \in \widehat{\mathcal{W}}(\mathbb{Q}^2)$, such that $\hat{\pi}$ passes through the point of coalescence between π_1 and π_2. Such an event has probability 0, since conditioned on π_1 and π_2 up to the time of their coalescence, $\hat{\pi}$ is an independent Brownian motion with zero probability of hitting a given space-time point – the point of coalescence between π_1 and π_2.

(c) For some $z \in \mathbb{R}^2$, $m_{\text{in}}(z) \ge 1$ and $m_{\text{out}}(z) \ge 3$. In this case, $\hat{m}_{\text{out}}(z) \ge 2$ and $\hat{m}_{\text{in}}(z) \ge 2$, which is equivalent to Case (b) by the symmetry between $\mathcal{W}$ and $\widehat{\mathcal{W}}$.

(d) For some $z \in \mathbb{R}^2$, $m_{\text{out}}(z) \geq 4$. In this case, $\hat{m}_{\text{in}}(z) \geq 3$, which is equivalent to Case (a).

We leave the verification of statements (i)–(iv) as an exercise. □

Exercise 6.2.12 *Verify statements (i)–(iv) in Theorem 6.2.11.*

6.3 The Brownian Net

The Brownian net generalizes the Brownian web by allowing paths to branch. The existence of such an object is again motivated by its discrete analogue, the collection of discrete time branching-coalescing simple symmetric random walks on $\mathbb{Z}$. Figure 6.4 gives an illustration: from each $(x,n) \in \mathbb{Z}^2_{\text{even}}$, an arrow is drawn from (x,n) to either $(x-1,n+1)$ or $(x+1,n+1)$ with probability $(1-\varepsilon)/2$ each, representing whether the walk starting at x at time n should move to $x-1$ or $x+1$ at time $n+1$; and with probability ε, arrows are drawn from (x,n) to both $(x-1,n+1)$ and $(x+1,n+1)$, so that the walk starting at (x,n) branches into two walks, with one moving to $x-1$ and the other to $x+1$ at time $n+1$.

If we consider the collection of all upward random walk paths obtained by following the arrows, then the natural question is: when space-time is scaled diffusively, could this random collection of paths have a nontrivial limit? To have an affirmative answer to this question, it is necessary to choose the branching probability ε to depend suitably on the diffusive scaling parameter. More precisely:

Q.1 If space-time is rescaled by $S_\varepsilon(x,t) := (\varepsilon x, \varepsilon^2 t)$, and the branching probability is chosen to be $b\varepsilon$ for some $b > 0$, then what is the scaling limit of the collection of branching-coalescing random walk paths as $\varepsilon \downarrow 0$?

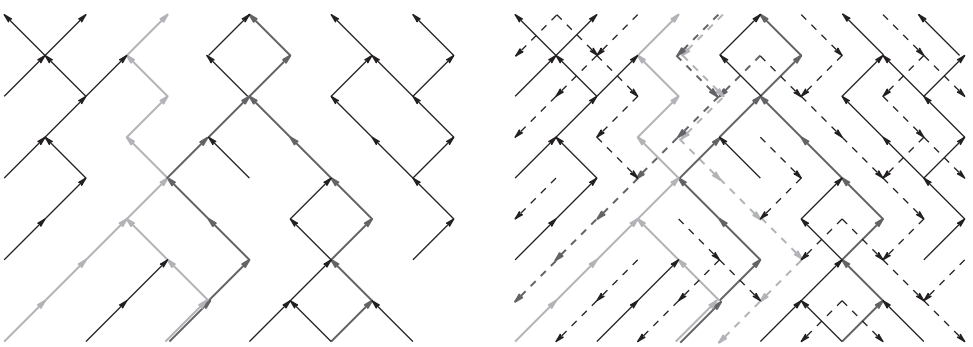

Figure 6.4. Discrete space-time branching-coalescing random walks on $\mathbb{Z}^2_{\text{even}}$, and its dual on $\mathbb{Z}^2_{\text{odd}}$.

The limit will be what we call the *Brownian net* $\mathcal{N}_b$ *with branching parameter* b. For simplicity, we will focus on the case $b = 1$, with the limit called the *standard Brownian net* $\mathcal{N}$. The fact that a nontrivial scaling limit exists with the above choice of the branching probability is hinted by the following observation. Instead of considering the collection of all random walk paths, let us first restrict our attention to two special subsets: the set of *leftmost*, respectively *rightmost*, random walk paths where the random walk always follows the arrow to the left, respectively right, whenever it encounters a branching point (see Figure 6.4). Note that the collection of leftmost paths is a collection of coalescing random walks with drift $-\varepsilon$, which ensures that each path under the diffusive scaling $S_\varepsilon(x,t) = (\varepsilon x, \varepsilon^2 t)$ converges to a Brownian motion with drift -1. Therefore we expect the collection of leftmost paths to converge to a variant of the Brownian web, $\mathcal{W}^{\mathrm{l}}$, which consists of coalescing Brownian motions with drift -1. Similarly, we expect the collection of rightmost paths to converge to a limit $\mathcal{W}^{\mathrm{r}}$, which consists of coalescing Brownian motions with drift $+1$. Of course, $\mathcal{W}^{\mathrm{l}}$ and $\mathcal{W}^{\mathrm{r}}$ are coupled in a nontrivial way.

The above observation explains the choice of the branching probability, and we see that any limit of the branching-coalescing random walk paths must contain the two coupled Brownian webs $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$. The questions that remain are:

(A) How to characterize the joint law of $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$?
(B) Can we construct the scaling limit of branching-coalescing random walks from $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$?

To answer (A), it suffices to characterize the joint distribution $l_{z_1} \in \mathcal{W}^{\mathrm{l}}(z_1), \ldots, l_{z_k} \in \mathcal{W}^{\mathrm{l}}(z_k)$ and $r_{z'_1} \in \mathcal{W}^{\mathrm{r}}(z'_1), \ldots, r_{z'_{k'}} \in \mathcal{W}^{\mathrm{r}}(z'_{k'})$ for a finite collection of $(z_i)_{1\le i\le k}$ and $(z'_i)_{1\le i\le k'}$ in $\mathbb{R}^2$. An examination of their discrete analogue suggests that:

- the paths $(l_{z_1}, \ldots, l_{z_k}, r_{z'_1}, \ldots, r_{z'_{k'}})$ evolve independently when they are apart;
- the *leftmost paths* $(l_{z_1}, \ldots, l_{z_k})$ coalesce when they meet, and the same is true for the *rightmost paths* $(r_{z'_1}, \ldots, r_{z'_{k'}})$;
- a pair of leftmost and rightmost paths $(l_{z_i}, r_{z'_j})$ solves the following pair of stochastic differential equations (SDEs):

$$\begin{aligned} \mathrm{d}L_t &= 1_{\{L_t \neq R_t\}} \mathrm{d}B^{\mathrm{l}}_t + 1_{\{L_t = R_t\}} \mathrm{d}B^{\mathrm{s}}_t - \mathrm{d}t, \\ \mathrm{d}R_t &= 1_{\{L_t \neq R_t\}} \mathrm{d}B^{\mathrm{r}}_t + 1_{\{L_t = R_t\}} \mathrm{d}B^{\mathrm{s}}_t + \mathrm{d}t, \end{aligned} \tag{6.11}$$

  where the leftmost path L and the rightmost path R are driven by independent Brownian motions B^{l} and B^{r} when they are apart, and driven by the same Brownian motion B^{s} (independent of B^{l} and B^{r}) when they coincide; furthermore, L and R are subject to the constraint that $L_t \le R_t$ for all $t \ge T := \inf\{u \ge \sigma_L \vee \sigma_R : L_u \le R_u\}$, with σ_L and σ_R being the starting times of L and R.

It turns out that the SDE (6.11) has a unique weak solution, and the above properties uniquely determine the joint law of $(l_{z_1}, \ldots, l_{z_k}, r_{z'_1}, \ldots, r_{z'_{k'}})$, which we will call *left-right coalescing Brownian motions*. Extending the starting points to a countable dense set in $\mathbb{R}^2$, and then taking closure of the resulting set of leftmost, respectively rightmost, paths, a.s. determines $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, which we will call the *left-right Brownian web*.

To answer (B), observe that in the discrete case, all random walk paths can be constructed by hopping back and forth between leftmost and rightmost random walk paths. This suggests a similar approach to construct the scaling limit of the set of all branching-coalescing random walk paths, which we will call *the Brownian net* $\mathcal{N}$. More precisely, to construct $\mathcal{N}$ from $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, we simply consider the set of all paths that can be obtained by *hopping* a finite number of times between paths in $\mathcal{W}^{\mathrm{l}}$ and $\mathcal{W}^{\mathrm{r}}$, and then take its closure.

The above considerations led to the original construction of the Brownian net $\mathcal{N}$ in [86], called the *hopping construction*.

From Figure 6.4, it is easily seen that the branching-coalescing random walks on $\mathbb{Z}^2_{\mathrm{even}}$ a.s. uniquely determine a dual collection of branching-coalescing random walks on $\mathbb{Z}^2_{\mathrm{odd}}$, running backward in time. Furthermore, the two systems are equally distributed apart from a rotation in space-time by 180^o and a lattice shift. Therefore in the scaling limit, we expect the left-right Brownian web $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ to have a dual $(\widehat{\mathcal{W}}^{\mathrm{l}}, \widehat{\mathcal{W}}^{\mathrm{r}})$, which determines a dual Brownain net $\widehat{\mathcal{N}}$. As for the Brownian web, such a duality provides a powerful tool. In particular, it leads to a second construction of the Brownian net, called the *wedge construction* in [86].

Besides the hopping and wedge constructions of the Brownian net, there are two more constructions, called the *mesh construction*, also developed in [86], and the *marking construction* developed by Newman, Ravishankar and Schertzer in [73], where the Brownian net was conceived independently from [86]. The mesh construction is based on the observation that, given the left-right Brownian web $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, there exist space-time regions (called *meshes*) with their left boundaries being rightmost paths, and their right boundaries being leftmost paths. Such unusual configurations make these meshes forbidden regions, where no paths can enter. The mesh construction asserts that the Brownian net consists of all paths which do not enter meshes.

In contrast to the hopping construction, which is an *outside-in* approach where the Brownian net is constructed from its outermost paths – the leftmost and rightmost paths – the *marking construction* developed in [73] is an *inside-out* approach, where one starts from a Brownian web and then constructs the Brownian net by adding branching points. In the discrete setting, this amounts to turning coalescing random walks into branching-coalescing random walks by changing

each lattice point independently into a branching point with probability ε. In the continuum setting, this turns out to require *Poisson marking* the set of $(1,2)$ points of the Brownian web (cf. Theorem 6.2.11) and turning them into branching points, so that the incoming Brownian web path can continue along either of the two outgoing Brownian web paths. We will introduce the marking construction in detail in Section 6.4, where we will study couplings between the Brownian web and net.

In the rest of this section, we will define the left-right Brownian web $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, give the hopping, wedge and mesh constructions of the Brownian net, and study various properties of the Brownian net, including the branching-coalescing point set, the backbone of the Brownian net, and special points of the Brownian net.

6.3.1 The Left-right Brownian Web and its Dual

The discussions above show that the key object in the construction of the Brownian net $\mathcal{N}$ is the left-right Brownian web $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, which should be the diffusive scaling limit of the collections of leftmost and rightmost branching-coalescing random walk paths with branching probability ε. In turn, the key ingredient in the construction of $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ is the pair of left-right SDEs in (6.11), which can be shown to be well-posed.

Proposition 6.3.1 (The left-right SDE) *For each initial state $(L_0, R_0) \in \mathbb{R}^2$, there exists a unique weak solution to the SDE* (6.11) *subject to the constraint that $L_t \leq R_t$ for all $t \geq T := \inf\{s \geq 0 : L_s = R_s\}$. Furthermore, almost surely, if $I := \{t \geq T : L_t = R_t\} \neq \emptyset$, then I is nowhere dense perfect set with positive Lebesgue measure.*

Proof Sketch. We sketch the basic idea and refer to [86, Proposition 2.1 and 3.1] for details. Assume w.l.o.g. that $L_0 = R_0 = 0$. Define

$$T_t := \int_0^t 1_{\{L_s < R_s\}} \mathrm{d}s, \qquad S_t := \int_0^t 1_{\{L_s = R_s\}} \mathrm{d}s,$$

and

$$\tilde{B}^{\mathrm{l}}_{T_t} := \int_0^t 1_{\{L_s < R_s\}} \mathrm{d}B^{\mathrm{l}}_s, \quad \tilde{B}^{\mathrm{r}}_{T_t} := \int_0^t 1_{\{L_s < R_s\}} \mathrm{d}B^{\mathrm{r}}_s, \quad \tilde{B}^{\mathrm{s}}_{S_t} := \int_0^t 1_{\{L_s = R_s\}} \mathrm{d}B^{\mathrm{s}}_s.$$

Then (L_t, R_t) solves

$$\begin{aligned} L_t &= \tilde{B}^{\mathrm{l}}_{T_t} + \tilde{B}^{\mathrm{s}}_{S_t} - t, \\ R_t &= \tilde{B}^{\mathrm{r}}_{T_t} + \tilde{B}^{\mathrm{s}}_{S_t} + t, \end{aligned} \tag{6.12}$$

and the difference $D_t := R_t - L_t$ satisfies

$$D_t = (\tilde{B}^{\mathrm{r}} - \tilde{B}^{\mathrm{r}})_{T_t} + 2t = \sqrt{2}\tilde{B}_{T_t} + 2T_t + 2S_t \tag{6.13}$$

with the constraint that D is nonnegative, where $\tilde{B} := \frac{1}{\sqrt{2}}(\tilde{B}^{\mathrm{r}} - \tilde{B}^{\mathrm{l}})$ is also a standard Brownian motion.

Since $T_t = \int_0^t 1_{\{L_s < R_s\}} \mathrm{d}s = \int_0^t 1_{\{D_s > 0\}} \mathrm{d}s$, it is easily seen that T_t must be continuous and strictly increasing in t. Therefore $\tau := T_t$ admits an inverse $T^{-1}\tau = t$. Rewriting the equation (6.13) for D with respect to the variable τ, the time D spent at the origin, we obtain

$$\tilde{D}_\tau := \frac{1}{\sqrt{2}} D_{T^{-1}\tau} = \tilde{B}_\tau + \sqrt{2}\tau + \sqrt{2}S_{T^{-1}\tau}, \tag{6.14}$$

where $\tilde{D}_\tau$ can be regarded as a transformation of $\tilde{B}_\tau + \sqrt{2}\tau$ by adding an increasing function $\tilde{S}_\tau := \sqrt{2}S_{T^{-1}\tau}$, which increases only when $\tilde{D}_\tau = 0$ such that $\tilde{D}_\tau$ stays nonnegative. Such an equation is known as a Skorohod equation, with $\tilde{D}_t$ being the Skorohod reflection of $\tilde{B}_\tau + \sqrt{2}\tau$ at the origin. Such a Skorohod equation admits a pathwise unique solution [59, Section 3.6.C], with

$$\tilde{S}_\tau := \sqrt{2}S_{T^{-1}\tau} = -\inf_{0 \le s \le \tau} (\tilde{B}_\tau + \sqrt{2}\tau),$$

which is in fact also the local time at the origin for the drifted Brownian motion $\tilde{B}_\tau + \sqrt{2}\tau$ reflected at the origin.

Having determined $S_{T^{-1}\tau}$ and $D_{T^{-1}\tau}$ almost surely from $\tilde{B} = \frac{1}{\sqrt{2}}(\tilde{B}^{\mathrm{r}} - \tilde{B}^{\mathrm{l}})$, to recover D_t, we only need to make a time change from τ back to $t := T_t + S_t = \tau + S_{T^{-1}\tau}$. Note that this time change has no effect when D is away from 0, but adds positive Lebesgue time when D is at 0. Therefore in contrast to $D_{T^{-1}\tau}$, which is a drifted Brownian motion reflected instantaneously at the origin, D_t is the same Brownian motion *sticky reflected* at the origin (see e.g., [102] and the references therein for further details on sticky reflected Brownian motions). Similarly, from $T_t = \tau$ and $S_t = S_{T^{-1}\tau}$, we can construct (L_t, R_t) in (6.12). From the same arguments, it is also easily seen that $I := \{t \ge 0 : L_t = R_t\}$ is almost surely a nowhere dense perfect set with positive Lebesgue measure. □

Having characterized the interaction of a single pair of leftmost and rightmost paths, we can now construct a collection of *left-right coalescing Brownian motions* $(l_{z_1}, \ldots, l_{z_k}, r_{z'_1}, \ldots, r_{z'_{k'}})$ with the properties that: (1) the paths evolve independently when they do not coincide; (2) $(l_{z_1}, \ldots, l_{z_k})$, respectively $(r_{z'_1}, \ldots, r_{z'_{k'}})$, is distributed as a collection of coalescing Brownian motions with drift -1, respectively $+1$; (3) every pair (l_{z_i}, r_{z_j}) is a weak solution to the left-right SDE (6.11). The construction can be carried out inductively. Assume w.l.o.g. that the paths all start at the same time. Then

- Let the paths evolve independently until the first time two paths meet.
- If this pair of paths are of the same type, then let them coalesce and iterate the construction with one path less than before.
- If this pair of paths are of different types, then let them evolve as a left-right pair solving the SDE (6.11), and let all other paths evolve independently, until the first time two paths (other than paths in the same left-right pair) meet.
- If the two meeting paths are of the same type, then let them coalesce and iterate the construction.
- If the two meeting paths are of different types, then let them form a left-right pair (breaking whatever pair relations they were in), and iterate the construction.

It is easily seen that this iterative construction terminates after a finite number of steps, when either a single path, or a single pair of leftmost and rightmost paths remains.

We are now ready to characterize the *left-right Brownian web* $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$.

Theorem 6.3.2 (The left-right Brownian web and its dual) *There exists an $\mathcal{H}^2$-valued random variable $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, called the (standard) left-right Brownian web, whose distribution is uniquely determined by the following properties:*

(i) *The left Brownian web $\mathcal{W}^{\mathrm{l}}$ (resp. right Brownian web $\mathcal{W}^{\mathrm{r}}$) is distributed as a Brownian web $\mathcal{W}$ tilted with drift -1 (resp. $+1$), i.e., $\mathcal{W}^{\mathrm{l}}$ (resp. $\mathcal{W}^{\mathrm{r}}$) has the same distribution as the image of $\mathcal{W}$ under the space-time transformation $(x,t) \to (x-t,t)$ (resp. $(x,t) \to (x+t,t)$).*

(ii) *For any finite deterministic set of points $z_1, \ldots, z_k, z'_1, \ldots, z'_{k'} \in \mathbb{R}^2$, the collection of paths $(l_{z_1}, \cdots, l_{z_k}; r_{z'_1}, \cdots, r_{z'_{k'}})$ is distributed as a family of left-right coalescing Brownian motions.*

Furthermore, almost surely there exists a dual left-right Brownian web $(\widehat{\mathcal{W}}^{\mathrm{l}}, \widehat{\mathcal{W}}^{\mathrm{r}}) \in \widehat{\mathcal{H}}^2$, such that $(\mathcal{W}^{\mathrm{l}}, \widehat{\mathcal{W}}^{\mathrm{l}})$ (resp. $(\mathcal{W}^{\mathrm{r}}, \widehat{\mathcal{W}}^{\mathrm{r}})$) is distributed as $(\mathcal{W}, \widehat{\mathcal{W}})$ tilted with drift -1 (resp. $+1$), and $(-\widehat{\mathcal{W}}^{\mathrm{l}}, -\widehat{\mathcal{W}}^{\mathrm{r}})$ has the same distribution as $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$.

Proof Sketch. The existence and uniqueness of a left-right Brownian web satisfying properties (i)–(ii) follow the same argument as for the Brownian web in Theorem 6.2.3. The almost sure existence of a dual left-right Brownian web follows from the duality of the Brownian web. The fact that $(\widehat{\mathcal{W}}^{\mathrm{l}}, \widehat{\mathcal{W}}^{\mathrm{r}})$ has the same distribution as $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, except for rotation around the origin in space-time by 180^o, can be derived by taking the diffusive scaling limits of their discrete counterparts, where such an equality in distribution is trivial (see [86, Theorem 5.3] for further details). □

Exercise 6.3.3 *Show that almost surely, no path in $\mathcal{W}^{\mathrm{l}}$ can cross any path in $\mathcal{W}^{\mathrm{r}}$ or $\widehat{\mathcal{W}}^{\mathrm{r}}$ from left to right, where π_1 is said to cross π_2 from left to right if there exist $s < t$ such that $\pi_1(s) < \pi_2(s)$ and $\pi_1(t) > \pi_2(t)$. Similarly, paths in $\mathcal{W}^{\mathrm{r}}$ cannot cross paths in $\mathcal{W}^{\mathrm{l}}$ or $\widehat{\mathcal{W}}^{\mathrm{l}}$ from right to left.*

6.3.2 The Hopping Construction of the Brownian Net

We are now ready to construct the Brownian net $\mathcal{N}$ by allowing paths to hop back and forth between paths in the left-right Brownian web $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, where a path π obtained by hopping from π_1 to π_2 at time t, with $\pi_1(t) = \pi_2(t)$, is defined by $\pi := \pi_1$ on $[\sigma_{\pi_1}, t]$ and $\pi := \pi_2$ on $[t, \infty)$.

Given π_1 and π_2 with $\pi_1(t) = \pi_2(t)$, which are in the scaling limit $\mathcal{N}$ of branching-coalescing random walk paths, it is not guaranteed that π obtained by hopping from π_1 to π_2 at time t is also in $\mathcal{N}$. Indeed, even though $\pi_1(t) = \pi_2(t)$, π_1 and π_2 may still arise as limits of random walk paths which do not meet, for which hopping is not possible. Therefore π has no approximating analogue among the branching-coalescing random walk paths, and hence π may not be in $\mathcal{N}$. One remedy is to allow hopping from π_1 to π_2 at time t only if the two paths *cross at time t*, i.e., $\sigma_{\pi_1}, \sigma_{\pi_2} < t$, there exist times $t^- < t^+$ with $(\pi_1(t^-) - \pi_2(t^-))(\pi_1(t^+) - \pi_2(t^+)) < 0$, and $t = \inf\{s \in (t^-, t^+) : (\pi_1(t^-) - \pi_2(t^-))(\pi_1(s) - \pi_2(s)) < 0\}$. We call t the *crossing time* between π_1 and π_2. If $\pi_1, \pi_2 \in \mathcal{N}$ and they cross at time t, then it is easily seen by discrete approximations that the path π obtained by hopping from π_1 to π_2 at time t must also be in $\mathcal{N}$.

Given a set of paths K, let $\mathcal{H}_{\mathrm{cross}}(K)$ denote the set of paths obtained by hopping a finite number of times among paths in K at crossing times. The Brownian net $\mathcal{N}$ can then be constructed by setting $\mathcal{N} := \overline{\mathcal{H}_{\mathrm{cross}}(\mathcal{W}^{\mathrm{l}} \cup \mathcal{W}^{\mathrm{r}})}$. Here is the hopping characterization of the Brownian net from [86, Theorem 1.3].

Theorem 6.3.4 (Hopping characterization of the Brownian net) *There exists an $(\mathcal{H}, \mathcal{B}_{\mathcal{H}})$-valued random variable $\mathcal{N}$, called the standard Brownian net, whose distribution is uniquely determined by the following properties:*

(i) *For each $z \in \mathbb{R}^2$, $\mathcal{N}(z)$ a.s. contains a unique left-most path l_z and right-most path r_z.*
(ii) *For any finite deterministic set of points $z_1, \ldots, z_k, z'_1, \ldots, z'_{k'} \in \mathbb{R}^2$, the collection of paths $(l_{z_1}, \ldots, l_{z_k}, r_{z'_1}, \ldots, r_{z'_{k'}})$ is distributed as a family of left-right coalescing Brownian motions.*
(iii) *For any deterministic countable dense sets $\mathcal{D}^{\mathrm{l}}, \mathcal{D}^{\mathrm{r}} \subset \mathbb{R}^2$,*

$$\mathcal{N} = \overline{\mathcal{H}_{\mathrm{cross}}(\{l_z : z \in \mathcal{D}^{\mathrm{l}}\} \cup \{r_z : z \in \mathcal{D}^{\mathrm{r}}\})} \quad \text{a.s.} \tag{6.15}$$

Proof. The uniqueness in law of a random variable $\mathcal{N}$ satisfying the above properties is easily verified by the same argument as that for the Brownian web in Theorem 6.2.3. For existence, we can just define $\mathcal{N} := \overline{\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})}$, where $(\mathcal{W}^{\text{l}}, \mathcal{W}^{\text{r}})$ is the standard left-right Brownian web.

To show that $\mathcal{N}$ satisfies properties (i)–(iii), first note that $\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})$ satisfies properties (i)–(ii), where for each deterministic z, $\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})$ contains a leftmost element and a rightmost element, which are just $l_z \in \mathcal{W}^{\text{l}}(z)$ and $r_z \in \mathcal{W}^{\text{r}}(z)$. We claim that taking closure of $\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})$ does not change the leftmost and rightmost element starting from any given $z = (x, t)$.

Indeed, if $\mathcal{N}(z)$ contains any path π with $\pi(t') < l_z(t') - \varepsilon$ for some $t' > t$ and $\varepsilon > 0$, then there exists a sequence $\pi_n \in \mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})$ starting from z_n, such that $\pi_n \to \pi$ and $\pi_n(t') \le l_z(t') - \varepsilon$ for all n large. Since $l_{z_n} \in \mathcal{W}^{\text{l}}(z_n)$ is the leftmost path in $\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})$ starting from z_n, we have $l_{z_n} \le \pi_n$, and hence also $l_{z_n}(t') \le l_z(t') - \varepsilon$ for all n large. However, this is impossible because $l_{z_n} \to l_z$, and hence the time of coalescence between l_{z_n} and l_z tends to t as $n \to \infty$ by Corollary 6.2.8.

We can therefore conclude that $\mathcal{N} := \overline{\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})}$ also satisfies properties (i)–(ii). Since any path in $\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}}$ can be approximated by paths in $\mathcal{W}^{\text{l}}(\mathcal{D}^{\text{l}}) \cup \mathcal{W}^{\text{r}}(\mathcal{D}^{\text{r}})$ in the strong sense as in Corollary 6.2.8, it is easily seen that $\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}}) \subset \overline{\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}}(\mathcal{D}^{\text{l}}) \cup \mathcal{W}^{\text{r}}(\mathcal{D}^{\text{r}}))}$, and hence $\mathcal{N}$ also satisfies property (iii).

What we have left out in the proof is the a.s. precompactness of $\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})$, which is needed for $\mathcal{N}$ to qualify as an $(\mathcal{H}, \mathcal{B}_{\mathcal{H}})$-valued random variable. We leave this as an exercise. □

Exercise 6.3.5 *Show that almost surely, $\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})$ is precompact.*

Exercise 6.3.6 *Show that a.s., no path in $\mathcal{N}$ can cross any path in $\mathcal{W}^{\text{r}}$ or $\widehat{\mathcal{W}}^{\text{r}}$ from left to right, or cross any path in $\mathcal{W}^{\text{l}}$ or $\widehat{\mathcal{W}}^{\text{l}}$ from right to left, where the definition of crossing is, as in Exercise 6.3.3.*

Remark 6.3.7 Of course, we need to justify that the Brownian net characterized in Theorem 6.3.4 is indeed the scaling limit of branching-coalescing random walks, as motivated at the start of Section 6.3. Such a convergence result was established in [86, Section 5.3].

6.3.3 The Wedge Construction of the Brownian Net

The wedge and mesh constructions of the Brownian net $\mathcal{N}$ are both based on the observation that there are certain forbidden regions in space-time where Brownian net paths cannot enter. It turns out that the Brownian net can also be characterized as the set of paths that do not enter these forbidden regions. In the wedge construction, these forbidden regions, called *wedges*, are defined from the dual

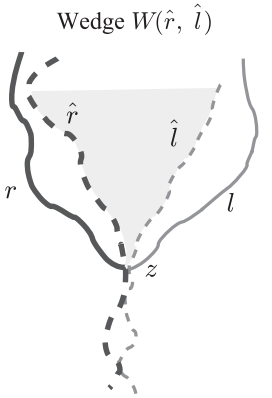

Figure 6.5. A wedge $W(\hat{r},\hat{l})$ with bottom point z.

left-right Brownian web $(\widehat{\mathcal{W}}^{\mathrm{l}},\widehat{\mathcal{W}}^{\mathrm{r}})$, while in the mesh construction, these forbidden regions, called *meshes*, are defined from the left-right Brownian web $(\mathcal{W}^{\mathrm{l}},\mathcal{W}^{\mathrm{r}})$.

Definition 6.3.8 (Wedges) *Let $(\mathcal{W}^{\mathrm{l}},\mathcal{W}^{\mathrm{r}},\widehat{\mathcal{W}}^{\mathrm{l}},\widehat{\mathcal{W}}^{\mathrm{r}})$ be the standard left-right Brownian web and its dual. For any $\hat{r}\in\widehat{\mathcal{W}}^{\mathrm{r}}$ and $\hat{l}\in\widehat{\mathcal{W}}^{\mathrm{l}}$ that are ordered with $\hat{r}(s)<\hat{l}(s)$ at the time $s:=\hat{\sigma}_{\hat{r}}\wedge\hat{\sigma}_{\hat{l}}$, let $T:=\sup\{t<s:\hat{r}(t)=\hat{l}(t)\}$ (possibly equals $-\infty$) be the first hitting time of $\hat{r}$ and $\hat{l}$. We call the open set (see Figure 6.5)*

$$W=W(\hat{r},\hat{l}):=\{(x,u)\in\mathbb{R}^2:T<u<s,\ \hat{r}(u)<x<\hat{l}(u)\} \tag{6.16}$$

a wedge *of $(\widehat{\mathcal{W}}^{\mathrm{l}},\widehat{\mathcal{W}}^{\mathrm{r}})$ with left and right boundary $\hat{r}$ and $\hat{l}$ and bottom point $z:=(\hat{r}(T),T)=(\hat{l}(T),T)$. A path π is said to* enter W from outside *if there exist $\sigma_\pi\le s<t$ such that $(\pi(s),s)\notin\overline{W}$ and $(\pi(t),t)\in W$.*

Theorem 6.3.9 (Wedge characterization of the Brownian net) *Let $(\mathcal{W}^{\mathrm{l}},\mathcal{W}^{\mathrm{r}},\widehat{\mathcal{W}}^{\mathrm{l}},\widehat{\mathcal{W}}^{\mathrm{r}})$ be the standard left-right Brownian web and its dual. Then almost surely,*

$$\mathcal{N}=\{\pi\in\Pi:\pi\text{ does not enter any wedge of }(\widehat{\mathcal{W}}^{\mathrm{l}},\widehat{\mathcal{W}}^{\mathrm{r}})\text{ from outside}\} \tag{6.17}$$

is the standard Brownian net associated with $(\mathcal{W}^{\mathrm{l}},\mathcal{W}^{\mathrm{r}})$, i.e., $\mathcal{N}=\overline{\mathcal{H}_{\mathrm{cross}}(\mathcal{W}^{\mathrm{l}}\cup\mathcal{W}^{\mathrm{r}})}$.

Remark 6.3.10 The wedge characterization can also be applied to the Brownian web $\mathcal{W}$ with both $\widehat{\mathcal{W}}^{\mathrm{l}}$ and $\widehat{\mathcal{W}}^{\mathrm{r}}$ replaced by the dual Brownian web $\widehat{\mathcal{W}}$. Indeed, $\mathcal{W}$ can be seen as a degenerate Brownian net $\mathcal{N}_b$ with branching parameter $b=0$, where $\mathcal{N}_b$ can be constructed in the same way as the standard Brownian net $\mathcal{N}$ with $b=1$, except that the left-right coalescing Brownian motions in $(\mathcal{W}^{\mathrm{l}},\mathcal{W}^{\mathrm{r}})$ now have drift $\mp b$ respectively. For $\mathcal{W}$, the wedge characterization is stronger than requiring paths not to cross any path in $\widehat{\mathcal{W}}$ (cf. Remark 6.2.6), because it also prevents paths from entering a wedge from outside through its bottom point.

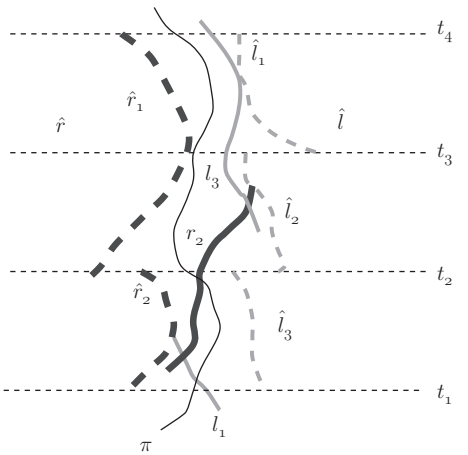

Figure 6.6. Steering a hopping path in the 'fish-trap' $(\hat{r},\hat{l})$.

Proof of Theorem 6.3.9. First we show that no path in $\overline{\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})}$ can enter any wedge of $(\widehat{\mathcal{W}}^{\text{l}}, \widehat{\mathcal{W}}^{\text{r}})$ from outside. If this is false, then there must be some path $\pi \in \mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})$ which enters a wedge $W(\hat{r},\hat{l})$ from outside. There are two possibilities: either π enters W from outside by crossing one of its two boundaries, which is impossible by Exercise 6.3.6; or π enters W from outside through its bottom point z. However, by the same argument as why a point of coalescence between two dual Brownian web paths cannot be hit by a forward Brownian web path (cf. Theorem 6.2.11), no path in $\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}}$ can enter the bottom point of a wedge W, and hence neither can any path in $\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})$. This verifies the desired inclusion.

We now show the converse inclusion that any path not entering wedges from outside must be in $\overline{\mathcal{H}_{\text{cross}}(\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}})}$. Let π be such a path. The strategy to approximate π by hopping paths is illustrated in Figure 6.6.

To approximate π in a given time interval, say $[s,t]$, we first partition $[s,t]$ into sub-intervals of equal length, $[t_{i-1},t_i]$, for $1 \le i \le N$. Fix an $\varepsilon > 0$. From the top time t_N, we consider the wedge formed by $\hat{r}_1 \in \widehat{\mathcal{W}}^{\text{r}}$ and $\hat{l}_1 \in \widehat{\mathcal{W}}^{\text{l}}$ starting respectively at $\pi(t_N) - \varepsilon$ and $\pi(t_N) + \varepsilon$ at time t_N. Note that $\hat{r}_1$ and $\hat{l}_1$ cannot meet during the interval $[t_{N-1},t_N]$, otherwise π would be entering the wedge $W(\hat{r}_1,\hat{l}_1)$ from outside. At time t_{N-1}, we check whether $\hat{r}_1 < \pi(t_{N-1}) - \varepsilon$, and if it is the case, then we start $\hat{r}_2$ at $\pi(t_{N-1}) - \varepsilon$ at time t_{N-1}. Similarly we check whether $\hat{l}_1 > \pi(t_{N-1}) + \varepsilon$, and if it is the case, then start $\hat{l}_2$ at $\pi(t_{N-1}) + \varepsilon$ at time t_{N-1}. In any event, we still have a pair of dual left-right paths which enclose π on the interval $[t_{N-2},t_{N-1}]$, which are within ε distance of π at the top time t_{N-1}. This procedure is iterated until the time interval $[t_0,t_1]$, and it constructs a 'fish-trap' of

dual left-right paths, which can now be used to construct a forward hopping path that stays inside the 'fish-trap'.

Indeed, a forward path $l \in \mathcal{W}^{\mathrm{l}}$ cannot cross dual paths $\hat{l}_i \in \widehat{\mathcal{W}^{\mathrm{l}}}$, which form the right boundary of the 'fish-trap'. When l hits the left-boundary of the 'fish-trap', we can then hop to a path $r \in \mathcal{W}^{\mathrm{r}}$ until it hits the right-boundary of the 'fish-trap'. Iterating this procedure then constructs a hopping path that stays inside the 'fish-trap'. The almost sure equicontinuity of paths in $\mathcal{W}^{\mathrm{l}} \cup \mathcal{W}^{\mathrm{r}} \cup \widehat{\mathcal{W}^{\mathrm{l}}} \cup \widehat{\mathcal{W}^{\mathrm{r}}}$ ensures that only a finite number of hoppings is needed to reach time t_N, and the supnorm distance on the interval $[s,t]$ between π and any path inside the 'fish-trap' can be made arbitrarily small by choosing N large and ε small. Therefore π can be approximated arbitrarily well by paths in $\mathcal{H}_{\mathrm{cross}}(\mathcal{W}^{\mathrm{l}} \cup \mathcal{W}^{\mathrm{r}})$. For further details, see [86, Lemma 4.7]. □

6.3.4 The Mesh Construction of the Brownian Net

Definition 6.3.11 (Meshes) *Let $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ be the standard left-right Brownian web. If for a given $z = (x,t) \in \mathbb{R}^2$, there exist $l \in \mathcal{W}^{\mathrm{l}}(z)$ and $r \in \mathcal{W}^{\mathrm{r}}(z)$ such that $r(s) < l(s)$ on $(t, t+\varepsilon)$ for some $\varepsilon > 0$, then denoting $T := \inf\{s > t : r(s) = l(s)\}$, we call the open set (see Figure 6.7)*

$$M = M(r,l) := \{(y,s) \in \mathbb{R}^2 : t < s < T,\ r(s) < y < l(s)\} \tag{6.18}$$

a mesh *of $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ with left and right boundaries r and l and bottom point z. A path π is said to* enter *M if there exist $\sigma_\pi < s < t$ such that $(\pi(s), s) \notin M$ and $(\pi(t), t) \in M$.*

Theorem 6.3.12 (Mesh characterization of the Brownian net) *Let $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ be the standard left-right Brownian web. Then almost surely,*

$$\mathcal{N} = \{\pi \in \Pi : \pi \text{ does not enter any mesh of } (\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})\} \tag{6.19}$$

is the standard Brownian net associated with $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, i.e., $\mathcal{N} = \overline{\mathcal{H}_{\mathrm{cross}}(\mathcal{W}^{\mathrm{l}} \cup \mathcal{W}^{\mathrm{r}})}$.

Remark 6.3.13 The mesh characterization can also be applied to the Brownian web $\mathcal{W}$, where the bottom point of the mesh must be of either type $(0,2)$, $(1,2)$, or $(0,3)$ in Theorem 6.2.11.

We note that there is a subtle difference between *a path entering a mesh M from outside*, vs *a path entering a mesh M*. In particular, a path entering M (but not entering M from outside) could start inside $\overline{M}$, hit the boundary of M at a later time, and then move inside M. The heart of the proof of Theorem 6.3.12 consists in ruling out such scenarios for Brownian net paths, for which meshes

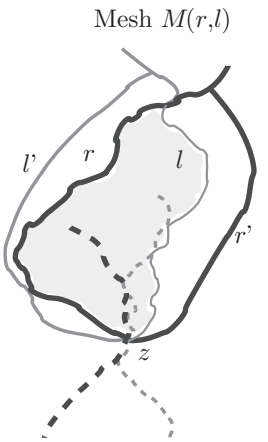

Figure 6.7. A mesh $M(r,l)$ with bottom point z.

play an essential role. As a by-product, one can show the following result (see [86, Proposition 1.8]), which is stronger than the assertions of Exercise 6.3.6.

Proposition 6.3.14 (Containment by left-most and right-most paths) *Let* $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ *be the standard left-right Brownian web, and let* $\mathcal{N}$ *be the Brownian net associated with it. Then almost surely, there exist no* $\pi \in \mathcal{N}$ *and* $l \in \mathcal{W}^{\mathrm{l}}$ *such that* $l(s) \le \pi(s)$ *and* $\pi(t) < l(t)$ *for some* $\sigma_\pi \vee \sigma_l < s < t$. *An analogue statement holds for right-most paths.*

The proofs of Proposition 6.3.14 and Theorem 6.3.12 are fairly involved and we refer the details to [86, Theorem 1.7 and Proposition 1.8].

The wedge and mesh characterizations of the Brownian net have the following interesting corollary.

Proposition 6.3.15 (Brownian net is closed under hopping) *Let* $\mathcal{N}$ *be the standard Brownian net. Then:*

(i) *Almost surely for any* $\pi_1, \pi_2 \in \mathcal{N}$ *with* $\pi_1(t) = \pi_2(t)$ *and* $\sigma_{\pi_1}, \sigma_{\pi_2} < t$ *for some* t, *the path* π *defined by* $\pi := \pi_1$ *on* $[\sigma_{\pi_1}, t]$ *and* $\pi := \pi_2$ *on* $[t, \infty)$ *is in* $\mathcal{N}$.

(ii) *For any deterministic* t, *almost surely for any* $\pi_1, \pi_2 \in \mathcal{N}$ *with* $\pi_1(t) = \pi_2(t)$ *and* $\sigma_{\pi_1}, \sigma_{\pi_2} \le t$, *the path* π *defined by* $\pi := \pi_1$ *on* $[\sigma_{\pi_1}, t]$ *and* $\pi := \pi_2$ *on* $[t, \infty)$ *is in* $\mathcal{N}$.

Exercise 6.3.16 *Use the mesh characterization to prove Proposition 6.3.15 (i), and the use then wedge characterization together with Proposition 6.3.18 below to prove Proposition 6.3.15 (ii).*

Here is an even more striking corollary of the mesh characterization [86, Proposition 1.13].

Proposition 6.3.17 (Image set property) *Let $\mathcal{N}$ be the Brownian net. For $T \in [-\infty,\infty)$, let $\mathcal{N}_T := \{\pi \in \mathcal{N} : \sigma_\pi = T\}$ and $N_T := \{(\pi(t),t) \in \mathbb{R}^2 : \pi \in \mathcal{N}_T, t \in (T,\infty)\}$. Then almost surely for any $T \in [-\infty,\infty)$, any path $\pi \in \Pi$ with $\sigma_\pi = T$ and $\{(\pi(t),t) : t > T\} \subset N_T$ is a path in $\mathcal{N}$.*

Note that for the Brownian web, this property is easily seen to hold.

6.3.5 The Branching-coalescing Point Set

Similar to the definition of the coalescing point set from the Brownian web in Section 6.2.4, we can define the so-called *branching-coalescing point set* from the Brownian net $\mathcal{N}$ as follows.

Given the Brownian net $\mathcal{N}$, and a closed set $A \subset \mathbb{R}$, define the *branching-coalescing point set* by

$$\xi_t^A := \{y \in \mathbb{R} : y = \pi(t) \text{ for some } \pi \in \mathcal{N}(A \times \{0\})\}, \qquad t \geq 0. \tag{6.20}$$

In other words, ξ_t^A is the set of points in $\mathbb{R}$ that lie on some path in $\mathcal{N}$ that start from A at time 0.

Using the wedge characterization of the Brownian net, we can compute the density of $\xi_t^{\mathbb{R}}$. As $t \downarrow 0$, we see in Proposition 6.3.18 below that the density diverges at the same rate $1/\sqrt{\pi t}$ as for the coalescing point set in Proposition 6.2.7, which indicates that coalescence plays the dominant role for small times, while the density converges to the constant 2 as $t \uparrow \infty$, which results from the balance between branching and coalescence for large times.

Proposition 6.3.18 (Density of branching-coalescing point set) *Let $\xi_\cdot^{\mathbb{R}}$ be the branching-coalescing point set defined from the Brownian net $\mathcal{N}$ as in* (6.20)*. Then for all $t > 0$ and $a < b$,*

$$\mathbb{E}[|\xi_t^{\mathbb{R}} \cap [a,b]|] = (b-a)\Big(\frac{e^{-t}}{\sqrt{\pi t}} + 2\Phi(\sqrt{2t})\Big), \tag{6.21}$$

where $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-\frac{y^2}{2}} \mathrm{d}y$.

Exercise 6.3.19 *Prove Proposition 6.3.18 by adapting the proof of Proposition 6.2.7 and showing that, almost surely, $\xi_t^{\mathbb{R}} \cap (a,b) \neq \emptyset$ if and only if $\hat{r} \in \widehat{\mathcal{W}}^{\mathrm{r}}$ starting at (a,t), and $\hat{l} \in \widehat{\mathcal{W}}^{\mathrm{l}}$ starting at (b,t), do not meet above time 0.*

Remark 6.3.20 Surprisingly, we can even identify the law of $\xi_t^{\mathbb{R}}$ as $t \uparrow \infty$, which is a Poisson point process on $\mathbb{R}$ with intensity 2. Furthermore, $\xi_\cdot$ is reversible with

respect to the law of the Poisson point process on $\mathbb{R}$ with intensity 2. Formulated in terms of the Brownian net, this amounts to the statement that $\mathcal{N}(*,-\infty)$, the collection of Brownian net paths started at time $-\infty$ and called the *backbone of the Brownian net* in [86], has the same distribution as $-\mathcal{N}(*,-\infty)$, the set of paths obtained by reflecting the graph of each path in $\mathcal{N}(*,-\infty)$ across the origin in space-time. These results were established in [86, Section 9] by first observing their analogues for a discrete system of branching-coalescing random walks, and then passing to the continuum limit.

Remark 6.3.21 Proposition 6.3.18 implies that for each $t > 0$, almost surely $\xi_t^{\mathbb{R}}$ is a locally finite point set. However, almost surely there exists a dense set of random times at which $\xi_t^{\mathbb{R}}$ contains no isolated point, and is in particular uncountable (see [91, Proposition 3.14]).

Remark 6.3.22 As noted in Remark 6.2.10, started from the whole real line, the coalescing point set forms a Pfaffian point process at each time $t > 0$. It will be interesting to investigate whether the branching-coalescing point set $\xi_t^{\mathbb{R}}$ also admits an explicit characterization as a Pfaffian point process.

6.3.6 Special Points of the Brownian Net

Similar to the classification of special points for the Brownian web formulated in Theorem 6.2.11, we can give an almost sure classification of all points in $\mathbb{R}^2$ according to the configuration of paths in the Brownian net entering and leaving the point. Such an analysis was carried out in [91], where it was shown that a.s. there are 20 types of points, in contrast to the 7 types for the Brownian web.

Since Brownian net paths must be contained between paths in the left Brownian web and right Brownian web, as stated in Proposition 6.3.14, the classification of special points is in fact carried out mainly for $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$. First we introduce a notion of equivalence between paths entering and leaving a point, which is weaker than that introduced for the Brownian web in Section 6.2.5.

Definition 6.3.23 (Equivalence of paths entering and leaving a point) *We say $\pi_1, \pi_2 \in \Pi$ are equivalent paths entering $z = (x,t) \in \mathbb{R}^2$, denoted by $\pi_1 \sim^z_{\mathrm{in}} \pi_2$, if π_1 and π_2 enter z and $\pi_1(t-\varepsilon_n) = \pi_2(t-\varepsilon_n)$ for a sequence $\varepsilon_n \downarrow 0$. We say π_1, π_2 are equivalent paths leaving z, denoted by $\pi_1 \sim^z_{\mathrm{out}} \pi_2$, if π_1 and π_2 leave z and $\pi_1(t+\varepsilon_n) = \pi_2(t+\varepsilon_n)$ for a sequence $\varepsilon_n \downarrow 0$.*

When applied to paths in the Brownian web, the above notion of equivalence implies the equivalence introduced in Section 6.2.5, which is why we have abused the notation and used the same symbols $\sim^z_{\mathrm{in}}$ and $\sim^z_{\mathrm{out}}$. Although $\sim^z_{\mathrm{in}}$ and $\sim^z_{\mathrm{out}}$ are not equivalence relations on the space of all paths Π, they are easily seen to be

equivalence relations on $\mathcal{W}^{\mathrm{l}} \cup \mathcal{W}^{\mathrm{r}}$. This allows us to classify z according to the equivalence classes of paths entering and leaving z, which are necessarily ordered.

To denote the type of a point, we first list the incoming equivalence classes of paths from left to right, and then, separated by a comma, the outgoing equivalence classes of paths from left to right. If an equivalence class contains only paths in $\mathcal{W}^{\mathrm{l}}$, resp. $\mathcal{W}^{\mathrm{r}}$, we will label it by l, resp. r, while if it contains both paths in $\mathcal{W}^{\mathrm{l}}$ and in $\mathcal{W}^{\mathrm{r}}$, we will label it by p, standing for pair. For points with (up to equivalence) one incoming and two outgoing paths, a subscript l, resp. r, means that all incoming paths belong to the left one, resp. right one, of the two outgoing equivalence classes; a subscript s indicates that incoming paths in $\mathcal{W}^{\mathrm{l}}$ belong to the left outgoing equivalence class, while incoming paths in $\mathcal{W}^{\mathrm{r}}$ belong to the right outgoing equivalence class. If at a point there are no incoming paths in $\mathcal{W}^{\mathrm{l}} \cup \mathcal{W}^{\mathrm{r}}$, then we denote this by o or n, where o indicates that there are no incoming paths in the Brownian net $\mathcal{N}$, while n indicates that there are incoming paths in $\mathcal{N}$ (but none in $\mathcal{W}^{\mathrm{l}} \cup \mathcal{W}^{\mathrm{r}}$).

For example, a point is of type $(\mathrm{p}, \mathrm{lp})_{\mathrm{r}}$ if at this point there is one equivalence class of incoming paths in $\mathcal{W}^{\mathrm{l}} \cup \mathcal{W}^{\mathrm{r}}$ and there are two outgoing equivalence classes. The incoming equivalence class is of type p while the outgoing equivalence classes are of type l and p, from left to right. All incoming paths in $\mathcal{W}^{\mathrm{l}} \cup \mathcal{W}^{\mathrm{r}}$ continue as paths in the outgoing equivalence class of type p.

Since the dual left-right Brownian web $(\widehat{\mathcal{W}}^{\mathrm{l}}, \widehat{\mathcal{W}}^{\mathrm{r}})$ can be used to define a dual Brownian net $\widehat{\mathcal{N}}$, the type of z w.r.t. $(\widehat{\mathcal{W}}^{\mathrm{l}}, \widehat{\mathcal{W}}^{\mathrm{r}})$ and $\widehat{\mathcal{N}}$ can be defined in the same way, after rotating their graphs in $\mathbb{R}^2$ by 180^o around the origin.

We now state the classification result from [91, Theorem 1.7], while omitting its proof.

Theorem 6.3.24 (Special points of the Brownian net) *Let $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ be the standard left-right Brownian web, let $(\widehat{\mathcal{W}}^{\mathrm{l}}, \widehat{\mathcal{W}}^{\mathrm{r}})$ be its dual, and let $\mathcal{N}$ and $\widehat{\mathcal{N}}$ be the associated Brownian net and its dual. Then almost surely, each point in $\mathbb{R}^2$ is of one of the following 20 types in $\mathcal{N}/\widehat{\mathcal{N}}$:*

(1) $(\mathrm{o},\mathrm{p})/(\mathrm{o},\mathrm{p})$, $(\mathrm{o},\mathrm{pp})/(\mathrm{p},\mathrm{p})$, $(\mathrm{p},\mathrm{p})/(\mathrm{o},\mathrm{pp})$, $(\mathrm{o},\mathrm{ppp})/(\mathrm{pp},\mathrm{p})$, $(\mathrm{pp},\mathrm{p})/(\mathrm{o},\mathrm{ppp})$, $(\mathrm{p},\mathrm{pp})_{\mathrm{l}}/(\mathrm{p},\mathrm{pp})_{\mathrm{l}}$, $(\mathrm{p},\mathrm{pp})_{\mathrm{r}}/(\mathrm{p},\mathrm{pp})_{\mathrm{r}}$;
(2) $(\mathrm{p},\mathrm{pp})_{\mathrm{s}}/(\mathrm{p},\mathrm{pp})_{\mathrm{s}}$, called *separation points;*
(3) $(\mathrm{l},\mathrm{p})/(\mathrm{o},\mathrm{lp}), (\mathrm{o},\mathrm{lp})/(\mathrm{l},\mathrm{p}), (\mathrm{r},\mathrm{p})/(\mathrm{o},\mathrm{pr}), (\mathrm{o},\mathrm{pr})/(\mathrm{r},\mathrm{p})$;
(4) $(\mathrm{l},\mathrm{pp})_{\mathrm{r}}/(\mathrm{p},\mathrm{lp})_{\mathrm{r}}, (\mathrm{p},\mathrm{lp})_{\mathrm{r}}/(\mathrm{l},\mathrm{pp})_{\mathrm{r}}, (\mathrm{r},\mathrm{pp})_{\mathrm{l}}/(\mathrm{p},\mathrm{pr})_{\mathrm{l}}, (\mathrm{p},\mathrm{pr})_{\mathrm{l}}/(\mathrm{r},\mathrm{pp})_{\mathrm{l}}$;
(5) $(\mathrm{l},\mathrm{lp})\mathrm{r}/(\mathrm{l},\mathrm{lp})_{\mathrm{r}}, (\mathrm{r},\mathrm{pr})\mathrm{l}/(\mathrm{r},\mathrm{pr})_{\mathrm{l}}$;
(6) $(\mathrm{n},\mathrm{p})/(\mathrm{o},\mathrm{lr}), (\mathrm{o},\mathrm{lr})/(\mathrm{n},\mathrm{p})$;

and all of these types occur. For each deterministic time $t \in \mathbb{R}$, almost surely, each point in $\mathbb{R} \times \{t\}$ is of either type $(\mathrm{o},\mathrm{p})/(\mathrm{o},\mathrm{p}), (\mathrm{o},\mathrm{pp})/(\mathrm{p},\mathrm{p})$, *or* $(\mathrm{p},\mathrm{p})/(\mathrm{o},\mathrm{pp})$, *and*

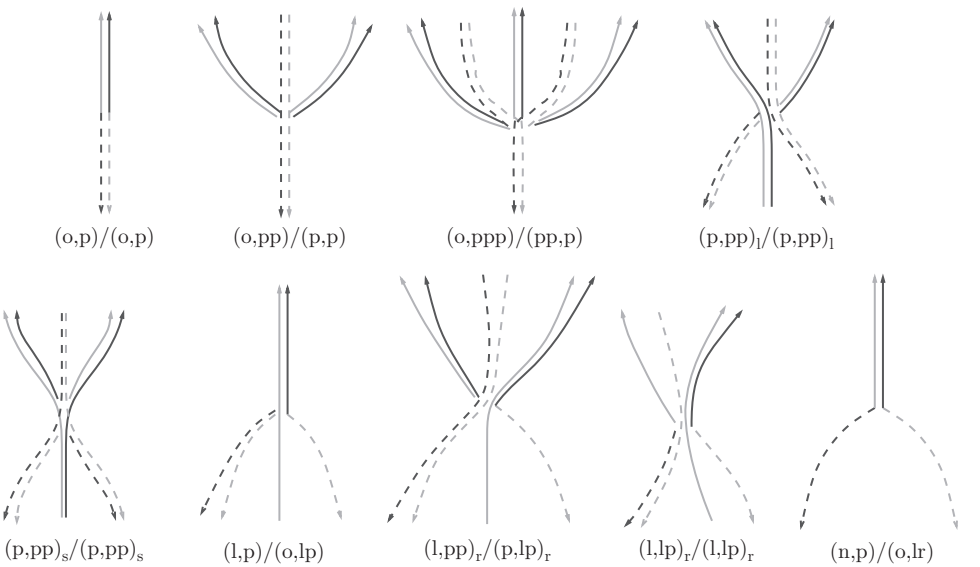

Figure 6.8. Special points of the Brownian net, modulo symmetry.

all of these types occur. A deterministic point $(x,t) \in \mathbb{R}^2$ *is almost surely of type* (o,p)/(o,p).

Remark 6.3.25 Note that the points listed in item (1) are analogues of the seven types of points of the Brownian web in Figure 6.3, where a path of the Brownian web is replaced by a pair of paths in $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$. Modulo symmetry, this gives rise to four distinct types of points. The types of points listed within each item from (2)–(6) are related to each other by symmetry. Therefore modulo symmetry, there are nine types of special points for the Brownian net, as illustrated in Figure 6.8.

Remark 6.3.26 A basic ingredient in the proof of Theorem 6.3.24 is the characterization of the interaction between paths in $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ and paths in $(\widehat{\mathcal{W}}^{\mathrm{l}}, \widehat{\mathcal{W}}^{\mathrm{r}})$. The interaction between paths in $\mathcal{W}^{\mathrm{l}}$ and $\widehat{\mathcal{W}}^{\mathrm{l}}$, and similar between paths in $\mathcal{W}^{\mathrm{r}}$ and $\widehat{\mathcal{W}}^{\mathrm{r}}$, is given by Skorohod reflection, as mentioned in Remark 6.2.5. It turns out that paths in $\mathcal{W}^{\mathrm{r}}$ interact with paths in $\widehat{\mathcal{W}}^{\mathrm{l}}$ also via Skorohod reflection, except that given $\hat{l} \in \widehat{\mathcal{W}}^{\mathrm{l}}$, if $r \in \mathcal{W}^{\mathrm{r}}$ initially starts on the left of $\hat{l}$, then it is Skorohod reflected to the left of $\hat{l}$ until the collision local time between r and $\hat{l}$ exceeds an independent exponential random variable, at which time r crosses over to the right of $\hat{l}$ and is Skorohod reflected to the right of $\hat{l}$ from that time on (see [91, Lemma 2.1]).

Remark 6.3.27 Although Theorem 6.3.24 classifies points in $\mathbb{R}^2$ mostly according to the configuration of paths in the left-right Brownian web, it can be used to

deduce the configuration of paths in the Brownian net entering and leaving each type of point. It turns out that $\sim^z_{\text{in}}$ is in fact an equivalence relation among paths in $\mathcal{N}$, and the same holds for $\sim^z_{\text{out}}$ for z of any type other than (o, lr). In particular, when z is not of type (o,lr), each equivalence class of paths in $\mathcal{W}^{\text{l}} \cup \mathcal{W}^{\text{r}}$ entering or leaving z can be enlarged to an equivalence class of paths in $\mathcal{N}$. In particular, by Proposition 6.3.14, all Brownian net paths in an equivalence class of type p, which contains a pair of equivalent paths $(l, r) \in (\mathcal{W}^{\text{l}}, \mathcal{W}^{\text{r}})$, must be bounded between l and r when sufficiently close to z. This applies in particular to points of type $(\text{p}, \text{pp})_{\text{s}}$, the *separation points*, as well as points of type (pp,p), which we call the *meeting points*. See [91, Section 1.4] for further details.

Exercise 6.3.28 *Show that the set of separation points of $\mathcal{N}$ is a.s. countable and dense in $\mathbb{R}^2$.*

6.4 Coupling the Brownian Web and Net

We now introduce a coupling between the Brownian web $\mathcal{W}$ and the Brownian net $\mathcal{N}$, which is again best motivated from their discrete analogues, the coalescing and branching-coalescing random walks. In particular, this gives the fourth construction of the Brownian net, the *marking construction* mentioned before Section 6.3.1, which was developed by Newman, Ravishankar and Schertzer in [73].

Given a realization of branching-coalescing random walks as illustrated in Figure 6.9 (a), where each lattice point in $\mathbb{Z}^2_{\text{even}}$ is a branching point with probability ε, we can construct a collection of coalescing random walks by simply forcing the random walk to go either left or right with probability $1/2$ each, independently at each branching point. Interestingly, these branching points have analogues in the continuum limit $\mathcal{N}$, which are the *separation points*, i.e., points of type $(\text{p}, \text{pp})_{\text{s}}$ in Theorem 6.3.24. Therefore, given a realization of the Brownian net $\mathcal{N}$, we can sample a Brownian web $\mathcal{W}$ by forcing the Brownian web paths to go either left or right with probability $1/2$ each, independently at each separation point. The complication is that each path in the Brownian net will encounter infinitely many separation points on any finite time interval. But fortunately, the number of separation points that is relevant for determining the path's position at a given time t is almost surely locally finite.

Conversely, given a realization of coalescing random walks as illustrated in Figure 6.9 (b), we can construct a collection of branching-coalescing random walks by turning each lattice point in $\mathbb{Z}^2_{\text{even}}$ independently into a branching point with probability ε. The key observation by Newman et al. in [73] is that each lattice point in $\mathbb{Z}^2_{\text{even}}$ is a point where a random walk path and some dual random walk

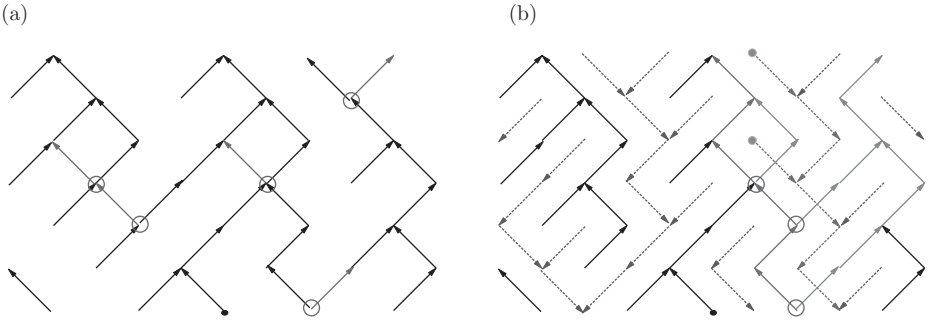

Figure 6.9. (a) Sampling coalescing random walks from branching-coalescing random walks. (b) Turning coalescing random walks into branching-coalescing random walks.

path are 1 unit of distance apart, and turning that lattice point into a branching point adds a new random walk path that crosses the dual random walk path (see Figure 6.9 (b), where the dotted lines are the dual paths, and the branching points are circled). In the diffusive scaling limit, most of the added branching points have no effect because they lead to small excursions from the coalescing random walk paths that vanish in the limit. The branching points that have an effect in the limit are points at which a forward random walk path and a dual random walk path – started macroscopically apart – come within 1 unit of distance from each other. In the scaling limit, these become precisely the (1,2) points of the Brownian web in Theorem 6.2.11, where a Brownian web path meets a dual Brownian web path. In the discrete setting, these points of close encounter between forward and backward random walk paths are turned independently into branching points with probability ε. In the continuum limit, this leads to Poisson marking of the points of collision between Brownian web paths and dual Brownian web paths, with the intensity measure given by the intersection local time measure between the forward and dual paths. The Brownian net can then be constructed by allowing the Brownian web paths to branch at these Poisson marked (1,2) points, which are precisely the separation points in the resulting Brownian net.

To formulate precisely the coupling between the Brownian web and net motivated by the above heuristic discussions, we first introduce the necessary background. By Exercise 6.3.28, the set of separation points S is a.s. a countable set, and it was pointed out in Remark 6.3.27 that any path $\pi \in \mathcal{N}$ entering a separation point $z = (x,t)$ must do so bounded between a pair of equivalent paths $(l,r) \in (\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ entering z, and similarly when leaving z, it must also be enclosed by one of the two pairs of outgoing equivalent paths $(l_1,r_1),(l_2,r_2) \in$

$(\mathcal{W}^{\mathrm{l}}(z), \mathcal{W}^{\mathrm{r}}(z))$, ordered from left to right. We can then define

$$\mathrm{sgn}_{\pi}(z) := \begin{cases} -1 & \text{if } l_1 \le \pi \le r_1 \text{ on } [t,\infty), \\ +1 & \text{if } l_2 \le \pi \le r_2 \text{ on } [t,\infty). \end{cases} \tag{6.23}$$

For points of type $(1,2)$ in the Brownian web $\mathcal{W}$, we can similarly define

$$\mathrm{sgn}_{\mathcal{W}}(z) := \begin{cases} -1 & \text{if z is of type } (1,2)_{\mathrm{l}} \text{ in } \mathcal{W}, \\ +1 & \text{if z is of type } (1,2)_{\mathrm{r}} \text{ in } \mathcal{W}. \end{cases} \tag{6.24}$$

By setting the sign of paths entering each separation point independently to be ± 1 with probability $1/2$, we will recover the Brownian web as a subset of the Brownian net.

To construct the Brownian net from the Brownian web, a key object is the local time measure on points of intersection between paths in $\mathcal{W}$ and paths in $\widehat{\mathcal{W}}$., i.e., points of type $(1,2)$ in $\mathcal{W}$. Its existence was proved in [73, Proposition 3.1], which we quote below (see also [92, Proposition 3.4]).

Proposition 6.4.1 (Intersection local time) *Let $(\mathcal{W}, \hat{\mathcal{W}})$ be the Brownian web and its dual. Then a.s. there exists a unique measure ℓ, concentrated on the set of points of type $(1,2)$ in $\mathcal{W}$, such that for each $\pi \in \mathcal{W}$ and $\hat{\pi} \in \hat{\mathcal{W}}$,*

$$\begin{aligned} &\ell(\{z = (x,t) \in \mathbb{R}^2 : \sigma_{\pi} < t < \hat{\sigma}_{\hat{\pi}},\ \pi(t) = x = \hat{\pi}(t)\}) \\ &= \lim_{\varepsilon \downarrow 0} \varepsilon^{-1} |\{t \in \mathbb{R} : \sigma_{\pi} < t < \hat{\sigma}_{\hat{\pi}},\ |\pi(t) - \hat{\pi}(t)| \le \varepsilon\}|. \end{aligned} \tag{6.25}$$

The measure ℓ is a.s. non-atomic and σ-finite. We let ℓ_{l} and ℓ_{r} denote the restrictions of ℓ to the sets of points of type $(1,2)_{\mathrm{l}}$ and $(1,2)_{\mathrm{r}}$, respectively.

By Poisson marking (1,2) points of $\mathcal{W}$ with intensity measure $\ell = \ell_{\mathrm{l}} + \ell_{\mathrm{r}}$, we obtain a countable set of (1,2) points. Allowing paths in $\mathcal{W}$ to branch at these points then leads to the Brownian net.

We can now formulate the coupling between the Brownian web and net (see [92, Theorems 4.4 and 4.6]).

Theorem 6.4.2 (Coupling between the Brownian web and net) *Let $\mathcal{W}$ be the standard Brownian web, and $\mathcal{N}$ the standard Brownian net. Let S denote the set of separation points of $\mathcal{N}$. Then there exists a coupling between $\mathcal{W}$ and $\mathcal{N}$ such that:*

(i) *Almost surely $\mathcal{W} \subset \mathcal{N}$, and each separation point $z \in S$ in $\mathcal{N}$ is of type $(1,2)$ in $\mathcal{W}$.*

(ii) *Conditional on $\mathcal{N}$, the random variables* $(\mathrm{sgn}_{\mathcal{W}}(z))_{z\in S}$ *are i.i.d. with* $\mathbb{P}[\mathrm{sgn}_{\mathcal{W}}(z) = \pm 1 \,|\, \mathcal{N}] = 1/2$*, and a.s.,*

$$\mathcal{W} = \{\pi \in \mathcal{N} : \mathrm{sgn}_{\pi}(z) = \mathrm{sgn}_{\mathcal{W}}(z)\ \forall z \in S \text{ s.t. } \pi \text{ enters } z\}. \tag{6.26}$$

(iii) *Conditional on $\mathcal{W}$, the sets* $S_{\mathrm{l}} := \{z \in S : \mathrm{sgn}_{\mathcal{W}}(z) = -1\}$ *and* $S_{\mathrm{r}} := \{z \in S : \mathrm{sgn}_{\mathcal{W}}(z) = +1\}$ *are independent Poisson point sets with intensities ℓ_{l} and ℓ_{r}, respectively, and a.s.,*

$$\mathcal{N} = \lim_{\Delta_n \uparrow S} \mathrm{hop}_{\Delta_n}(\mathcal{W}) \tag{6.27}$$

for any sequence of finite sets Δ_n increasing to $S = S_{\mathrm{l}} \cup S_{\mathrm{r}}$, where $\mathrm{hop}_{\Delta_n}(\mathcal{W})$ *is the set of paths obtained from $\mathcal{W}$ by allowing paths entering any (1,2) point $z \in \Delta_n$ to continue along either of the two outgoing paths at z.*

Remark 6.4.3 Theorem 6.4.2 can be generalized to the case where $\mathcal{W}^{\mathrm{l}}$, $\mathcal{W}$, and $\mathcal{W}^{\mathrm{r}}$ may be tilted with different drifts, as long as the drifts remain ordered (see [92, Theorem 6.15]). We only need to modify the probability $\mathbb{P}(\mathrm{sgn}_{\mathcal{W}}(z) = \pm 1|\mathcal{N})$ in (ii), while the intensity measures ℓ_{l} and ℓ_{r} in (iii) need to be multiplied by constants depending on the drifts of $\mathcal{W}^{\mathrm{l}}$, $\mathcal{W}$ and $\mathcal{W}^{\mathrm{r}}$.

Remark 6.4.4 Theorem 6.4.2 (ii) shows how to sample a Brownian web $\mathcal{W}$ from a Brownian net $\mathcal{N}$ by forcing the paths to continue either left or right independently at each separation point. What if we sample independently another Brownian web $\mathcal{W}'$, conditioned on $\mathcal{N}$? How can we characterize the joint distribution of $\mathcal{W}$ and $\mathcal{W}'$? It turns out that $(\mathcal{W}, \mathcal{W}')$ forms a pair of so-called *sticky Brownian webs*, where Brownian motions in $\mathcal{W}$ and $\mathcal{W}'$ undergo sticky interaction. Such an object was first introduced in [56]. For further details, see [92, Section 3.3 and Lemma 6.16].

Proof Sketch. Theorem 6.4.2 is proved in [92, Section 6] via discrete approximation, using the fact that a similar coupling to that in Theorem 6.4.2 (ii) and (iii) holds in the discrete system. The proof is too complex and lengthy to be included here. Instead, we outline below some key ingredients.

The fact that a coupling exists between the Brownian web and net follows by taking the scaling limit of the coupled coalescing and branching-coalescing random walks, where i.i.d. signs are assigned to each branching point to determine the path of the coalescing walks.

To show that such a coupling satisfies Theorem 6.4.2 (ii), the key is to show that the branching-coalescing random walks, together with the branching points, converge to the Brownian net together with its separation points. Furthermore, under such a convergence, we can match the branching points with the separation points and assign them the same i.i.d. signs, such that coalescing random walk

paths converge to paths in the Brownian net as long as the discrete and continuum paths follow the same signs at the respective branching and separation points.

The main difficulty with the above approach is that the set of separation points of $\mathcal{N}$ is dense in $\mathbb{R}^2$, so it is unclear in what sense the branching points should converge to the separation points. The solution rests on the observation that, for Brownian net paths starting at some time S, their positions at a later time U depend a.s. only on the signs of these paths at a locally finite set of separation points in the time interval (S,U) (called (S,U)-*relevant separation points*). As we refine our knowledge of the paths at more times in the interval (S,U), more separation points become relevant. Furthermore, the relevant separation points form a locally finite directed graph, called the *finite graph representation* of the Brownian net, which determines a coarse-grained structure of the Brownian net. Therefore a natural notion of convergence is to show that for any $S < U$, the (S,U)-relevant separation points and the associated finite graph representation arise as limits of similar structures in the branching-coalescing random walks. Once such a convergence is verified, Theorem 6.4.2 can then be easily verified.

To show that the coupling between the Brownian web and net satisfies the first statement in Theorem 6.4.2 (iii), the key is to show that when restricted to the set of intersection points between a pair of forward and dual Brownian web paths $(\pi, \hat{\pi})$, the set of separation points is distributed as a Poisson point process with intensity measure given by the intersection local time measure between π and $\hat{\pi}$. These separation points are in fact relevant separation points with respect to the starting times of π and $\hat{\pi}$, and hence must arise as the limit of relevant branching points in the discrete system. On the other hand, the corresponding relevant branching points form a Bernoulli point process with the intensity measure given by the counting measure on the set of space-time points where a pair of forward and dual coalescing random walk paths approximating $(\pi, \hat{\pi})$ come within distance 1. With space-time scaled by $(\varepsilon, \varepsilon^2)$, this counting measure can be shown to converge to the intersection local time measure between π and $\hat{\pi}$, and hence the limit of the branching points (i.e., the set of separation points restricted to the intersections between π and $\hat{\pi}$) must be a Poisson point process with intensity measure given by the intersection local time measure.

To prove the second part of Theorem 6.4.2 (iii), i.e., (6.27), note that under the coupling between the Brownian web and net, the Brownian web is embedded in the Brownian net and the Poisson marked (1,2) points of the web are exactly the separation points of the net. Using the coarse-grained structures of the Brownian net given by the finite graph representations, it can then be shown that turning the separation points into branching points for the Brownian web gives the Brownian net. □

The proof sketch above shows that the key ingredient in the proof of Theorem 6.4.2 is the notion of relevant separation points and the associated finite graph representation of the Brownian net. Since they shed new light on the structure of the Brownian net, we will discuss these notions in detail in the rest of this section.

6.4.1 Relevant Separation Points of the Brownian Net

As noted in Exercise 6.3.28, the set of separation points (i.e., points of type $(\mathrm{p}, \mathrm{pp})_{\mathrm{s}}$ in Theorem 6.3.24) is almost surely countable and dense in $\mathbb{R}^2$. In fact almost surely each path in the Brownian net encounters infinitely many separation points in any open time interval. However, for given deterministic times $S < U$, there is only a locally finite set of separation points which are relevant for deciding where paths in the Brownian net started at time S end up at time U. More precisely, as illustrated in Figure 6.10, we define (cf. [91, Section 2.3]):

Definition 6.4.5 (Relevant separation points) *A separation point $z = (x,t)$ of the Brownian net $\mathcal{N}$ is called* (S,U)-relevant *for some $-\infty \le S < t < U \le \infty$, if there exists $\pi \in \mathcal{N}$ such that $\sigma_\pi = S$ and $\pi(t) = x$, and there exist $l \in \mathcal{W}^{\mathrm{l}}(z)$ and $r \in \mathcal{W}^{\mathrm{r}}(z)$ such that $l < r$ on (t,U).*

Note that because z is a separation point, l and r are continuations of incoming paths at z. By Proposition 6.3.15, the paths obtained by hopping from π to either l or r at z are both paths in the Brownian net, which are distinct on the time interval (t,U). We do not require $l(U) < r(U)$ because an (S,U)-relevant separation point defined as above also turns out to be a relevant separation point w.r.t. the dual Brownian net $\widehat{\mathcal{N}}$.

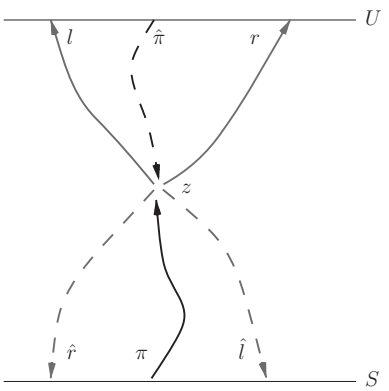

Figure 6.10. An (S,U)-relevant separation point.

Exercise 6.4.6 *Use the wedge characterization of the Brownian net and the steering argument illustrated in Figure 6.6 to show that almost surely, for each $-\infty \le S < U \le \infty$, a separation point $z = (x,t)$ with $S < t < U$ is (S,U)-relevant in $\mathcal{N}$ if and only if $-z$ is $(-U,-S)$-relevant in the dual Brownian net $-\widehat{\mathcal{N}}$ rotated 180^o around the origin in $\mathbb{R}^2$.*

A crucial property of the (S,U)-relevant separation points is that they are almost surely locally finite, which follows from the following density calculation (cf. [91, Proposition 2.9]).

Proposition 6.4.7 (Density of relevant separation points) *Let $\mathcal{N}$ be a standard Brownian net. Then for each deterministic $-\infty \le S < U \le \infty$, if $R_{S,U}$ denotes the set of (S,U)-relevant separation points, then*

$$\mathbb{E}[|R_{S,U} \cap A|] = 2 \int_A \Psi(t-S)\Psi(U-t)\,\mathrm{d}x\,\mathrm{d}t \qquad \text{for all Borel-measurable } A \subset \mathbb{R} \times (S,U), \tag{6.28}$$

where

$$\Psi(t) := \frac{e^{-t}}{\sqrt{\pi t}} + 2\Phi(\sqrt{2t}), \quad 0 < t \le \infty, \quad \text{and} \quad \Phi(x) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-y^2/2}\mathrm{d}y. \tag{6.29}$$

In particular, if $-\infty < S,\ U < \infty$, then $R_{S,U}$ is a.s. a locally finite subset of $\mathbb{R} \times [S,U]$.

Proof Sketch. For $S < s < u < U$, let

$$E_s := \{\pi(s) : \pi \in \mathcal{N}, \sigma_\pi = S\} \qquad \text{and} \qquad F_u := \{\hat{\pi}(s) : \hat{\pi} \in \widehat{\mathcal{N}}, \hat{\sigma}_{\hat{\pi}} = U\}.$$

Observe that a.s., each (S,U)-relevant separation point z in the strip $\mathbb{R} \times [s,u]$ can be traced back along some Brownian net path $\pi \in \mathcal{N}$ to a position $x \in E_s$. Furthermore, if $(l,r) \in (\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ is the pair of left-right Brownian web paths starting at (x,s), then we must have $F_u \cap (l(u), r(u)) \neq \emptyset$, and z is bounded between l and r on the time interval $[s,u]$. Therefore each (S,U)-relevant separation point in $\mathbb{R} \times [s,u]$ can be approximated by some (x,s) with x in

$$Q_{s,u} := \{x \in E_s :\ F_u \cap (l(u), r(u)) \neq \emptyset, \text{ where } (l,r) \in (\mathcal{W}^{\mathrm{l}}(x,s), \mathcal{W}^{\mathrm{r}}(x,s))\}.$$

The density of the set $Q_{s,u}$ can be easily computed. As we partition (S,U) into disjoint intervals $[t_i, t_{i+1})$ of size $1/n$ with $n \uparrow \infty$, each (S,U)-relevant separation point z can be approximated by some (x_n, t_{i_n}) with $x_n \in Q_{t_{i_n}, t_{i_n+1}}$ for some i_n, such that $(x_n, t_{i_n}) \to z$ as $n \to \infty$. The upper bound on the density of $R_{S,U}$ then follows

by Fatou's Lemma. For the lower bound, it suffices to consider points in $Q_{t_i,t_{i+1}}$ which are separated from each other by some $\delta > 0$ which is sent to zero after taking the limit $n \to \infty$. □

Exercise 6.4.8 *Prove Proposition 6.4.7 by rigorously implementing the strategy outlined above.*

6.4.2 Finite Graph Representation of the Brownian Net

Having defined (S,U)-relevant separation points and established their local finiteness, we are now ready to introduce the *finite graph representation*, which is a directed graph with the interior vertices given by the (S,U)-relevant separation points, and the directed edges determined by how Brownian net paths go from one relevant separation point to the next. Such a directed graph gives a coarse-grained representation of the Brownian net, and plays a key role in studying the coupling between the Brownian web and net in Theorem 6.4.2.

Let $-\infty < S < U < \infty$ be deterministic times, let $R_{S,U}$ be the set of (S,U)-relevant separation points of $\mathcal{N}$ and set

$$\begin{aligned} R_S &:= \mathbb{R} \times \{S\}, \\ R_U &:= \{(x,U) : x \in \mathbb{R},\ \exists \pi \in \mathcal{N} \text{ with } \sigma_\pi = S \text{ s.t. } \pi(U) = x\}. \end{aligned} \tag{6.30}$$

We will make the set $R := R_S \cup R_{S,U} \cup R_U$ into a directed graph, with the directed edges representing the Brownian net paths.

First we identify the directed edges leading out from each $z = (x,t) \in R_{S,U}$. By Remark 6.3.27 and Proposition 6.3.14, any $\pi \in \mathcal{N}$ with $\sigma_\pi = S$ that enters an (S,U)-relevant separation point z must leave z enclosed by one of the two outgoing pairs of equivalent left-right paths $(l_1,r_1),(l_2,r_2) \in (\mathcal{W}^{\mathrm{l}}(z),\mathcal{W}^{\mathrm{r}}(z))$, with $l_1 \sim^z_{\mathrm{out}} r_1$ and $l_2 \sim^z_{\mathrm{out}} r_2$. Take the pair (l_1,r_1) for instance. If $l_1(U) \neq r_1(U)$, then there must be a last separation point z' along the pair (l_1,r_1), which is also an (S,U)-relevant separation point because hopping from π to either l_1 or r_1 at time t still gives a Brownian net path. If $l_1(U) = r_1(U)$, then we just set $z' = (l_1(U),U) = (r_1(U),U) \in R_U$. In either case, we note that all paths that leave z bounded between l_1 and r_1 must continue to do so until they reach z'. Therefore we draw a directed edge from z to z', denoted by $z \to_{l_1,r_1} z'$, representing all Brownian net paths that go from z to z' while bounded between l_1 and r_1. Similarly, paths that leave z while bounded between l_2 and r_2 must all lead to some $z'' \in R_{S,U} \cup R_U$, which is represented by a directed edge $z \to_{l_2,r_2} z''$.

Next we identify the directed edges leading out from $z = (x,S) \in R_S$. By Theorem 6.3.24, at a deterministic time S, almost surely each $z \in R_S$ is of type (o, p), (p, p), or (o, pp). If z is of type (o, p) or (p, p), then a single equivalent

pair $l \sim^{z}_{\text{out}} r$ starts from z, and all Brownian net paths leaving z must be bounded between l and r, and leads to some $z' \in R_{S,U} \cup R_U$, which we represent by a directed edge $z \to_{l,r} z'$. Similarly when z is of type (o, pp), two directed edges start from z.

Given the above directed graph with vertex set $R_S \cup R_{S,U} \cup R_U$, it is not difficult to see that each Brownian net path from time S to U corresponds to a directed path from R_S to R_U, and conversely, each directed path from R_S to R_U can be associated with a family of Brownian net paths from time S to U. We summarize the basic properties below (see [92, Proposition 6.5]).

Proposition 6.4.9 (Finite graph representation) *Let $\mathcal{N}$ be a Brownian net with associated left-right Brownian web $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, and let $-\infty < S < U < \infty$ be deterministic times. Let $R := R_S \cup R_{S,U} \cup R_U$ and directed edges $\to_{l,r}$ be defined as above. Then, a.s. (see Figure 6.11):*

(a) *For each $z \in R_S$ that is not of type* (o, pp)*, there exist unique $l \in \mathcal{W}^{\mathrm{l}}(z)$, $r \in \mathcal{W}^{\mathrm{r}}(z)$ and $z' \in R$ such that $z \to_{l,r} z'$.*

(b) *For each $z = (x,t)$ such that either $z \in R_{S,U}$ or $z \in R_S$ is of type* (o, pp)*, there exist unique $l, l' \in \mathcal{W}^{\mathrm{l}}(z)$, $r, r' \in \mathcal{W}^{\mathrm{r}}(z)$ and $z', z'' \in R$ such that $l \le r' < l' \le r$ on $(t, t+\varepsilon)$ for some $\varepsilon > 0$, $z \to_{l,r'} z'$ and $z \to_{l',r} z''$. For $z \in R_{S,U}$ one has $z' \neq z''$. For $z \in R_S$ of type* (o, pp)*, one has $z' \neq z''$ if and only if there exists $\hat{\pi} \in \widehat{\mathcal{N}}$ with $\hat{\sigma}_{\hat{\pi}} = U$ such that $\hat{\pi}$ enters z.*

(c) *For each $\pi \in \mathcal{N}$ with $\sigma_\pi = S$, there exist $z_i = (x_i, t_i) \in R$ $(i = 0, \ldots, n)$ and $l_i \in \mathcal{W}^{\mathrm{l}}(z_i)$, $r_i \in \mathcal{W}^{\mathrm{l}}(z_i)$ $(i = 0, \ldots, n-1)$ such that $z_0 \in R_S$, $z_n \in R_U$, $z_i \to_{l_i, r_i} z_{i+1}$ and $l_i \le \pi \le r_i$ on $[t_i, t_{i+1}]$ $(i = 0, \ldots, n-1)$.*

(d) *If $z_i = (x_i, t_i) \in R$ $(i = 0, \ldots, n)$ and $l_i \in \mathcal{W}^{\mathrm{l}}(z_i)$, $r_i \in \mathcal{W}^{\mathrm{l}}(z_i)$ $(i = 0, \ldots, n-1)$ satisfy $z_0 \in R_S$, $z_n \in R_U$, and $z_i \to_{l_i, r_i} z_{i+1}$ $(i = 0, \ldots, n-1)$, then there exists a $\pi \in \mathcal{N}$ with $\sigma_\pi = S$ such that $l_i \le \pi \le r_i$ on $[t_i, t_{i+1}]$.*

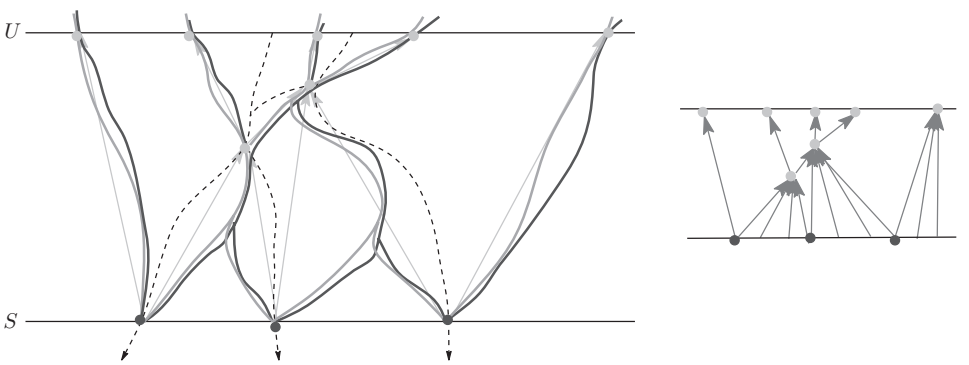

Figure 6.11. Finite graph representation.

Exercise 6.4.10 *Use the finite graph representation to show that: conditioned on the Brownian net $\mathcal{N}$ with the set of separation points S, almost surely for any $(\alpha_z)_{z\in S} \in \{\pm 1\}^S$,*

$$W = \{\pi \in \mathcal{N} : \mathrm{sgn}_\pi(z) = \alpha_z \, \forall z \in S \text{ s.t. } \pi \text{ enters } z\}$$

(cf. (6.26)*) is a closed subset of $\mathcal{N}$, containing at least one path starting from each $(x,t) \in \mathbb{R}^2$.*

6.5 Scaling Limits of Random Walks in i.i.d. Space-time Environment

We will now show how the Brownian web and net can be used to construct the continuum limits of one-dimensional random walks in i.i.d. random space-time environments.

A random walk X in an i.i.d. random space-time environment is defined as follows. Let $\omega := (\omega_z)_{z\in\mathbb{Z}^2_{\text{even}}}$ be i.i.d. $[0,1]$-valued random variables with common distribution μ. We view ω as a random space-time environment, such that conditioned on ω, if the random walk is at position x at time t, then in the next time step the walk jumps to $x+1$ with probability $\omega_{(x,t)}$ and jumps to $x-1$ with probability $1-\omega_{(x,t)}$ (see Figure 6.12). Let $P^\omega_{(x,s)}$ denote probability for the walk starting at x at time s in the environment ω, and let $\mathbb{P}$ denote probability for ω.

The question is: if we scale space and time by ε and ε^2 respectively, is it possible to choose a law μ_ε for an environment $\omega^\varepsilon := (\omega^\varepsilon_z)_{z\in\mathbb{Z}^2_{\text{even}}}$ such that the walk in the environment ω^ε converges as $\varepsilon \downarrow 0$ to a limiting random motion in a continuum space-time random environment? Different ways of looking at the random walk in random environment suggest that the answer is yes.

One alternative way of characterizing the law of random walk in random environment is to consider the law of the family of random transition probability kernels

$$K^\omega_{s,t}(x,y) := P^\omega_{(x,s)}(X(t) = y), \qquad (x,s),(y,t) \in \mathbb{Z}^2_{\text{even}},\ s \le t. \tag{6.31}$$

Still another way is to specify the *n-point motions*, i.e., the law of n random walks $\vec{X} := (X_1(t),\ldots,X_n(t))_{t\ge 0}$ sampled independently from the same environment ω, and then averaged with respect to the law of ω. Note that $\vec{X}$ is a Markov chain with transition probability kernel

$$K^{(n)}_{0,t}(\vec{x},\vec{y}) = \int \prod_{i=1}^{n} K^\omega_{0,t}(x_i,y_i)\mathbb{P}(\mathrm{d}\omega), \quad (x_i,0),(y_i,t) \in \mathbb{Z}^2_{\text{even}} \text{ for } 1 \le i \le n. \tag{6.32}$$

Furthermore, the n-point motions are *consistent* in the sense that the law of any k-element subset of $(X_1(\,\cdot\,),\ldots,X_n(\,\cdot\,))$ is governed by that of the k-point motion.

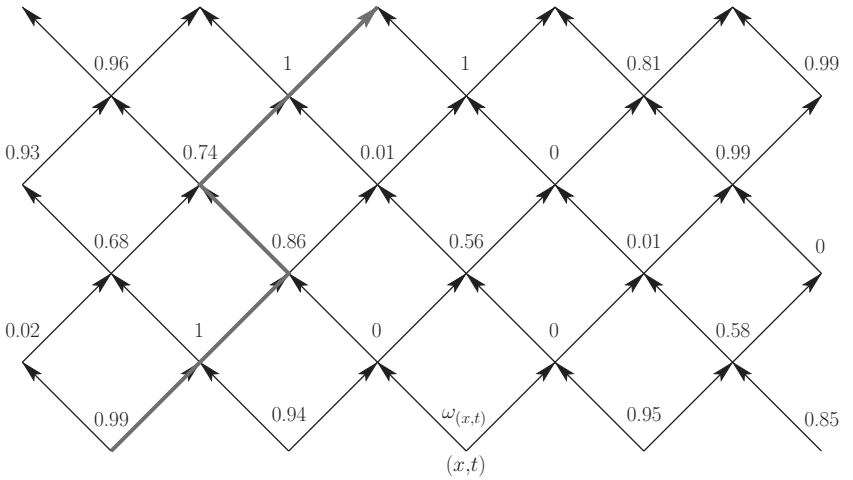

Figure 6.12. Random walk on $\mathbb{Z}^2_{\text{even}}$ in a random environment ω.

Note that the moments of $\omega_{(0,0)}$ are determined by $(K^{(n)})_{n\in\mathbb{N}}$, since

$$\mathbb{E}[\omega^n_{(0,0)}] = K^{(n)}_{0,1}(\vec{0},\vec{1}).$$

Therefore the law of $\omega_{(0,0)}$, as well as that of ω and K^ω, are uniquely determined by the law of the n-point motions $(K^{(n)})_{n\in\mathbb{N}}$.

Evidence for the existence of a continuum limit for random walks in i.i.d. space-time random environments came from the convergence of the n-point motions, which was first established by Le Jan and Lemaire [62] for i.i.d. Beta-distributed random environments, and subsequently extended to general environments by Howitt and Warren [55]. By the theory of *stochastic flow of kernels* developed by Le Jan and Raimond [63], this implies that the family of random probability kernels K^ω in (6.32) also converges in a suitable sense to a continuum limit. Motivated by these results, it was then shown in [92] that under the same convergence criterion as in [55], not only the n-point motions converge, but also the random environments themselves converge to a continuum space-time limit which can be constructed explicitly from the Brownian web and net.

In Section 6.5.1 below, we will first briefly recall the theory of *stochastic flows of kernels* from [63], and then review Howitt and Warren's convergence result [55] for the n-point motions of random walks in i.i.d. space-time random environments. In Section 6.5.2, we will then show how to construct the limiting continuum space-time environment from the Brownian web and net, which arises naturally if one looks at the discrete environments in the right way. Lastly in Section 6.5.3, we will discuss some properties of the continuum random motion in random environment.

6.5.1 Stochastic Flows of Kernels and the Howitt-Warren Flows

In (6.31) and (6.32), we saw that the family of random transition probability kernels $K^{\omega}_{s,t}(x,\cdot)$ and the family of consistent n-point motions provide alternative characterizations of a random walk in i.i.d. space-time random environment. It turns out that in general, without knowing the existence of any underlying random environment, there is a correspondence between a consistent family of n-point motions on a Polish space, and a family of random probability kernels called a *stochastic flow of kernels*. This was the main result of Le Jan and Raimond in [63, Theorem 2.1], which motivated the study of concrete examples of consistent n-point motions in [64, 62, 55]. We now recall the notion of a *stochastic flow of kernels*.

For any Polish space E, let $\mathcal{B}(E)$ denote the Borel σ-field on E and let $\mathcal{M}_1(E)$ denote the space of probability measures on E, equipped with the weak topology and Borel σ-algebra. A *random probability kernel*, defined on some probability space $(\Omega,\mathcal{F},\mathbb{P})$, is a measurable function $K : \Omega \times E \to \mathcal{M}_1(E)$. Two random probability kernels K, K' are said to equal in finite-dimensional distributions if for each $x_1,\dots,x_n \in E$, the n-tuple of random probability measures $(K(x_1,\cdot),\dots,K(x_n,\cdot))$ is equally distributed with $(K'(x_1,\cdot),\dots,K'(x_n,\cdot))$. Two or more random probability kernels are called independent if their finite-dimensional distributions are independent.

Definition 6.5.1 (Stochastic flow of kernels) *A* stochastic flow of kernels *on a Polish space E is a collection $(K_{s,t})_{s\le t}$ of random probability kernels on E such that*

(i) *For all $s \le t \le u$ and $x \in E$, a.s. $K_{s,s}(x,A) = \delta_x(A)$ and $\int_E K_{s,t}(x,\mathrm{d}y)K_{t,u}(y,A) = K_{s,u}(x,A)$ for all $A \in \mathcal{B}(E)$.*

(ii) *For each $t_0 < \cdots < t_n$, the random probability kernels $(K_{t_{i-1},t_i})_{i=1,\dots,n}$ are independent.*

(iii) *$K_{s,t}$ and $K_{s+u,t+u}$ are equal in finite-dimensional distributions for each real $s \le t$ and u.*

We have omitted two weak continuity conditions on $(K_{s,t}(x,\cdot)_{s\le t,x\in E}$ from [63, Definition 2.3], which are automatically satisfied by a version of K if the n-point motions defined via (6.32) are a family of Feller processes.

Remark 6.5.2 Although we motivated the notion of stochastic flow of kernels from random walks in i.i.d. space-time random environments, the kernels K in Definition 6.5.1 cannot be associated with an underlying random environment

unless there exists a version of K such that Definition 6.5.1 (i) is strengthened to:

$$\text{(i}') \qquad A.s., \ \forall s \le t \le u \text{ and } x \in E, \ \ K_{s,s}(x,\cdot) = \delta_x(\,\cdot\,) \text{ and } \int_E K_{s,t}(x,\mathrm{d}y) K_{t,u}(y,\cdot) = K_{s,u}(x,\cdot).$$

In general, it is not known whether such a version always exists.

Le Jan and Raimond showed in [63, Theorem 2.1] that every consistent family of Feller n-point motions corresponds to a stochastic flow of kernels. Using Dirichlet form construction of Markov processes, they then constructed as an example in [64] a consistent family of n-point motions on the circle which are a special type of sticky Brownian motions. Subsequently, Le Jan and Lemaire showed in [62] that the n-point motions of random walks in i.i.d. Beta-distributed space-time random environments converge to the n-point motions constructed in [64]. Howitt and Warren [55] then found the general condition for the convergence of n-point motions of random walks in i.i.d. space-time random environments, and they characterized the limiting n-point motions, which are also sticky Brownian motions, in terms of well-posed martingale problems.

We next recall Howitt and Warren's result [55], or rather, a different formulation of their result as presented in [92, Appendix A] with discrete instead of continuous time random walks, and with a reformulation of the martingale problem characterizing the limiting sticky Brownian motions.

Theorem 6.5.3 (Convergence of n-point motions to sticky Brownian motions) *For $\varepsilon > 0$, let μ_ε be the common law of an i.i.d. space-time random environment $(\omega_z^\varepsilon)_{z \in \mathbb{Z}^2_{\text{even}}}$, such that*

$$\begin{aligned} &\text{(i)} \quad \varepsilon^{-1} \int (2q-1)\mu_\varepsilon(\mathrm{d}q) \underset{\varepsilon \to 0}{\longrightarrow} \beta, \\ &\text{(ii)} \quad \varepsilon^{-1} q(1-q)\mu_\varepsilon(\mathrm{d}q) \underset{\varepsilon \to 0}{\Longrightarrow} \nu(\mathrm{d}q), \end{aligned} \tag{6.33}$$

for some $\beta \in \mathbb{R}$ and finite measure ν on $[0,1]$, where $\Rightarrow$ denotes weak convergence.

Let $\vec{X}^\varepsilon := (X_1^\varepsilon, \ldots, X_n^\varepsilon)$ be n independent random walks in ω^ε with $\varepsilon \vec{X}^\varepsilon(0) \to \vec{X}(0) = (X_1(0), \ldots, X_n(0)) \in \mathbb{R}^n$ as $\varepsilon \downarrow 0$. Then $(\varepsilon \vec{X}^\varepsilon(\varepsilon^2 t))_{t \ge 0}$ converges weakly to a family of sticky Brownian motions $(\vec{X}(t))_{t \ge 0} = (X_1(t), \ldots, X_n(t))_{t \ge 0}$, whose law is the unique solution of the following Howitt-Warren martingale problem with drift β and characteristic measure ν*:*

(i) *$\vec{X}$ is a continuous, square-integrable semi-martingale with initial condition $\vec{X}(0)$;*

(ii) *The covariance process between X_i and X_j is given by*

$$\langle X_i, X_j \rangle(t) = \int_0^t 1_{\{X_i(s)=X_j(s)\}} \mathrm{d}s, \qquad t \geq 0,\ i,j = 1,\dots,n; \tag{6.34}$$

(iii) *For each non-empty $\Delta \subset \{1,\dots,n\}$, let*

$$f_\Delta(\vec{x}) := \max_{i\in\Delta} x_i \qquad \text{and} \qquad g_\Delta(\vec{x}) := |\{i \in \Delta : x_i = f_\Delta(\vec{x})\}| \qquad (\vec{x} \in \mathbb{R}^n). \tag{6.35}$$

Then

$$f_\Delta(\vec{X}(t)) - \int_0^t \beta_+(g_\Delta(\vec{X}(s)))\mathrm{d}s \tag{6.36}$$

is a martingale with respect to the filtration generated by $\vec{X}$, where

$$\beta_+(1) := \beta \quad \text{and} \quad \beta_+(m) := \beta + 2\int \nu(\mathrm{d}q) \sum_{k=0}^{m-2} (1-q)^k \quad \text{for } m \geq 2. \tag{6.37}$$

Remark 6.5.4 The Howitt-Warren sticky Brownian motions evolve independently when they are apart, and experience sticky interaction when they meet. In particular, when $n = 2$, $X_1(t) - X_2(t)$ is a Brownian motion with stickiness at the origin, which is just a time changed Brownian motion such that its local time at the origin has been turned into real time, modulo a constant multiple that determines the stickiness. More generally, for Howitt-Warren sticky Brownian motions started at $X_1(0) = \cdots = X_n(0)$, the set of times with $X_1(t) = X_2(t) = \cdots = X_n(t)$ is a nowhere dense set with positive Lebesgue measure. The measure ν determines a two-parameter family of constants $\theta_{k,l} = \int q^k(1-q)^l \frac{\nu(\mathrm{d}q)}{q(1-q)}$, $k,l \geq 1$, which can be regarded as the rate (in a certain excursion theoretic sense) at which $(X_1,\cdots,X_n)$ split into two groups, $(X_1,\cdots,X_k)$ and $(X_{k+1},\cdots,X_{k+l})$, with $k+l=n$.

It is easily seen that the n-point motions defined by the Howitt-Warren martingale problem in Theorem 6.5.3 form a consistent family, and it is Feller by [55, Proposition 8.1]. Therefore by the aforementioned result of Le Jan and Raimond [63, Theorem 2.1], there exists a stochastic flow of kernels $(K_{s,t})_{s\leq t}$ on $\mathbb{R}$, unique in finite-dimensional distributions, such that the n-point motions of $(K_{s,t})_{s\leq t}$ in the sense of (6.32) are given by the unique solutions of the Howitt-Warren martingale problem. Therefore we define as follows:

Definition 6.5.5 (Howitt-Warren flow) *We call the stochastic flow of kernels, whose n-point motions solve the Howitt-Warren martingale problem in Theorem 6.5.3 for some $\beta \in \mathbb{R}$ and finite measure ν on $[0,1]$, the* Howitt-Warren flow *with* drift β *and* characteristic measure ν.

Remark 6.5.6 We single out three special classes of Howitt-Warren flows: (1) $\nu = 0$, for which the n-point motions are coalescing Brownian motions, and the flow is known as the *Arratia flow*; (2) $\nu(\mathrm{d}x) = a\mathrm{d}x$ for some $a > 0$, which we will call the *Le Jan-Raimond flow* since it was first constructed in [64] via Dirichlet forms and subsequently shown in [62] to arise as limits of random walks in i.i.d. Beta-distributed environments; (3) $\nu = a\delta_0 + b\delta_1$ for some $a, b \geq 0$ with $a + b > 0$, called the *erosion flow*, which was studied in [56].

As noted in Remark 6.5.2, it is not known a priori whether there exists a version of the Howitt-Warren flow K such that almost surely, $(K_{s,t}(x,\cdot))_{s<t,x\in\mathbb{R}}$ are truly transition probability kernels of a random motion in a random environment. It is then natural to ask whether such an underlying random environment indeed exists for the Howitt-Warren flows, and if yes, whether the environment can be explicitly characterized as in the discrete case. This is where the Brownian web and Brownian net enter the picture.

6.5.2 The Space-time Random Environment for the Howitt-Warren Flows

How can we construct a continuum space-time random environment such that the Howitt-Warren flow with drift β and characteristic measure ν is indeed the family of transition probability kernels of a random motion in this random environment? The answer again lies in discrete approximation.

Special case: ν satisfies $\int q^{-1}(1-q)^{-1}\nu(\mathrm{d}q) < \infty$. In this case, the continuum random environment can be constructed from the Brownian net. For $\varepsilon > 0$, define a probability measure μ_ε on $[0,1]$ by

$$\mu_\varepsilon := b\varepsilon\bar{\nu} + \tfrac{1}{2}(1-(b+c)\varepsilon)\delta_0 + \tfrac{1}{2}(1-(b-c)\varepsilon)\delta_1$$
$$\text{where} \quad b := \int \frac{\nu(\mathrm{d}q)}{q(1-q)}, \ c := \beta - \int (2q-1)\frac{\nu(\mathrm{d}q)}{q(1-q)}, \quad \bar{\nu}(\mathrm{d}q) := \frac{\nu(\mathrm{d}q)}{bq(1-q)}. \tag{6.38}$$

When ε is sufficiently small, such that $1-(b+|c|)\varepsilon \geq 0$, μ_ε is a probability measure on $[0,1]$ and is easily seen to satisfy (6.33) as $\varepsilon \downarrow 0$. Therefore by Theorem 6.5.3, μ_ε determines the law of an i.i.d. random environment $\omega^\varepsilon := (\omega_z^\varepsilon)_{z\in\mathbb{Z}^2_{\text{even}}}$, whose associated n-point motions converge to that of the Howitt-Warren flow with drift β and characteristic measure ν.

Note that for small ε, most of the ω_z^ε are either zero or one. We can thus encode ω^ε in two steps. First we identify a collection of branching-coalescing random walks determined by ω^ε, where at each $z = (x,t) \in \mathbb{Z}^2_{\text{even}}$, the walk moves to $(x+1, t+1)$ if $\omega_z^\varepsilon = 1$, moves to $(x-1, t+1)$ if $\omega_z^\varepsilon = 0$, and is a branching point

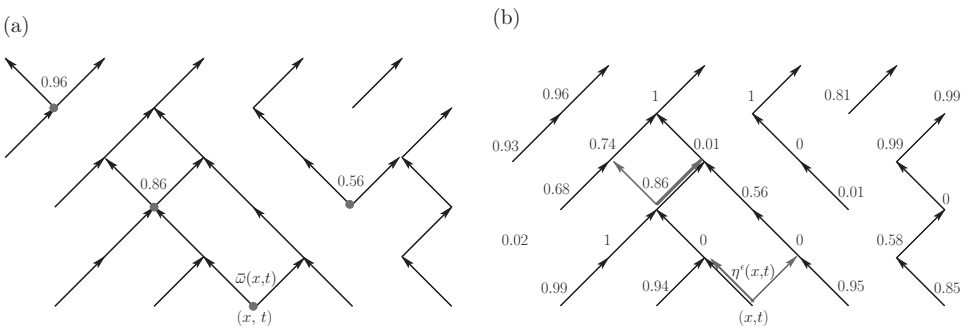

Figure 6.13. Representation of the random environment $(\omega_z^\varepsilon)_{z\in\mathbb{Z}^2_{\text{even}}}$ in terms of: (a) a marked discrete net $(N^\varepsilon, \bar{\omega}^\varepsilon)$; (b) a marked discrete web $(W^\varepsilon, \eta^\varepsilon)$.

if $\omega_z^\varepsilon \in (0,1)$ since the walk has strictly positive probability of moving to either $(x+1,t+1)$ or $(x-1,t+1)$. Given the set of branching-coalescing random walk paths, which we denote by N^ε and call a *discrete net*, we can then specify the value of ω_z^ε at each branching point by sampling i.i.d. random variables $\bar{\omega}_z^\varepsilon$ with common law $\bar{\nu}$ (see Figure 6.13(a)). The pair $(N^\varepsilon, \bar{\omega}^\varepsilon)$ then gives an alternative representation of the random environment ω, wherein a walk must navigate along N^ε, and when it encounters a branching point z, it jumps either left or right with probability $1-\bar{\omega}_z^\varepsilon$, respectively $\bar{\omega}_z^\varepsilon$.

The above setup is essentially the same as the coupling between branching-coalescing random walks and coalescing random walks discussed in Section 6.4, which corresponds to taking $(\bar{\omega}_z^\varepsilon)_{z\in\mathbb{Z}^2_{\text{even}}}$ to be i.i.d. $\{0,1\}$-valued random variables in the current setting. As we rescale space and time by ε and ε^2 respectively, we note that N^ε converges to a variant of the Brownian net $\mathcal{N}_{\beta_-,\beta_+}$, constructed from a pair of left-right Brownian webs as in Section 6.3 with respective drifts

$$\beta_- = \beta - 2\int \nu(\mathrm{d}q)(1-q)^{-1}, \qquad \beta_+ = \beta + 2\int \nu(\mathrm{d}q)\, q^{-1}. \tag{6.39}$$

The branching points of N^ε converge to the set of separation points of $\mathcal{N}_{\beta_-,\beta_+}$, denoted by S. Since the law of $\bar{\omega}_z^\varepsilon$ at the branching points is $\bar{\nu}$, independent of ε, we should then assign i.i.d. random variables ω_z with the same law $\bar{\nu}$ to each separation point $z \in S$. The pair $(\mathcal{N}_{\beta_-,\beta_+}, (\omega_z)_{z\in S})$ then gives the desired continuum space-time random environment, wherein a random motion π must navigate along $\mathcal{N}_{\beta_-,\beta_+}$, and independently at each separation point z it encounters, it chooses $\mathrm{sgn}_\pi(z)$ (see (6.23)) to be $+1$ with probability ω_z and -1 with probability $1-\omega_z$.

We next give a precise formulation of how the subclass of Howitt-Warren flows with $\int q^{-1}(1-q)^{-1}\nu(\mathrm{d}q) < \infty$ can be obtained from the random environment $(\mathcal{N}_{\beta_-,\beta_+}, (\omega_z)_{z\in S})$ (cf. [92, Theorem 4.7]). For its proof, see [92]. To construct a version of the Howitt-Warren flow which a.s. satisfies the Chapman-Kolmogorov

equation (condition (i′) in Remark 6.5.2), we will sample a collection of coalescing paths $\mathcal{W}$ given the environment $(\mathcal{N}_{\beta_-,\beta_+},(\omega_z)_{z\in S})$.

Theorem 6.5.7 (Constructing Howitt-Warren flows in a Brownian net) *Let $\beta\in\mathbb{R}$ and let ν be a finite measure on $[0,1]$ with $\int q^{-1}(1-q)^{-1}\nu(\mathrm{d}q)<\infty$. Let $\mathcal{N}_{\beta_-,\beta_+}$ be a Brownian net with drifts β_-,β_+ defined as in* (6.39)*, and let S be its set of separation points. Conditional on $\mathcal{N}_{\beta_-,\beta_+}$, let $\omega:=(\omega_z)_{z\in S}$ be i.i.d. $[0,1]$-valued random variables with law $\bar{\nu}$ defined as in* (6.38)*. Conditional on $(\mathcal{N}_{\beta_-,\beta_+},\omega)$, let $(\alpha_z)_{z\in S}$ be independent $\{-1,+1\}$-valued random variables such that $\mathbb{P}[\alpha_z=1\,|\,(\mathcal{N}_{\beta_-,\beta_+},\omega)]=\omega_z$. Then*

$$\mathcal{W}:=\{\pi\in\mathcal{N}_{\beta_-,\beta_+}:\mathrm{sgn}_\pi(z)=\alpha_z\ \forall z\in S\ \textit{s.t. }\pi\ \textit{enters }z\} \tag{6.40}$$

is distributed as a Brownian web with drift β. For any $z=(x,t)\in\mathbb{R}^2$, if z is of type (1,2)*, then let $\pi_z^{\uparrow}\in\mathcal{W}(z)$ be any path in $\mathcal{W}$ entering z restricted to the time interval $[t,\infty)$; otherwise let $\pi_z^{\uparrow}$ be the rightmost path in $\mathcal{W}(z)$. Then*

$$K^{\uparrow}_{s,t}(x,\cdot):=\mathbb{P}[\pi^{\uparrow}_{(x,s)}(t)\in\cdot\,|\,(\mathcal{N}_{\beta_-,\beta_+},\omega)],\qquad s\le t,\ x\in\mathbb{R} \tag{6.41}$$

defines a version of the Howitt-Warren flow with drift β and characteristic measure ν, which satisfies condition (i′) *in Remark 6.5.2.*

General case: ν is any finite measure on $[0,1]$. Let $(\mu_\varepsilon)_{\varepsilon>0}$ satisfy (6.33) and let $\omega^\varepsilon:=(\omega_z^\varepsilon)_{z\in\mathbb{Z}^2_{\mathrm{even}}}$ be an i.i.d. random environment with common law μ_ε. Without assuming $\int q^{-1}(1-q)^{-1}\nu(\mathrm{d}q)<\infty$, it may no longer be possible to capture the continuum random environment by a Brownian net, because either β_- or β_+ in (6.39) or both can be infinity. Instead, we follow the alternative view on the coupling between branching-coalescing and coalescing random walks in Section 6.4, where we first sample the coalescing walks and then introduce the branching points. The same approach can be applied here to encode the random environment ω in two steps.

First we construct a collection of coalescing random walks by sampling coalescing walks in the same environment ω and then average over the law of ω^ε. More precisely, from each $z=(x,t)\in\mathbb{Z}^2_{\mathrm{even}}$, the walk moves to either $(x+1,t+1)$ or $(x-1,t+1)$, represented by $\alpha_z^\varepsilon:=1$ or $\alpha_z^\varepsilon=-1$, with respective probability $\int_0^1 q\mu_\varepsilon(\mathrm{d}q)$ and $\int_0^1(1-q)\mu_\varepsilon(\mathrm{d}q)$. Let W^ε denote this collection of coalescing random walks, which we call a *discrete web* and which has a natural dual $\widehat{W}^\varepsilon$ as illustrated in Figure 6.1. Next, we identify the law of ω^ε conditioned on W^ε.

Note that conditioned on W^ε, $(\omega_z^\varepsilon)_{z\in\mathbb{Z}^2_{\mathrm{even}}}$ are independent with conditional distribution

$$\mu_\varepsilon^{\mathrm{l}}:=\frac{(1-q)\mu_\varepsilon(\mathrm{d}q)}{\int(1-q)\mu_\varepsilon(\mathrm{d}q)}\quad\text{if }\alpha_z^\varepsilon=-1;\qquad \mu_\varepsilon^{\mathrm{r}}:=\frac{q\mu_\varepsilon(\mathrm{d}q)}{\int q\mu_\varepsilon(\mathrm{d}q)}\quad\text{if }\alpha_z^\varepsilon=1. \tag{6.42}$$

Therefore conditioned on W^ε, if we sample i.i.d. random variables $(\eta_z^\varepsilon)_{\{z:\alpha_z^\varepsilon=-1\}}$ with common law $\mu_\varepsilon^{\mathrm{l}}$ and i.i.d. random variables $(1-\eta_z^\varepsilon)_{\{z:\alpha_z^\varepsilon=1\}}$ with common law $\mu_\varepsilon^{\mathrm{r}}$, then $(W^\varepsilon, \eta^\varepsilon)$ provides an alternative representation for the environment ω^ε, where at each $z \in \mathbb{Z}^2_{\mathrm{even}}$, a walk in the random environment follows the same jump as in W^ε with probability $1-\eta_z^\varepsilon$, and jumps in the opposite direction with probability η_z^ε (see Figure 6.13(b)).

The above setup is an extension of the coupling between coalescing and branching-coalescing random walks discussed in Section 6.4, which corresponds to taking $(\eta_z^\varepsilon)_{z\in\mathbb{Z}^2_{\mathrm{even}}}$ to be i.i.d. $\{0,1\}$-valued random variables with mean ε. As we rescale space and time by ε and ε^2 respectively, we note that W^ε converges to a Brownian web $\mathcal{W}_0$ with drift β. On the other hand, η^ε can be seen to converge to a marked Poisson point process on the set of intersection points between paths in $\mathcal{W}_0$ and paths in its dual $\widehat{\mathcal{W}_0}$, i.e., the $(1,2)$ points of $\mathcal{W}_0$. Indeed, we can regard $(\eta_z^\varepsilon)_{\{z:\alpha_z^\varepsilon=-1\}}$ as a marked point process on the set of space-time points in $\mathbb{Z}^2_{\mathrm{even}}$ where a random walk path in W^ε is exactly one unit of distance to the left of a dual random walk path in $\widehat{W}^\varepsilon$. As noted in Section 6.4, when space-time is rescaled by $(\varepsilon,\varepsilon^2)$ and measure is rescaled by ε, the counting measure on these points of collision between forward and dual random walk paths converges to the intersection local time measure ℓ_{l} on points of type $(1,2)_{\mathrm{l}}$ in the Brownian web. Since for every $u \in (0,1)$, by (6.33), the mark η_z^ε at each $z \in \mathbb{Z}^2_{\mathrm{even}}$ with $\alpha_z^\varepsilon = -1$ satisfies

$$\begin{aligned}\mathbb{P}(\eta_z^\varepsilon \in (u,1]) &= \frac{\int_u^1 (1-q)\mu_\varepsilon(\mathrm{d}q)}{\int_0^1 (1-q)\mu_\varepsilon(\mathrm{d}q)} = \frac{\varepsilon\int_u^1 q^{-1}\varepsilon^{-1}q(1-q)\mu_\varepsilon(\mathrm{d}q)}{\int_0^1 (1-q)\mu_\varepsilon(\mathrm{d}q)} \\ &\sim \varepsilon\int_u^1 \frac{2}{q}\nu(\mathrm{d}q) \quad \text{as } \varepsilon \downarrow 0,\end{aligned}$$

it follows that $(\eta_z^\varepsilon)_{\{z:\alpha_z^\varepsilon=-1\}}$ converges to a marked Poisson point process $\mathcal{M}_{\mathrm{l}} \subset \mathbb{R}^2 \times [0,1]$ with intensity measure $\ell_{\mathrm{l}}(\mathrm{d}z) \otimes \frac{2}{q}1_{\{q>0\}}\nu(\mathrm{d}q)$. Similarly, $(1-\eta_z^\varepsilon)_{\{z:\alpha_z^\varepsilon=+1\}}$ converges to a marked Poisson point process $\mathcal{M}_{\mathrm{r}} \subset \mathbb{R}^2 \times [0,1]$ with intensity measure $\ell_{\mathrm{r}}(\mathrm{d}z) \otimes \frac{2}{1-q}1_{\{q<1\}}\nu(\mathrm{d}q)$.

The triple $(\mathcal{W}_0, \mathcal{M}_{\mathrm{l}}, \mathcal{M}_{\mathrm{r}})$ then gives the desired continuum space-time random environment, wherein a random motion π must navigate along paths in $\mathcal{W}_0$, and independently at each marked $(1,2)_{\mathrm{l}}$ point $(z,\eta_z) \in \mathcal{M}_{\mathrm{l}}$ or marked $(1,2)_{\mathrm{r}}$ point $(z,1-\eta_z) \in \mathcal{M}_{\mathrm{r}}$ it encounters, with probability η_z, π chooses orientation $-\mathrm{sgn}_{\mathcal{W}_0}(z)$, i.e., π switches to the second outgoing path in $\mathcal{W}_0(z)$ instead of continuing along the incoming path. This description is correct when $\nu(\{0\}) = \nu(\{1\}) = 0$. However when $\nu(\{0\}) > 0$, we have excluded $\ell_{\mathrm{l}}(\mathrm{d}z) \otimes \frac{2}{q}1_{\{q=0\}}\nu(\mathrm{d}q)$ from the intensity measure of $\mathcal{M}_{\mathrm{l}}$, which has a non-negligible effect on paths

sampled in the random environment. To understand this effect, we can approximate $\nu(\{0\})\delta_0(\mathrm{d}q)$ by $\nu(\{0\})\delta_h(\mathrm{d}q)$ with $h \downarrow 0$. This leads to Poisson marking $(1,2)_{\mathrm{l}}$ points of $\mathcal{W}_0$ with intensity measure $\frac{2}{h}\nu(\{0\})\ell_{\mathrm{l}}(\mathrm{d}z)$, and whenever a random motion π in this environment encounters such a Poisson point z, it chooses its orientation at z to be $-\mathrm{sgn}_{\mathcal{W}_0}(z)$ with probability h. The net effect is that, at each point z of a Poisson point process B_{l} with intensity measure $2\nu(\{0\})\ell_{\mathrm{l}}(\mathrm{d}z)$, π chooses its orientation at z to be $-\mathrm{sgn}_{\mathcal{W}_0}(z)$ with probability 1. As $h \downarrow 0$, the resulting effect is that, when we sample an independent motion π' in the same random environment, then an independent copy of B_{l} – call it B'_{l} – must be sampled such that whenever π' encounters some $z \in B'_{\mathrm{l}}$, it chooses its orientation at z to be $-\mathrm{sgn}_{\mathcal{W}_0}(z)$ with probability 1.

We will now give a precise formulation of how the Howitt-Warren flow can be obtained from a random environment $(\mathcal{W}_0, \mathcal{M}_{\mathrm{l}}, \mathcal{M}_{\mathrm{r}})$ (cf. [92, Theorem 3.7]). As in Theorem 6.5.7, we will sample a collection of coalescing Brownian motions $\mathcal{W}$ given the environment $(\mathcal{W}_0, \mathcal{M}_{\mathrm{l}}, \mathcal{M}_{\mathrm{r}})$.

Theorem 6.5.8 (Construction of Howitt-Warren flows) *Let $\beta \in \mathbb{R}$ and let ν be a finite measure on $[0,1]$. Let $\mathcal{W}_0$ be a Brownian web with drift β. Let $\mathcal{M}$ be a marked Poisson point process on $\mathbb{R}^2 \times [0,1]$ with intensity measure*

$$\ell_{\mathrm{l}}(\mathrm{d}z) \otimes \frac{2}{q} 1_{\{q>0\}} \nu(\mathrm{d}q) + \ell_{\mathrm{r}}(\mathrm{d}z) \otimes \frac{2}{1-q} 1_{\{q<1\}} \nu(\mathrm{d}q).$$

Conditional $(\mathcal{W}_0, \mathcal{M})$, let α_z be independent $\{-1,+1\}$-valued random variables with $\mathbb{P}[\alpha_z = +1 | (\mathcal{W}_0, \mathcal{M})] = \omega_z$ for each $(z,\omega_z) \in \mathcal{M}$, and let

$$A := \{z : (z,\omega_z) \in \mathcal{M}, \alpha_z \neq \mathrm{sgn}_{\mathcal{W}_0}(z)\}.$$

Let B be an independent Poisson point set with intensity $2\nu_{\mathrm{l}}(\{0\})\ell_{\mathrm{l}} + 2\nu_{\mathrm{r}}(\{1\})\ell_{\mathrm{r}}$. Define

$$\mathcal{W} := \lim_{\Delta_n \uparrow A \cup B} \mathrm{switch}_{\Delta_n}(\mathcal{W}_0) \tag{6.43}$$

for any sequence of finite sets $\Delta_n \uparrow A \cup B$, where $\mathrm{switch}_{\Delta_n}(\mathcal{W}_0)$ is the set of paths obtained from $\mathcal{W}_0$ by redirecting all paths in $\mathcal{W}$ entering any $z \in \Delta_n$ in such a way that z of type $(1,2)_{\mathrm{l}}$ in $\mathcal{W}$ becomes type $(1,2)_{\mathrm{r}}$ in $\mathcal{W}_0$ and vice versa. Then $\mathcal{W}$ is equally distributed with $\mathcal{W}_0$, and with $\pi_z^{\uparrow}$ defined as in Theorem 6.5.7,

$$K^{\uparrow}_{s,t}(x,\cdot) := \mathbb{P}[\pi^{\uparrow}_{(x,s)}(t) \in \cdot \,|\, (\mathcal{W}_0, \mathcal{M})] \qquad s \leq t,\ x \in \mathbb{R}, \tag{6.44}$$

defines a version of the Howitt-Warren flow with drift β and characteristic measure ν, which satisfies condition (i′) *in Remark 6.5.2.*

Exercise 6.5.9 *Use Theorem 6.4.2 to show that $\mathcal{W}$ as defined in* (6.43) *does not depend on the choice of $\Delta_n \uparrow A \cup B$, and $\mathcal{W}$ is equally distributed with $\mathcal{W}_0$.*

6.5.3 Properties of the Howitt-Warren Flow

The construction of the continuum space-time random environment underlying the Howitt-Warren flows allows the study of almost sure properties of the flow $K^{\uparrow} := (K^{\uparrow}_{s,t}(x,\cdot))_{s<t,x\in\mathbb{R}}$ (cf. [92, Section 2]). In particular, we can study almost sure path properties of the following measure-valued process induced by the Howitt-Warren flow, called the *Howitt-Warren process*:

$$\rho_t(\mathrm{d}y) = \int K^{\uparrow}_{0,t}(x,\mathrm{d}y)\rho_0(\mathrm{d}x), \tag{6.45}$$

where ρ_0 can be taken to be any locally finite measure on $\mathbb{R}$ in the class

$$\mathcal{M}_{\mathrm{g}}(\mathbb{R}) := \Big\{\rho : \int e^{-cx^2}\rho(\mathrm{d}x) < \infty \text{ for all } c > 0\Big\}, \tag{6.46}$$

where $\rho_n \in \mathcal{M}_{\mathrm{g}}(\mathbb{R})$ is defined to converge to $\rho \in \mathcal{M}_{\mathrm{g}}(\mathbb{R})$ if $\int f(x)\rho_n(\mathrm{d}x) \to \int f(x)\rho(\mathrm{d}x)$ for all $f \in C_c(\mathbb{R})$, which is just convergence in the vague topology, plus $\int e^{-cx^2}\rho_n(\mathrm{d}x) \to \int e^{-cx^2}\rho(\mathrm{d}x)$ for all $c > 0$.

A first consequence is that $(\rho_t)_{t\geq 0}$ is a Markov process with continuous sample path, and continuous dependence on the initial condition and starting time, which makes it a Feller process.

Theorem 6.5.10 (Howitt-Warren process) *Let $(K^{\uparrow}_{s,t}(x,\cdot))_{s\leq t,x\in\mathbb{R}}$ be the version of Howitt-Warren flow with drift β and characteristic measure ν, defined in Theorem 6.5.7 or 6.5.8. Let $(\rho_t)_{t\geq 0}$ be the Howitt-Warren process defined from $K^{\uparrow}$ as in* (6.45) *with $\rho_0 \in \mathcal{M}_{\mathrm{g}}$. Then*

(i) *$(\rho_t)_{t\geq 0}$ is an $\mathcal{M}_{\mathrm{g}}(\mathbb{R})$-valued Markov process with almost sure continuous sample paths;*

(ii) *If $(\rho^{\langle n\rangle}_t)_{t\geq s_n}$ are Howitt-Warren processes defined from $K^{\uparrow}$ with deterministic initial condition $\rho^{\langle n\rangle}_{s_n}$ at time s_n, with $s_n \to 0$, then for any $t > 0$ and $t_n \to t$,*

$$\rho^{\langle n\rangle}_{s_n} \underset{n\to\infty}{\Longrightarrow} \rho_0 \quad \textit{implies} \quad \rho^{\langle n\rangle}_{t_n} \underset{n\to\infty}{\Longrightarrow} \rho_t \qquad a.s., \tag{6.47}$$

where $\Rightarrow$ denotes convergence in $\mathcal{M}_{\mathrm{g}}(\mathbb{R})$.

If $\int q^{-1}(1-q)^{-1}\nu(\mathrm{d}q) < \infty$, then the above statements hold with $\mathcal{M}_{\mathrm{g}}(\mathbb{R})$ replaced by $\mathcal{M}_{\mathrm{loc}}(\mathbb{R})$, the space of locally finite measures on $\mathbb{R}$ equipped with the vague topology.

We can also identify almost surely the support of ρ_t for all $t \geq 0$.

Theorem 6.5.11 (Support of Howitt-Warren process) *Let $(\rho_t)_{t\geq 0}$ be a Howitt-Warren process with drift β and characteristic measure, and the initial condition ρ_0 has compact support. Let β_-, β_+ be defined as in* (6.39).

(a) *If $-\infty < \beta_- < \beta_+ < \infty$, then a.s. for all $t > 0$, the support of ρ_t satisfies*

$$\mathrm{supp}(\rho_t) = \{\pi(t) : \pi \in \mathcal{N}_{\beta_-,\beta_+}, \pi(0) \in \mathrm{supp}(\rho_0)\}, \tag{6.48}$$

where $\mathcal{N}_{\beta_-,\beta_+}$ is the Brownian net with drift parameters β_-, β_+ as in Theorem 6.5.7.

(b) *If $\beta_- = -\infty$ and $\beta_+ < \infty$, then a.s. $\mathrm{supp}(\rho_t) = (-\infty, r_t] \cap \mathbb{R}$ for all $t > 0$, where $r_t := \sup(\mathrm{supp}(\rho_t))$. An analogue statement holds when $\beta_- > -\infty$ and $\beta_+ = \infty$.*

(c) *If $\beta_- = -\infty$ and $\beta_+ = \infty$, then a.s. $\mathrm{supp}(\rho_t) = \mathbb{R}$ for all $t > 0$.*

Theorem 6.5.11 (a) shows that when the Howitt-Warren flow can be constructed from a Brownian net, then at deterministic times, ρ_t is almost surely atomic. This result can be extended to general Howitt-Warren flows. However, almost surely there exist random times when ρ_t contains no atoms, and the only exceptions are the Arratia flow with $\nu = 0$, and the *erosion flows*, which have characteristic measures of the form $\nu = a\delta_0 + b\delta_1$ with $a, b \geq 0$ and $a + b > 0$.

Theorem 6.5.12 (Atomicness vs non-atomicness) *Let $(\rho_t)_{t\geq 0}$ be a Howitt-Warren process with drift β and characteristic measure ν.*

(a) *For each $t > 0$, ρ_t is a.s. purely atomic.*

(b) *If $\int_{(0,1)} \nu(\mathrm{d}q) > 0$, then a.s. there exists a dense set of random times $t > 0$ when ρ_t contains no atoms.*

(c) *If $\int_{(0,1)} \nu(\mathrm{d}q) = 0$, then a.s. ρ_t is purely atomic at all $t > 0$.*

We can also study ergodic properties of the Howitt-Warren process. In particular, for any Howitt-Warren process other than the measure-valued process generated by the coalescing Arratia flow, there is a unique spatially ergodic stationary law for $(\rho_t)_{t\geq 0}$, which is also the weak limit of ρ_t as $t \to \infty$ if the law of the initial condition ρ_0 is spatially ergodic with finite mean density.

Theorem 6.5.13 (Ergodic properties) *Let $(\rho_t)_{t\geq 0}$ be a Howitt-Warren process with drift $\beta \in \mathbb{R}$, characteristic measure $\nu \neq 0$, and initial law $\mathcal{L}(\rho_0)$.*

(i) *If $\mathcal{L}(\rho_0)$ is ergodic w.r.t. $T_a\rho_0(\cdot) = \rho_0(a+\cdot)$ for all $a \in \mathbb{R}$, and $\mathbb{E}[\rho_0([0,1])] = 1$, then as $t \to \infty$, $\mathcal{L}(\rho_t)$ converges weakly to a limit Λ_1 which is also ergodic with respect to T_a for all $a \in \mathbb{R}$. Furthermore, if $\mathcal{L}(\rho_0) = \Lambda_1$, then $\mathbb{E}[\rho_0([0,1])] = 1$ and $\mathcal{L}(\rho_t) = \Lambda_1$ for all $t > 0$.*

(ii) *If $\mathcal{L}(\rho_0)$ is ergodic w.r.t. T_a for all $a \in \mathbb{R}$ and $\mathbb{E}[\rho_0([0,1])] = \infty$, then as $t \to \infty$, $\mathcal{L}(\rho_t)$ has no subsequential weak limits supported on $\mathcal{M}_{\mathrm{loc}}(\mathbb{R})$, the set of locally finite measures on $\mathbb{R}$.*

Remark 6.5.14 When $\nu(\mathrm{d}x) = 1_{[0,1]}(x)\mathrm{d}x$, it is known from [64, Proposition 9(b)] that Λ_1 is the law of a random measure $\rho^* = \sum_{(x,u)\in\mathcal{P}} u\delta_x$ for a Poisson point process $\mathcal{P}$ on $\mathbb{R}\times[0,\infty)$ with intensity measure $\mathrm{d}x \times u^{-1}e^{-u}\mathrm{d}u$.

The proof of the above results can be found in [92].

Remark 6.5.15 Howitt-Warren flows are the continuum analogues of discrete random walks in i.i.d. space-time random environments, and they share the same fluctuations on large space-time scales. In particular, for any Howitt-Warren process $(\rho_t)_{t\geq 0}$ (assuming drift $\beta = 0$ and $\rho_0(\mathrm{d}x) = \mathrm{d}x$ for simplicity), the rescaled current process $(\mathcal{I}(nt, \sqrt{n}x)/n^{1/4})_{t>0,x\in\mathbb{R}}$, where

$$\mathcal{I}(t,x) := \int_{\mathbb{R}}\int_R 1_{\{u<0\}}1_{\{v>x\}}K^{\uparrow}_{0,t}(u,\mathrm{d}v)\rho_0(\mathrm{d}u) - \int_{\mathbb{R}}\int_R 1_{\{u>0\}}1_{\{v<x\}}K^{\uparrow}_{0,t}(u,\mathrm{d}v)\rho_0(\mathrm{d}u),$$

converges to a universal Gaussian process as $n\to\infty$. This was shown in [106], and the same universal (Edwards-Wilkinson) fluctuations have been established earlier for random walks in i.i.d. space-time random environments (see e.g., [85]). In a different direction, the transition probabilities of a random walk in i.i.d. space-time random environments are believed to have the same universal (Tracy-Widom GUE) fluctuations as the point-to-point partition functions of a directed polymer. This was verified recently in [12] for special Beta-distributed random environments, and one expects similar results to hold for the Howitt-Warren flow $K^{\uparrow}_{0,t}(0,\mathrm{d}x)$.

6.6 Convergence to the Brownian Web and Net

In this section, we give general convergence criteria for the Brownian web that were originally formulated in [40, 72], and simplified criteria when paths do not cross each other. We will discuss strategies for verifying these criteria and focus in particular on the convergence of coalescing random walks to the Brownian web. Lastly we will formulate a set of convergence criteria for the Brownian net and verify them for branching-coalescing simple random walks.

6.6.1 General Convergence Criteria for the Brownian Web

We give here general convergence criteria for a sequence of random variables $(X_n)_{n\in\mathbb{N}}$, taking values in the space of compact sets of paths $\mathcal{H}$, to converge in distribution to the Brownian web $\mathcal{W}$. When X_n consists of non-crossing paths, these criteria can be greatly simplified, which will be discussed in the next subsection.

First we formulate a criterion which ensures tightness for the laws of $(X_n)_{n\in\mathbb{N}}$. We then formulate criteria which ensure that any subsequential limit of $(X_n)_{n\in\mathbb{N}}$ contains a copy of the Brownian web (lower bound), but nothing more (upper bound).

Tightness: To understand the tightness criterion we will formulate, let us first see what should be the tightness criterion for a sequence of path-valued random variables $(Y_n)_{n\in\mathbb{N}}$, i.e., $Y_n \in \Pi$ with the space of paths Π defined as in Section 6.2.1. Due to the compactification of $\mathbb{R}^2$ in Section 6.2.1, it suffices to show that when restricted to any finite space-time window $\Lambda_{L,T} = [-L,L] \times [-T,T]$, the law of the random paths $(Y_n)_{n\in\mathbb{N}}$ are tight in the sense that for any $\varepsilon > 0$, there exists a modulus of continuity $\phi : [0,1] \to \mathbb{R}$, which is increasing with $\phi(\delta) \downarrow 0$ as $\delta \downarrow 0$, such that uniformly in $n \in \mathbb{N}$,

$$\mathbb{P}(\forall s \text{ with } (Y_n(s),s) \in \Lambda_{L,T} \text{ and } t \in [s,s+1],\ |Y_n(t) - Y_n(s)| \le \phi(t-s)) \ge 1-\varepsilon. \tag{6.49}$$

The modulus of continuity ϕ allows the construction of an equicontinuous, and hence compact, set of paths. To construct a ϕ that satisfies (6.49), it suffices to show that for any $\eta > 0$,

$$\lim_{\delta\downarrow 0} \limsup_{n\to\infty} \mathbb{P}(\exists s < t < s+\delta \text{ with } (Y_n(s),s) \in \Lambda_{L,T},\ s.t.\ |Y_n(t) - Y_n(s)| > \eta) = 0. \tag{6.50}$$

Indeed, fix a sequence $\eta_m \downarrow 0$. Then for each $m \in \mathbb{N}$, by (6.50), we can find $\delta_m > 0$ sufficiently small such that uniformly in $n \in \mathbb{N}$, the probability in (6.50) is bounded by $\varepsilon/2^m$. We can then define $\phi(h) := \eta_m$ for $h \in (\delta_{m-1}, \delta_m]$, which is easily seen to satisfy (6.49).

When we consider a sequence of random compact sets of paths $(X_n)_{n\in\mathbb{N}}$, the tightness criterion is similar, except that we need to control the modulus of continuity uniformly for all paths in X_n (cf. Exercise 6.2.2). Therefore condition (6.50) should be modified to show that for any finite $\Lambda_{L,T}$ and any $\eta > 0$,

$$\begin{aligned}\lim_{\delta\downarrow 0} \limsup_{n\to\infty} \mathbb{P}(&\exists \pi \in X_n \text{ and } s < t < s+\delta \text{ with } (\pi(s),s) \\ &\in \Lambda_{L,T},\ s.t.\ |\pi(t) - \pi(s)| > \eta) = 0.\end{aligned} \tag{6.51}$$

To control the modulus of continuity of all paths in X_n simultaneously, it is convenient to divide $\Lambda_{L,T}$ into $16LT/\delta\eta$ sub-rectangles of dimension $\eta/2 \times \delta/2$, and bound the event in (6.51) by the union of events where $\Lambda_{L,T}$ in (6.51) is replaced by one of the $16LT/\delta\eta$ sub-rectangles of $\Lambda_{L,T}$. This leads to the following tightness criterion as formulated in [40, Proposition B1].

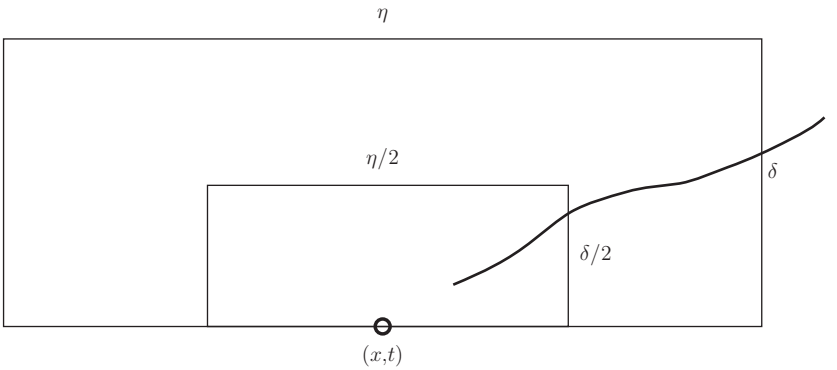

Figure 6.14. A path causing the event $A_{\delta,\eta}(x,t)$ to occur.

Proposition 6.6.1 (Tightness criterion) *The law of a sequence of $(\mathcal{H},\mathcal{B}_{\mathcal{H}})$-valued random variables $(X_n)_{n\in\mathbb{N}}$ is tight if*

$$\text{(T)} \qquad \forall L,T\in(0,\infty),\quad \lim_{\delta\downarrow 0}\delta^{-1}\limsup_{n\to\infty}\sup_{(x,t)\in[-L,L]\times[-T,T]}\mathbb{P}(X_n\in A_{\delta,\eta}(x,t))=0,$$

where $A_{\delta,\eta}(x,t)\in\mathcal{B}_{\mathcal{H}}$ consists of compact sets of paths $K\in\mathcal{H}$, such that K contains some path which intersects the rectangle $[x-\eta/4,x+\eta/4]\times[t,t+\delta/2]$, and at a later time, intersects the left or right boundary of the bigger rectangle $[x-\eta/2,x+\eta/2]\times[t,t+\delta]$ (see Figure 6.14).

Lower bound: Assuming that $(X_n)_{n\in\mathbb{N}}$ is a tight sequence of $\mathcal{H}$-valued random variables, we then need a criterion to ensure that any subsequential weak limit of $(X_n)_{n\in\mathbb{N}}$ contains almost surely a random subset which is distributed as the standard Brownian web. The following criterion serves this purpose, which is a form of convergence in finite-dimensional distributions.

(I) There exists $\pi_{n,z}\in X_n$ for each $z\in\mathbb{R}^2$, such that for any deterministic $z_1,\dots,z_k\in\mathbb{R}^2$, $(\pi_{n,z_i})_{1\le i\le k}$ converge in distribution to coalescing Brownian motions starting at $(z_i)_{1\le i\le k}$.

If $(X_n)_{n\in\mathbb{N}}$ is a tight sequence that satisfies condition (I), then for any deterministic countable dense set $\mathcal{D}\subset\mathbb{R}^2$, we note that by going to a further subsequence if necessary, any subsequential limit $\mathcal{X}$ of $(X_n)_{n\in\mathbb{N}}$ contains a collection of coalescing Brownian motions $\{\pi_z\}_{z\in\mathcal{D}}$ starting from each $z\in\mathcal{D}$. Since $\mathcal{W}:=\overline{\{\pi_z:z\in\mathcal{D}\}}$ is a standard Brownian web by Theorem 6.2.3, we obtain the desired lower bound that a.s. $\mathcal{W}\subset\mathcal{X}$.

Upper bound: Assuming that $(X_n)_{n\in\mathbb{N}}$ is a tight sequence of $\mathcal{H}$-valued random variables that satisfies condition (I), then it only remains to formulate a criterion to ensure that any subsequential weak limit $\mathcal{X}$ of $(X_n)_{n\in\mathbb{N}}$ a.s. contains no more paths than the Brownian web $\mathcal{W} \subset \mathcal{X}$. There are several approaches to this problem, depending partly on whether paths in X_n can cross each other or not.

One way is to control the expectation of the following family of counting random variables

$$\eta_{\mathcal{X}}(t,h;a,b) = |\{\pi(t+h) : \pi \in \mathcal{X}, \pi(t) \in [a,b]\}|, \qquad t \in \mathbb{R}, h > 0, a < b, \quad (6.52)$$

which considers all paths in $\mathcal{X}$ that intersect $[a,b]$ at time t and counts the number of distinct positions these paths occupy at time $t+h$. Thanks to the image set property of the Brownian web $\mathcal{W}$, which is a special case of Proposition 6.3.17 for the Brownian net, it is easily seen that if $\mathcal{X}$ contains strictly more paths than $\mathcal{W}$, than we must have $\eta_{\mathcal{X}}(t,h;a,b) > \eta_{\mathcal{W}}(t,h;a,b)$ for some rational t,h,a,b. Therefore to show $\mathcal{X} = \mathcal{W}$ a.s., it suffices to show that

$$\mathbb{E}[\eta_{\mathcal{X}}(t,h;a,b)] = \mathbb{E}[\eta_{\mathcal{W}}(t,h;a,b)] \qquad \text{for all } t \in \mathbb{R}, h > 0, a < b. \quad (6.53)$$

The following sufficient criteria have been formulated in [40] to ensure that (6.53) holds for any subsequential limit $\mathcal{X}$ of a sequence of $\mathcal{H}$-valued random variables $(X_n)_{n\in\mathbb{N}}$:

$$(\mathrm{B1}') \quad \forall h_0 > 0, \limsup_{n\to\infty} \sup_{h>h_0} \sup_{a,t\in\mathbb{R}} \mathbb{P}[\eta_{X_n}(t,h;a,a+\varepsilon) \geq 2] \underset{\varepsilon\to 0}{\longrightarrow} 0,$$

$$(\mathrm{B2}') \quad \forall h_0 > 0, \frac{1}{\varepsilon} \limsup_{n\to\infty} \sup_{h>h_0} \sup_{a,t\in\mathbb{R}} \mathbb{P}[X_n(t,h;a,\varepsilon) \neq X_n^-(t,h;a,\varepsilon) \cup X_n^+(t,h;a,\varepsilon)]$$

$$\underset{\varepsilon\to 0}{\longrightarrow} 0,$$

where $X_n(t,h;a,\varepsilon) \subset \mathbb{R}$ is the set of positions occupied at time $t+h$ by paths in X_n which intersect the interval $[a,a+\varepsilon]$ at time t, while $X_n^-(t,h;a,\varepsilon)$ (resp. $X_n^+(t,h;a,\varepsilon)$) is the subset of $X_n(t,h;a,\varepsilon)$ induced by paths in X_n which occupy the leftmost (resp. rightmost) position at time t among all paths in X_n that intersect $[a,a+\varepsilon]$ at time t. Condition (B1′) ensures that for each deterministic point $z \in \mathbb{R}^2$, any subsequential limit $\mathcal{X}$ contains a.s. at most one path starting from z. Condition (B2′) ensures that $\eta_{\mathcal{X}}(t,h;a,b)$ can be approximated by partitioning $[a,b]$ into equal-sized intervals and considering only paths in $\mathcal{X}$ starting at the boundaries of these intervals. Condition (B1′) and (I) together ensure that paths in $\mathcal{X}$ starting at these boundary points are distributed as coalescing Brownian motions. Condition (6.53) then follows.

We thus have the following convergence criteria for the Brownian web [40, Theorem 5.1].

Theorem 6.6.2 (Convergence criteria A) *Let $(X_n)_{n\in\mathbb{N}}$ be a sequence of $(\mathcal{H},\mathcal{B}_{\mathcal{H}})$-valued random variables satisfying conditions* (T), (I), (B1′), (B2′). *Then X_n converges in distribution to the standard Brownian web $\mathcal{W}$.*

Condition (B2′) turns out to be difficult to verify when paths in X_n can cross each other. This is in particular the case for non-nearest neighbor coalescing random walks on $\mathbb{Z}$, which led to the formulation in [72] of an alternative criterion in place of (B2′). The observation is that instead of $\eta_{\mathcal{X}}(t,h;a,b)$, we can consider the alternative family of counting random variables

$$\hat{\eta}_{\mathcal{X}}(t,h;a,b) = |\{\pi(t+h)\cap(a,b) : \pi\in\mathcal{X}, \pi(t)\in\mathbb{R}\}|, \qquad t\in\mathbb{R}, h>0, a<b, \quad (6.54)$$

which considers all paths in $\mathcal{X}$ starting before or at time t and counts the number of distinct positions these paths occupy in the interval (a,b) at time $t+h$. Given $\mathcal{X}\supset\mathcal{W}$, to show that $\mathcal{X}=\mathcal{W}$ a.s., it suffices to show that

$$\mathbb{E}[\hat{\eta}_{\mathcal{X}}(t,h;a,b)] = \mathbb{E}[\hat{\eta}_{\mathcal{W}}(t,h;a,b)] \qquad \text{for all } t\in\mathbb{R}, h>0, a<b. \quad (6.55)$$

This leads to the following alternative convergence criteria in [72, Theorem 1.4].

Theorem 6.6.3 (Convergence criteria B) *Let $(X_n)_{n\in\mathbb{N}}$ be a sequence of $(\mathcal{H},\mathcal{B}_{\mathcal{H}})$-valued random variables satisfying conditions* (T), (I), (B1′) *and the following condition*:

(E) *For any subsequential weak limit $\mathcal{X}$, $\mathbb{E}[\hat{\eta}_{\mathcal{X}}(t,h;a,b)] = \mathbb{E}[\hat{\eta}_{\mathcal{W}}(t,h;a,b)]$ $\forall t\in\mathbb{R}, h>0, a<b$.*

Then X_n converges in distribution to the standard Brownian web $\mathcal{W}$ as $n\to\infty$.

Condition (E) is actually much easier to verify than it appears. It turns out to be enough to establish the density bound

$$\limsup_{n\to\infty} \mathbb{E}[\hat{\eta}_{X_n}(t,h;a,b)] < \infty \qquad \forall t\in\mathbb{R}, h>0, a<b, \quad (6.56)$$

which by Fatou's Lemma implies that any subsequential weak limit $\mathcal{X}$ satisfies $\mathbb{E}[\hat{\eta}_{\mathcal{X}}(t,\varepsilon;a,b)]<\infty$ for all $t\in\mathbb{R}$, $\varepsilon>0$ and $a<b$. In particular,

$$\mathcal{X}^{t-}(t+\varepsilon) := \{\pi(t+\varepsilon) : \pi\in\mathcal{X}^{t-}\} \quad \text{with} \quad \mathcal{X}^{t-} := \{\pi\in\mathcal{X} : \sigma_\pi \le t\}, \quad (6.57)$$

the set of positions at time $t+\varepsilon$ generated by paths in $\mathcal{X}$ starting before or at time t, is almost surely a locally finite subset of $\mathbb{R}$. As a random subset of $\mathbb{R}$, $\mathcal{X}^{t-}(t+\varepsilon)$ arises as the limit of $X_n^{t-}(t+\varepsilon)$,[1] and we can use Skorohod's

[1] In practice, $\mathcal{X}^{t-}(t+\varepsilon)$ may contain positions which arise from limits of paths in X_n that start at times $t_n\downarrow t$. Therefore we should consider instead $X_n^{(t+\delta)-}(t+\varepsilon)$ for some $\delta>0$, so that its limit contains $\mathcal{X}^{t-}(t+\varepsilon)$.

representation theorem [9] to couple them such that almost surely, $X_n^{t-}(t+\varepsilon)$ converges to $\mathcal{X}^{t-}(t+\varepsilon)$ w.r.t. the Hausdorff metric on subsets of $\mathbb{R}$. For Markov processes such as coalescing random walks, we expect that the law of paths in X_n^{t-} restricted to the time interval $[t+\varepsilon,\infty)$ depends only on their positions $X_n^{t-}(t+\varepsilon)$ at time $t+\varepsilon$, and furthermore, conditions (B1′) and (I) can be applied to these restricted paths conditioned on $X_n^{t-}(t+\varepsilon)$. Therefore given $X_n^{t-}(t+\varepsilon)$ converging to $\mathcal{X}^{t-}(t+\varepsilon)$, we can apply conditions (B1′) and (I) to conclude that $\mathcal{X}^{t-}$ restricted to the time interval $[t+\varepsilon,\infty)$ is a collection of coalescing Brownian motions starting from the locally finite set $\mathcal{X}^{t-}(t+\varepsilon)$ at time $t+\varepsilon$, and hence

$$\mathbb{E}[\hat{\eta}_{\mathcal{X}}(t,h;a,b)] \le \mathbb{E}[\hat{\eta}_{\mathcal{W}}(t+\varepsilon,h-\varepsilon;a,b)] = \frac{b-a}{\sqrt{\pi(h-\varepsilon)}}.$$

Sending $\varepsilon \downarrow 0$ then establishes condition (E). For non-nearest neighbor coalescing random walks, the above strategy was carried out in [72].

6.6.2 Convergence Criteria for Non-crossing Paths

When $(X_n)_{n\in\mathbb{N}}$ almost surely consists of paths that do not cross each other, i.e., X_n contains no paths π_1,π_2 with $(\pi_1(s)-\pi_2(s))(\pi_1(t)-\pi_2(t)) < 0$ for some $s<t$, tightness in fact follows from condition (I) [40, Proposition B2].

Proposition 6.6.4 (Tightness criterion for non-crossing paths) *If for each $n \in \mathbb{N}$, X_n is an $\mathcal{H}$-valued random variable consisting almost surely of paths that do not cross each other, and $(X_n)_{n\in\mathbb{N}}$ satisfies condition* (I)*, then $(X_n)_{n\in\mathbb{N}}$ is a tight family.*

This result holds because when paths do not cross, the modulus of continuity of all paths in X_n can be controlled by the modulus of continuity of paths in X_n starting at a grid of space-time points, similar to Step (2) in the proof sketch for Theorem 6.2.3.

When X_n consists of non-crossing paths, conditions (B1′) and (B2′) can also be simplified to

$$\text{(B1)} \qquad \forall h>0,\ \limsup_{n\to\infty} \sup_{a,t\in\mathbb{R}} \mathbb{P}[\eta_{X_n}(t,h;a,a+\varepsilon)\ge 2] \underset{\varepsilon\to 0}{\longrightarrow} 0,$$

$$\text{(B2)} \qquad \forall h>0,\ \frac{1}{\varepsilon}\limsup_{n\to\infty} \sup_{a,t\in\mathbb{R}} \mathbb{P}[\eta_{X_n}(t,h;a,a+\varepsilon)\ge 3] \underset{\varepsilon\to 0}{\longrightarrow} 0.$$

Recall from the discussions before Theorem 6.6.2 that condition (B1) is to ensure that for each deterministic $z\in\mathbb{R}^2$, any subsequential limit $\mathcal{X}$ of $(X_n)_{n\in\mathbb{N}}$ almost surely contains at most one path starting from z. This property is easily seen to be implied by condition (I) (by the same argument as for Theorem 6.2.3 (a))

when X_n consists of non-crossing paths, which implies that $\mathcal{X}$ also consists of non-crossing paths. Therefore condition (B1) also becomes redundant, which leads to the following simplification of Theorems 6.6.2 (cf. [40, Theorem 2.2]) and 6.6.3.

Theorem 6.6.5 (Convergence criteria C) *Let $(X_n)_{n\in\mathbb{N}}$ be a sequence of $(\mathcal{H},\mathcal{B}_{\mathcal{H}})$-valued random variables which a.s. consist of non-crossing paths. If $(X_n)_{n\in\mathbb{N}}$ satisfies conditions* (I)*, and either* (B2) *or* (E)*, then X_n converges in distribution to the standard Brownian web.*

Condition (B2) is often verified by applying the FKG positive correlation inequality [42], together with a bound on the distribution of the time of coalescence between two paths (see e.g., Section 6.6.3 below). However, FKG is a strong property that is not satisfied by most models. In such cases, verifying condition (B2) can be difficult. Besides checking condition (E), another alternative is to use the dual (a.k.a. wedge) characterization of the Brownian web, as noted in Remark 6.3.10, to upper bound any subsequential weak limit of $(X_n)_{n\in\mathbb{N}}$.[2]

Indeed, if $(X_n)_{n\in\mathbb{N}}$ consists of non-crossing paths and satisfies condition (I), then we can construct a collection of dual paths $\widehat{X}_n$ which almost surely do not cross paths in X_n, and the starting points of paths in $\widehat{X}_n$ become dense in $\mathbb{R}^2$ as $n\to\infty$. The tightness of $(X_n)_{n\in\mathbb{N}}$ is easily seen to imply the tightness of $(\widehat{X}_n)_{n\in\mathbb{N}}$, and any subsequential weak limit $(\mathcal{X},\widehat{\mathcal{X}})$ of $(X_n,\widehat{X}_n)$ must satisfy the property that: for any deterministic countable dense set $\mathcal{D}\subset\mathbb{R}^2$, $\mathcal{X}(\mathcal{D})$ is distributed as a collection of coalescing Brownian motions, which by the non-crossing property a.s. uniquely determines $\widehat{\mathcal{X}}(\mathcal{D})$, which is distributed as a collection of dual coalescing Brownian motions. By the wedge characterization of the Brownian web in Remark 6.3.10, $\mathcal{X}$ equals $\overline{\mathcal{X}(\mathcal{D})}$, a standard Brownian web, if no path in $\mathcal{X}$ enters any wedge $W(\hat{\pi}_1,\hat{\pi}_2)$ of $\widehat{\mathcal{X}}(\mathcal{D})$ from outside (defined as in (6.16) with $\widehat{\mathcal{X}}(\mathcal{D})$ replacing both $\widehat{\mathcal{W}}^{\mathrm{l}}$ and $\widehat{\mathcal{W}}^{\mathrm{r}}$). This leads to

(U) For each $n\in\mathbb{N}$, there exists $\widehat{X}_n\in\widehat{\mathcal{H}}$ whose paths a.s. do not cross those of X_n and whose starting points are dense in $\mathbb{R}^2$ as $n\to\infty$, s.t. for any subsequential weak limit $(\mathcal{X},\widehat{\mathcal{X}})$ of$(X_n,\widehat{X}_n)$ and any deterministic countable dense $\mathcal{D}\subset\mathbb{R}^2$, a.s. paths in $\mathcal{X}$ do not enter any wedge of $\widehat{\mathcal{X}}(\mathcal{D})$ from outside.

We then have the following convergence result.

Theorem 6.6.6 (Convergence criteria D) *Let $(X_n)_{n\in\mathbb{N}}$ be a sequence of $(\mathcal{H},\mathcal{B}_{\mathcal{H}})$-valued random variables consisting of non-crossing paths and which satisfy*

[2] Recently a new approach to verify condition (B2) was proposed in [90], using a Lyapunov function type criterion on the gaps between three non-crossing paths, assuming that the gaps evolve jointly as a Markov process.

conditions (I) *and* (U). *Then* X_n *converges in distribution to the standard Brownian web.*

Remark 6.6.7 To verify condition (U), it suffices to show that paths in X_n do not enter wedges of $\widehat{X}_n$ from outside, and for any deterministic $z_1, z_2 \in \mathbb{R}^2$, not only do there exist paths $\hat{\pi}_{n,1}, \hat{\pi}_{n,2} \in \widehat{X}_n$ which converge to dual coalescing Brownian motions starting at z_1 and z_2 (which follows from condition (I) for $(X_n)_{n\in\mathbb{N}}$ and the non-crossing between paths in X_n and $\widehat{X}_n$), but also the time of coalescence between $\hat{\pi}_{n,1}$ and $\hat{\pi}_{n,2}$ converges to that of the coalescing Brownian motions. This ensures that the wedge $W(\hat{\pi}_{n,1}, \hat{\pi}_{n,2})$ converges and no path in the limit can enter the wedge through its bottom point. The latter can also be accomplished by showing that no limiting forward and dual paths can spend positive Lebesgue time together, as carried out in [81].

6.6.3 Convergence of Coalescing Simple Random Walks to the Brownian Web

We now illustrate how the convergence criteria in Theorems 6.6.5 and 6.6.6 for non-crossing paths can be verified for the discrete time coalescing simple random walks on $\mathbb{Z}$ (cf. [40, Theorem 6.1]).

Let X denote the collection of discrete time coalescing simple random walk paths on $\mathbb{Z}$, with one walk starting from every space-time lattice site $z \in \mathbb{Z}^2_{\text{even}}$. It is an easy exercise to show that X is a.s. precompact in the space of paths Π, and with a slight abuse of notation, we will henceforth denote the closure of X in (Π, d) also by X.

For each $\varepsilon \in (0,1)$, let $S_\varepsilon : \mathbb{R}^2 \to \mathbb{R}^2$ denote the diffusive scaling map

$$S_\varepsilon(x,t) = (\varepsilon x, \varepsilon^2 t). \tag{6.58}$$

For a path $\pi \in \Pi$, let $S_\varepsilon \pi$ denote the path whose graph is the image of the graph of π under S_ε. For a set of paths K, define $S_\varepsilon K := \{S_\varepsilon \pi : \pi \in K\}$.

Theorem 6.6.8 *Let* X *be the collection of coalescing simple random walk paths on* $\mathbb{Z}$ *defined as above. Then as* $\varepsilon \downarrow 0$, $X_\varepsilon := S_\varepsilon X$ *converges in distribution to the standard Brownian web* $\mathcal{W}$.

Proof sketch. We show how the various conditions in Theorems 6.6.5 and 6.6.6 can be verified.

(I): This condition follows by Donsker's invariance principle. Indeed, coalescing random walks can be constructed from independent random walks by the same procedure as the inductive construction of coalescing Brownian motions from independent Brownian motions in the proof for Theorem 6.2.3. Furthermore, this

construction is a.s. continuous w.r.t. the independent Brownian motions. Therefore (I) follows from Donsker's invariance principle for independent random walks and the Continuous Mapping Theorem.

(B2): We will verify this condition by applying the FKG inequality [42]. By translation invariance in space-time, it suffices to show that for any $t > 0$,

$$\lim_{\delta \downarrow 0} \delta^{-1} \limsup_{\varepsilon \downarrow 0} \mathbb{P}(\eta_{X_\varepsilon}(0,t;0,\delta) \geq 3) = 0.$$

Assume w.l.o.g. that $t = 1$. Formulated in terms of X, and letting $A := \delta^{-1}$ and $\sqrt{n} = \delta \varepsilon^{-1}$, it is equivalent to showing that

$$\lim_{A \to \infty} A \limsup_{n \to \infty} \mathbb{P}(\eta_X(0, A^2 n; 0, \sqrt{n}) \geq 3) = 0. \tag{6.59}$$

For $i \in \mathbb{Z}$, let π_i denote the random walk starting at $2i$ at time 0, and let $\tau_{i,j}$ denote the first meeting (coalescence) time between π_i and π_j. Then by a decomposition according to the first index $k \in \mathbb{N}$ with $\tau_{k-1,k} > A^2 n$, we have

$$\mathbb{P}(|\{\pi_i(A^2 n) : 0 \leq i \leq \sqrt{n}\}| \geq 3) = \sum_{k=1}^{\sqrt{n}-1} \mathbb{P}(\tau_{0,k-1} \leq A^2 n, \tau_{k-1,k} > A^2 n, \tau_{k,\sqrt{n}} > A^2 n).$$

Let us restrict to the event $E_k := \{\tau_{0,k-1} \leq A^2 n, \tau_{k-1,\sqrt{n}} > A^2 n\}$ and condition on π_{k-1} and $\pi_{\sqrt{n}}$. Note that the event $\{\tau_{k-1,k} > A^2 n\}$ is increasing (while $\{\tau_{k,\sqrt{n}} > A^2 n\}$ is decreasing) w.r.t. the increments of π_k, $(\pi_k(i) - \pi_k(i-1))_{1 \leq i \leq A^2 n} \in \{\pm 1\}^{A^2 n}$, where the product space $\{\pm 1\}^{A^2 n}$ is equipped with the partial order $\prec$ such that $(a_i)_{1 \leq i \leq A^2 n} \prec (b_i)_{1 \leq i \leq A^2 n}$ if $a_i \leq b_i$ for all i. Furthermore, on the event $\{\tau_{k-1,k}, \tau_{k,\sqrt{n}} > A^2 n\}$, π_k is distributed as an independent random walk with i.i.d. increments, whose law on $\{\pm 1\}^{A^2 n}$ satisfies the FKG inequality [42]. This implies that the events $\{\tau_{k-1,k} > A^2 n\}$ and $\{\tau_{k,\sqrt{n}} > A^2 n\}$ are negatively correlated under the law of π_k with i.i.d. increments. Denoting $\mathbb{P}_k$ for probability for π_k, we then have

$$\begin{aligned} &\sum_{k=1}^{\sqrt{n}-1} \mathbb{P}(\tau_{0,k-1} \leq A^2 n, \tau_{k-1,k} > A^2 n, \tau_{k,\sqrt{n}} > A^2 n) \\ &\quad = \sum_{k=1}^{\sqrt{n}-1} \mathbb{E}[1_{E_k} \mathbb{P}_k(\tau_{k-1,k} > A^2 n, \tau_{k,\sqrt{n}} > A^2 n)] \end{aligned}$$

$$\leq \sum_{k=1}^{\sqrt{n}-1} \mathbb{E}[1_{E_k} \mathbb{P}_k(\tau_{k-1,k} > A^2 n) \cdot \mathbb{P}_k(\tau_{k,\sqrt{n}} > A^2 n)]$$

$$\leq \sum_{k=1}^{\sqrt{n}-1} \mathbb{P}(\tau_{k-1,k} > A^2 n) \mathbb{P}(\tau_{k,\sqrt{n}} > A^2 n)$$

$$\leq \sqrt{n} \frac{C}{\sqrt{A^2 n}} \cdot \frac{C\sqrt{n}}{\sqrt{A^2 n}} \leq \frac{C^2}{A^2},$$

where for the second inequality, we used the fact that $\mathbb{P}_k(\tau_{k-1,k} > A^2 n)$ and $\mathbb{P}_k(\tau_{k,\sqrt{n}} > A^2 n)$ are respectively functions of π_{k-1} and $\pi_{\sqrt{n}}$, which are distributed as independent random walks on the event E_k, and in the last line, we used that

$$\mathbb{P}(\tau_{0,1} > A^2 n) \leq \frac{C}{A\sqrt{n}}, \tag{6.60}$$

$$\mathbb{P}(\tau_{k,\sqrt{n}} > A^2 n) \leq \mathbb{P}(\exists k < i \leq \sqrt{n} : \tau_{i-1,i} \geq A^2 n) \leq \frac{C\sqrt{n}}{A\sqrt{n}} = \frac{C}{A}, \tag{6.61}$$

which hold for any random walk on $\mathbb{Z}$ with finite variance [82, Proposition 32.4]. Condition (6.59) then follows.

(U): Recall from Figure 6.1 that the collection of coalescing simple random walk paths X uniquely determines a collection of dual coalescing simple random walk paths, which we denote by $\widehat{X}$. Clearly no path in X can enter any wedge of $\widehat{X}$ from outside. Furthermore, for any $z_1, z_2 \in \mathbb{R}^2$ and any choice of paths $\hat{\pi}_{\varepsilon,1}, \hat{\pi}_{\varepsilon,2} \in S_\varepsilon \widehat{X}$ which converge to dual coalescing Brownian motions starting at z_1, z_2, it is easily seen that the time of coalescence between $\hat{\pi}_{\varepsilon,1}$ and $\hat{\pi}_{\varepsilon,2}$ also converges to that between the limiting Brownian motions. Condition (U) then follows.

We will show how condition (E) is verified for general random walks in the next section. □

6.6.4 Convergence of General Coalescing Random Walks to the Brownian Web

Let X denote the collection of coalescing random walk paths on $\mathbb{Z}^2$ with one walk starting from each site in $\mathbb{Z}^2$, where the increments are i.i.d. with distribution μ with zero mean and finite variance $\sigma^2 := \sum_{x \in \mathbb{Z}} x^2 \mu(x)$, such that the walks are irreducible and aperiodic. Let $S_\varepsilon^\sigma : \mathbb{R}^2 \to \mathbb{R}^2$ be the diffusive scaling map

$$S_\varepsilon^\sigma(x,t) = (\varepsilon \sigma^{-1} x, \varepsilon^2 t), \tag{6.62}$$

and the action of S_ε^σ on paths and sets of paths is defined as S_ε in (6.58).

We have the following convergence result for general coalescing random walks on $\mathbb{Z}$, where paths may cross (see [72, Theorem 1.5] and [16, Theorem 1.2]).

Theorem 6.6.9 *Let X be the collection of coalescing random walk paths on $\mathbb{Z}$ defined as above. If the random walk increment distribution μ has zero mean, variance σ^2, and finite r-th moment $\sum_{x\in\mathbb{Z}} |x|^r \mu(x) < \infty$ for some $r > 3$, then as $\varepsilon \downarrow 0$, $S_\varepsilon^\sigma X$ converges in distribution to a standard Brownian web $\mathcal{W}$.*

Remark 6.6.10 It was pointed out in [72, Remark 4.1] that if $\sum_{x\in\mathbb{Z}} |x|^r \mu(x) = \infty$ for some $r < 3$, then tightness of the collection of coalescing random walks is lost due to the presence of arbitrarily large jumps originating from every space-time window on the diffusive scale.

Proof Sketch. For X consisting of random walk paths that can cross each other, it is not known how to verify condition (B2′) in Theorem 6.6.2. We will instead apply Theorem 6.6.3 and sketch how the conditions therein can be verified. Further details can be found in [72].

To verify condition (T), by translation invariance and reformulation in terms of the set of unscaled random walk paths X, it suffices to show that for any $\eta > 0$,

$$\lim_{\delta\downarrow 0} \delta^{-1} \limsup_{\varepsilon\downarrow 0} \mathbb{P}(X \in A_{2\delta\varepsilon^{-2},20\eta\varepsilon^{-1}}(0,0)) = 0, \tag{6.63}$$

where the event $E := \{X \in A_{2\delta\varepsilon^{-2},20\eta\varepsilon^{-1}}(0,0)\}$ (see Figure 6.14) is the event that X contains some path which intersects the rectangle $R_1 := [-\eta\varepsilon^{-1}, \eta\varepsilon^{-1}] \times [0, \delta\varepsilon^{-2}]$, and at a later time, intersects the left or right boundary of the bigger rectangle $R_2 := [-10\eta\varepsilon^{-1}, 10\eta\varepsilon^{-1}] \times [0, 2\delta\varepsilon^{-2}]$. The event E can occur either due to a random walk which starts outside the rectangle R_1 and has a jump that crosses R_1 horizontally, the probability of which can be easily shown to be negligible; or due to a random walk which starts from some $z \in R_1 \cap \mathbb{Z}^2$ and crosses the left, respectively the right, side of R_2, which events we denote respectively by E_z^- and E_z^+.

We can bound $\mathbb{P}(\cup_{z\in R_1\cap\mathbb{Z}^2} E_z^+)$ as follows. Let $(\pi_i)_{1\le i\le 4}$ be four random walks in X, starting respectively at $2.5\eta\varepsilon^{-1}$, $4.5\eta\varepsilon^{-1}$, $6.5\eta\varepsilon^{-1}$ and $8.5\eta\varepsilon^{-1}$ at time 0. Let B denote the event that each of these walks stays confined in a centered interval of size $\eta\varepsilon^{-1}$ up to time $2\delta\varepsilon^{-2}$ (see Figure 6.15). For $1 \le i \le 4$, let τ_i denote the first time $n \in \mathbb{N}$ when $\pi_z(n) \ge \pi_i(n)$, let τ_5 denote the first time π_z crosses the right side of R_2, and let C_i denote the event that π_z does not meet π_i before time τ_5. We

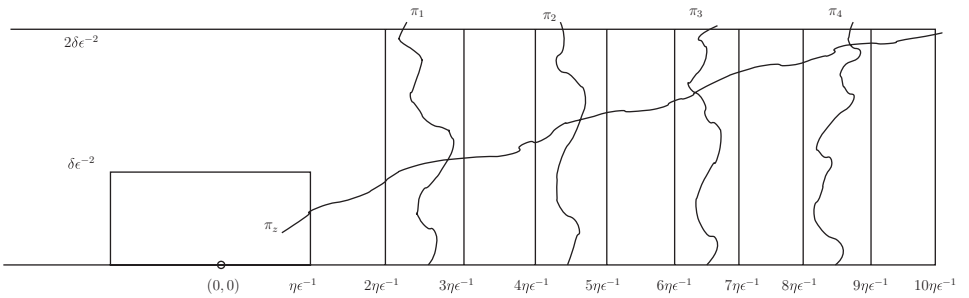

Figure 6.15. The random walks $\pi_1, \pi_2, \pi_3, \pi_4$ start respectively at $2.5\eta\epsilon^{-1}$, $4.5\eta\varepsilon^{-1}$, $6.5\eta\epsilon^{-1}$ and $8.5\eta\epsilon^{-1}$ at time 0 and each stays within an interval of size $\eta\epsilon^{-1}$ up to time $2\delta\epsilon^{-2}$. The walk π_z starts from z inside the rectangle $[-\eta\epsilon^{-1}, \eta\epsilon^{-1}] \times [0, \delta\epsilon^{-2}]$ and crosses the right side of the rectangle $[-10\eta\epsilon^{-1}, 10\eta\epsilon^{-1}] \times [0, 2\delta\epsilon^{-2}]$ before time $2\delta\epsilon^{-2}$ without meeting $(\pi_i)_{1\le i\le 4}$.

can then bound

$$\begin{aligned}
\lim_{\delta\downarrow 0}\delta^{-1}\limsup_{\varepsilon\downarrow 0}\mathbb{P}(\cup_{z\in R_1\cap\mathbb{Z}^2}E_z^+) &= \lim_{\delta\downarrow 0}\delta^{-1}\limsup_{\varepsilon\downarrow 0}\mathbb{P}(B^c) + \lim_{\delta\downarrow 0}\delta^{-1} \\
&\qquad \limsup_{\varepsilon\downarrow 0}\mathbb{P}(B\cap\cup_{z\in R_1\cap\mathbb{Z}^2}E_z^+) \\
&= \lim_{\delta\downarrow 0}\delta^{-1}\limsup_{\varepsilon\downarrow 0}\mathbb{P}(B\cap\cup_{z\in R_1\cap\mathbb{Z}^2}E_z^+) \\
&\le \lim_{\delta\downarrow 0}\delta^{-1}\limsup_{\varepsilon\downarrow 0} 2\eta\delta\varepsilon^{-3}\max_{z\in R_1\cap\mathbb{Z}^2}\mathbb{P}(B\cap E_z^+),
\end{aligned}$$

where in the second line we used Donsker's invariance principle and properties of Brownian motion. When the random walk increments are bounded by some $K < \infty$, it is easy to see that $\mathbb{P}(B \cap E_z^+) \le C\varepsilon^4$ uniformly in $z \in R_1 \cap E_z^+$, so that the limit above equals zero.

Indeed, by successively conditioning on π_z and $(\pi_i)_{1\le i\le 4}$ up to the stopping times τ_4, τ_3, τ_2 and τ_1, we note that π_z comes within distance $2K$ of π_i at time τ_i for each $1 \le i \le 4$, and π_z and π_i must separate by a distance of at least $\eta\varepsilon^{-1}$ before time τ_{i+1} without meeting. By the strong Markov property, these events are conditionally independent, and each event has a probability of order ε by [72, Lemma 2.4]. It then follows that $\mathbb{P}(B \cap E_z^+) \le C\varepsilon^4$. When the random walks have unbounded increments, it is necessary to control the overshoot $\pi_z(\tau_i) - \pi_i(\tau_i)$, so that the probability of π_z overshooting more than one π_i in one jump is negligible when taken union over all starting positions $z \in R_1 \cap \mathbb{Z}^2$. This can be done when the random walk increments have finite 5th moment [72]. To relax to finite r-th moment for some $r > 3$, a multiscale argument is needed to take advantage of the coalescence and reduction of the random walks instead of the crude union bound as above. This was carried out in [16].

Condition (I) can be verified by a similar argument to that for coalescing simple random walks in Theorem 6.6.8. The complication is that random walk paths can cross without coalescing. However, for random walks with finite second moments, at the time when their paths cross, the distance between the two walks is of order one uniformly w.r.t. their starting positions, and hence it takes another time interval of order one for the two walks to coalesce. Therefore the crossing time and the coalescence time between any pair of walks are indistinguishable in the diffusive scaling limit, and forcing coalescence of the random walks at crossing times gives a good approximation, for which Donsker's invariance principle and the Continuous Mapping Theorem can be applied.

Condition (B1′) amounts to showing that

$$\lim_{\eta \downarrow 0} \mathbb{P}(\eta_X(0,n;0,\eta\sqrt{n}) \geq 2) = 0, \tag{6.64}$$

which follows from the same bounds as in (6.60)–(6.61), since on the event $\{\eta_X(0,n;0,\eta\sqrt{n}) \geq 2\}$, we must have $\tau_{i-1,i} \geq n$ for some $1 \leq i \leq \eta\sqrt{n}$, where $\tau_{i-1,i}$ is the first meeting time between the two walks in X starting respectively at $i-1$ and i at time 0.

The key to verifying condition (E) is the density bound (6.56), which formulated in terms of the set of unscaled random walks X becomes

$$\mathbb{P}(0 \in \xi^{\mathbb{Z}}(n)) \leq \frac{C}{\sqrt{n}} \qquad \text{uniformly in } n \in \mathbb{N}, \tag{6.65}$$

where $\xi^A(n) := \{\pi(n) : \pi \in X, \pi(0) \in A\}$ for $A \subset \mathbb{Z}$. By translation invariance, for any $L \in \mathbb{N}$,

$$\begin{aligned}
\mathbb{P}(0 \in \xi^{\mathbb{Z}}(n)) &= \frac{1}{L}\mathbb{E}[\,|\xi^{\mathbb{Z}}(n) \cap [0,L)|\,] \\
&\leq \frac{1}{L}\sum_{k\in\mathbb{Z}} \mathbb{E}[\,|\xi^{\mathbb{Z}\cap[kL,(k+1)L)}(n) \cap [0,L)|\,] \\
&= \frac{1}{L}\sum_{k\in\mathbb{Z}} \mathbb{E}[\,|\xi^{\mathbb{Z}\cap[0,L)}(n) \cap [kL,(k+1)L)|\,] \\
&= \frac{1}{L}\mathbb{E}[\,|\xi^{\mathbb{Z}\cap[0,L)}(n)|\,] \leq \frac{1}{L}\mathbb{E}\left[1 + \sum_{i=1}^{L-1} 1_{\{\tau_{i-1,i}>n\}}\right] \leq \frac{1}{L} + \frac{C}{\sqrt{n}},
\end{aligned}$$

where in the last inequality, we applied (6.60). Letting $L \to \infty$ then gives (6.65). The rest of the proof of condition (E) then follows the line of argument sketched after Theorem 6.6.3. □

6.6.5 Convergence to the Brownian Net

Convergence to the Brownian net so far has only been established by Sun and Swart [86] for branching-coalescing simple random walk paths with asymptotically vanishing branching probability (see Figure 6.4), and recently by Etheridge, Freeman and Straulino [34] for the genealogies of a spatial Lambda-Fleming-Viot process. We identify below the key conditions and formulate them as convergence criteria for the Brownian net, which can be applied to random sets of paths with certain non-crossing properties. Finding effective and verifiable convergence criteria for random sets of paths which do not satisfy the non-crossing condition (C) below remains a major challenge.

Given a sequence of $(\mathcal{H}, \mathcal{B}_{\mathcal{H}})$-valued random variables $(X_n)_{n\in\mathbb{N}}$, we first impose a non-crossing condition.

(C) There exist subsets of *non-crossing* paths $W_n^{\mathrm{l}}, W_n^{\mathrm{r}} \subset X_n$, such that no path $\pi \in X_n$ crosses any $l \in W_n^{\mathrm{l}}$ from right to left, i.e., $\pi(s) > l(s)$ and $\pi(t) < l(t)$ for some $s < t$, and no path $\pi \in X_n$ crosses any $r \in W_n^{\mathrm{r}}$ from left to right.

The second condition is an analogue of condition (I), which ensures that $(W_n^{\mathrm{l}}, W_n^{\mathrm{r}})_{n\in\mathbb{N}}$ is a tight family by Proposition 6.6.4, and any subsequential weak limit contains a copy of the left-right Brownian web $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ defined as in Theorem 6.3.2.

($\mathrm{I}_{\mathcal{N}}$) There exist $l_{n,z} \in W_n^{\mathrm{l}}$ and $r_{n,z} \in W_n^{\mathrm{r}}$ for each $z \in \mathbb{R}^2$, such that for any deterministic $z_1, \ldots, z_k \in \mathbb{R}^2$, $(l_{n,z_1}, \ldots, l_{n,z_k}, r_{n,z_1}, \ldots, r_{n,z_k})$ converge in distribution to a collection of left-right coalescing Brownian motions starting at $(z_i)_{1\le i\le k}$, as in Theorem 6.3.2.

The tightness of $(W_n^{\mathrm{l}}, W_n^{\mathrm{r}})_{n\in\mathbb{N}}$ and condition (C) imply that $(X_n)_{n\in\mathbb{N}}$ is also a tight family, since (C) implies that almost surely, the modulus of continuity of paths in X_n (cf. (6.51)) can be bounded by the modulus of continuity of paths in $W_n^{\mathrm{l}} \cup W_n^{\mathrm{r}}$, whose starting points become dense in $\mathbb{R}^2$ as $n \to \infty$ by condition ($\mathrm{I}_{\mathcal{N}}$).

The next condition is

(H) A.s., X_n contains all paths obtained by hopping among paths in $W_n^{\mathrm{l}} \cup W_n^{\mathrm{r}}$ at crossing times, defined as in the hopping construction of the Brownian net in Theorem 6.3.4.

Condition (H) ensures that any subsequential weak limit of $(X_n)_{n\in\mathbb{N}}$ contains not only a copy of the left-right Brownian web $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, but also a copy of the Brownian net $\mathcal{N}$ constructed by hopping among paths in $\mathcal{W}^{\mathrm{l}} \cup \mathcal{W}^{\mathrm{r}}$ at crossing times.

Lastly, we formulate the analogue of condition (U) in Theorem 6.6.6, which gives an upper bound on any subsequential weak limit of $(X_n)_{n\in\mathbb{N}}$ via the dual wedge characterization of the Brownian net given in Theorem 6.3.9.

($\mathrm{U}_{\mathcal{N}}$) There exist $\widehat{W}_n^{\mathrm{l}}, \widehat{W}_n^{\mathrm{r}} \in \widehat{\mathcal{H}}$, whose starting points are dense in $\mathbb{R}^2$ as $n \to \infty$, such that a.s. paths in W_n^{l} and $\widehat{W}_n^{\mathrm{l}}$ (resp. paths in W_n^{r} and $\widehat{W}_n^{\mathrm{r}}$) do not cross, and for any subsequential weak limit $(\mathcal{X}, W^{\mathrm{l}}, W^{\mathrm{r}}, \widehat{W}^{\mathrm{l}}, \widehat{W}^{\mathrm{r}})$ of $(X_n, W_n^{\mathrm{l}}, W_n^{\mathrm{r}}, \widehat{W}_n^{\mathrm{l}}, \widehat{W}_n^{\mathrm{r}})$ and any deterministic countable dense $\mathcal{D} \subset \mathbb{R}^2$, a.s. paths in $\mathcal{X}$ do not enter any wedge of $(\widehat{W}^{\mathrm{l}}(\mathcal{D}), \widehat{W}^{\mathrm{r}}(\mathcal{D}))$ from outside.

Condition (H) implies that $(W^{\mathrm{l}}, W^{\mathrm{r}})$ contains a copy of the left-right Brownian web $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, and the non-crossing property implies that $(\widehat{W}^{\mathrm{l}}(\mathcal{D}), \widehat{W}^{\mathrm{r}}(\mathcal{D}))$ coincides with $(\widehat{\mathcal{W}}^{\mathrm{l}}(\mathcal{D}), \widehat{\mathcal{W}}^{\mathrm{r}}(\mathcal{D}))$ for the dual left-right Brownian web. By Theorem 6.3.9, the assumption that no path in $\mathcal{X}$ enters any wedge of $(\widehat{\mathcal{W}}^{\mathrm{l}}(\mathcal{D}), \widehat{\mathcal{W}}^{\mathrm{r}}(\mathcal{D}))$ from outside then implies that $\mathcal{X}$ is contained in the Brownian net constructed from $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$, which is the desired upper bound on $\mathcal{X}$.

We thus have the following convergence result.

Theorem 6.6.11 (Convergence criteria for the Brownian net) *Let $(X_n)_{n\in\mathbb{N}}$ be a sequence of $(\mathcal{H}, \mathcal{B}_{\mathcal{H}})$-valued random variables which satisfy conditions* (C), ($\mathrm{I}_{\mathcal{N}}$), (H) *and* ($\mathrm{U}_{\mathcal{N}}$) *above. Then X_n converges in distribution to the standard Brownian net $\mathcal{N}$.*

Remark 6.6.12 To verify condition ($\mathrm{U}_{\mathcal{N}}$), it suffices to show that paths in X_n do not enter wedges of $(\widehat{W}_n^{\mathrm{l}}, \widehat{W}_n^{\mathrm{r}})$ from outside, and when a sequence of pairs of paths in $(\widehat{W}_n^{\mathrm{l}}, \widehat{W}_n^{\mathrm{r}})$ converge to a pair of dual left-right coalescing Brownian motions, the associated first meeting times between the pair also converge, so that the associated wedges converge.

The list of conditions in Theorem 6.6.11 can be verified for the collection of branching-coalescing simple random walks on $\mathbb{Z}$, which answers **Q.1** at the start of Section 6.3.

Theorem 6.6.13 *For $\varepsilon \in (0,1)$, let X_ε be the collection of branching-coalescing simple random walk paths on $\mathbb{Z}^2_{\mathrm{even}}$ with branching probability ε, with walks starting from every site of $\mathbb{Z}^2_{\mathrm{even}}$ (see Figure 6.4). Let S_ε be defined as in* (6.58)*. Then $S_\varepsilon X_\varepsilon$ converges in distribution to the standard Brownian net $\mathcal{N}$ as $\varepsilon \downarrow 0$.*

Proof sketch. Conditions (C) and (H) hold trivially for the branching-coalescing simple random walks. Condition ($\mathrm{I}_{\mathcal{N}}$) was verified in [86, Section 5] using the fact that a single pair of leftmost and rightmost random walk paths solve a discrete analogue of the SDE (6.11) for a pair of left-right coalescing Brownian motions, and furthermore, the time when the pair of discrete paths meet converges to the

continuum analogue. By Remark 6.6.12, it is then easily seen that condition $(U_{\mathcal{N}})$ also holds, and hence Theorem 6.6.11 can be applied. □

6.7 Survey on Related Results

In this section, we survey interesting results connected to the Brownian web and net that have not been discussed so far, including alternative topologies, models whose scaling limits are connected to the Brownian web and net (including population genetic models, true self-avoiding walks, planar aggregation, drainage networks, supercritical oriented percolation), and the relation between the Brownian web and net, critical planar percolation and Tsirelson's theory of noise [95, 96].

6.7.1 Alternative Topologies

We review here several alternative choices of state spaces and topologies for the Brownian web, and compare them with the paths topology of Fontes et al. [40], as introduced in Section 6.2.1. In particular, we will review the *weak flow topology* of Norris and Turner [77], the *tube topology* of Berestycki, Garban and Sen [14] and the *marked metric measure spaces* used by Greven, Sun and Winter [52].

Another natural extension of the paths topology is to consider the space of compact sets of càdlàg paths equipped with the Hausdorff topology, where the space of càdlàg paths is equipped with the Skorohod metric after compactification of space-time as done in Figure 6.2. Such an extension has been carried out by Etheridge, Freeman and Straulino [34] in the study of scaling limits of spatial Lambda-Fleming-Viot processes.

6.7.1.1 Weak Flow Topology

In [77], Norris and Turner formulated a topology for stochastic flows, which includes the Arratia flow generated by coalescing Brownian motions.

Recall that a flow on a space E is a two-parameter family of functions $(\phi_{s,t})_{s\le t}$ from E to E, which satisfies the flow condition $\phi_{t,u} \circ \phi_{s,t} = \phi_{s,u}$ for any $s \le t \le u$. It is easily seen that the Brownian web almost surely defines a flow on $\mathbb{R}$. Indeed, if $(x,s) \in \mathbb{R}^2$ is a $(1,2)$ point in the Brownian web with one incoming and two outgoing paths, then $(\phi_{s,t}(x))_{t\ge s}$ should be defined as the continuation of the incoming Brownian path so that the flow condition is satisfied, and otherwise, $(\phi_{s,t}(x))_{t\ge s}$ can be defined to be any of the Brownian paths starting at (x,s). Note that there is no unique definition of this flow. Furthermore, it is not clear how to define a suitable topology on the space of flows in order to prove convergence of flows, as well as weak convergence of stochastic flows.

Norris and Turner addressed these issues by introducing the notion of a *weak flow*, which on $\mathbb{R}$, is a family $(\phi_{s,t})_{s\le t}$ with the properties that

(i) For all $s \le t$, $\phi_{s,t} \in \mathcal{D}$, the space of non-decreasing functions, so that the flow lines $(\phi_{s,t}(x))_{t\ge s}$, for $(x,s) \in \mathbb{R}^2$, do not cross.
(ii) If $\phi^+_{s,t}$ (resp. $\phi^-_{s,t}$) is the right (resp. left)-continuous version of $\phi_{s,t}$, then

$$\phi^-_{t,u} \circ \phi^-_{s,t} \le \phi^-_{s,u} \le \phi^+_{s,u} \le \phi^+_{t,u} \circ \phi^+_{s,t} \qquad \text{for all } s < t < u.$$

One can endow the space $\mathcal{D}$ with a metric $d_{\mathcal{D}}$ such that if $f, g \in \mathcal{D}$, and $f^\times, g^\times$ denote the graphs of f and g rotated clockwise by $\pi/4$ around the origin, then

$$d_{\mathcal{D}}(f,g) = \sum_{n=1}^{\infty} 2^{-n}(1 \wedge \sup_{|x|<n} |f^\times(x) - g^\times(x)|). \tag{6.66}$$

Note that when f and g are distribution functions, $|f^\times - g^\times|_\infty$ is just the Lévy metric on probability measures. The space of *continuous weak flows* $C^\circ(\mathbb{R}, \mathcal{D})$ then consists of continuous maps $\phi : \{(s,t) \in \mathbb{R}^2 : s \le t\} \to \mathcal{D}$ with $\phi_{s,s} = \mathrm{id}$ for all $s \in \mathbb{R}$, and it is endowed with the topology of uniform convergence on bounded subsets of $\{(s,t) \in \mathbb{R}^2 : s < t\}$.

To allow for non-continuous weak flows, Norris and Turner also introduced the space of *cádlág weak flows* $D^\circ(\mathbb{R}, \mathcal{D})$, which consists of maps ϕ from the space of bounded intervals to $\mathcal{D}$, such that $\phi_\emptyset = \mathrm{id}$, $\phi_{(s,t)} \to \phi_\emptyset$ as $t \downarrow s$ and $\phi_{(t,u)} \to \phi_\emptyset$ as $t \uparrow u$, while $\phi_{\{t\}}$ captures the jump discontinuity of the flow at time t. One can then equip $D^\circ(\mathbb{R}, \mathcal{D})$ with a Skorohod-type topology.

It was shown in [77] that $C^\circ(\mathbb{R}, \mathcal{D})$ and $D^\circ(\mathbb{R}, \mathcal{D})$ are Polish spaces, the Arratia flow is a continuous weak flow, and a family of Poisson local disturbance flows converge in distribution to the Arratia flow. Actually [77] considered only weak flows on the circle $\mathbb{S}$. To have non-decreasing maps, a map $f : \mathbb{S} \to \mathbb{S}$ is lifted to a map from $\mathbb{R}$ to $\mathbb{R}$ by identifying $\mathbb{S}$ with $[0, 2\pi)$ and extending f outside $[0, 2\pi)$ by setting $f(x + 2\pi) = f(x) + 2\pi$. The extension to weak flows on $\mathbb{R}$ has subsequently been carried out by Ellis in [32].

Compared with the paths topology introduced in Section 6.2.1, the advantage of the weak flow topology is that it is more natural for studying stochastic flows, and it allows for discontinuity in the flow lines (paths). The limitation is that it is restricted to flows with non-crossing paths. Proving weak convergence to the Arratia flow is very simple: it suffices to verify condition (I) in Section 6.6.1, which ensures that every weak limit point contains the coalescing Brownian motions. There is no need to upper bound the limiting set of paths, unlike the paths topology of Section 6.2.1, because the weak flow topology effectively discards all paths starting from the same space-time point other than the leftmost and rightmost paths, because of the metric $d_{\mathcal{D}}$ on $\mathcal{D}$.

6.7.1.2 Tube Topology

Just as the paths topology of Fontes et al. [40] was inspired by a similar topology introduced by Aizenman et al. [3, 4] for two-dimensional percolation, Berestycki, Garban and Sen took inspiration from another topology for two-dimensional percolation, the *quad topology* of Schramm and Smirnov [87], and introduced in [14] the *tube topology* for sets of continuous paths in $\mathbb{R}^d$. The basic idea is that the configuration of a set of paths is captured by the set of space-time tubes these paths cross.

Let $T^* := ([T^*], \partial_0 T^*, \partial_1 T^*) := ([0,1]^d \times [0,1], [0,1]^d \times \{0\}, [0,1]^d \times \{1\})$ be the unit tube in $\mathbb{R}^{d+1}$, with $\partial_0 T^*$ and $\partial_1 T^*$ being the lower and upper faces of the tube. A tube $T = ([T], \partial_0 T, \partial_1 T)$ in $\mathbb{R}^{d+1}$ is then defined to be the image of T^* under a homeomorphism $\phi : \mathbb{R}^{d+1} \to \mathbb{R}^{d+1}$ with the property that $\partial_0 T \subset \mathbb{R}^d \times \{t_0\}$ and $\partial_1 T \subset \mathbb{R}^d \times \{t_1\}$ for some $t_0 < t_1$, and $[T] \subset \mathbb{R}^d \times [t_0, t_1]$. We call t_0 the bottom time of T, and t_1 the top time. The space of all tubes, denoted by $\mathcal{T}$, can then be equipped with the Hausdorff metric $d_{\mathcal{T}}$, with

$$d_{\mathcal{T}}(T_1, T_2) := d_{\text{Haus}}([T_1], [T_2]) + d_{\text{Haus}}(\partial_0 T_1, \partial_0 T_2) + d_{\text{Haus}}(\partial_1 T_1, \partial_1 T_2), \quad (6.67)$$

where d_{Haus} is the Hausdorff distance between subsets of $\mathbb{R}^{d+1}$. The metric space $(\mathcal{T}, d_{\mathcal{T}})$ is then separable.

Given a continuous path $\pi : [t, \infty) \to \mathbb{R}^d$ starting at time t, it is said to cross a tube T with lower face $\partial_0 T \subset \mathbb{R}^d \times \{t_0\}$ and upper face $\partial_1 T \subset \mathbb{R}^d \times \{t_1\}$ for some $t_0 < t_1$, if $t \le t_0$, $(\pi(t_0), t_0) \in \partial_0 T$, $(\pi(t_1), t_1) \in \partial_1 T$), and $(\pi(s), s) \in [T]$ for all $s \in [t_0, t_1]$. Given a set K of continuous paths in $\mathbb{R}^d$, one can then identify the set of all tubes $\text{Cr}(K) \subset \mathcal{T}$ which are crossed by some path in K. Furthermore, if K is a compact set of continuous paths w.r.t. a metric on path space defined in the same way as in Section 6.2.1, then $\text{Cr}(K)$ is in fact a closed subset of $\mathcal{T}$. Therefore a random compact set of paths can be identified with a random closed subset of $\mathcal{T}$. The state space of closed subsets of the metric space $\mathcal{T}$ can then be equipped with the Fell topology, which makes it compact (see e.g., [66, Appendix B]).

One can actually narrow down the state space further. Observe that we can define a partial order $\le$ on $\mathcal{T}$, where we denote $T_1 \le T_2$ if whenever T_2 is crossed by some path π, T_1 is also crossed by π. We denote $T_1 < T_2$ if there are open neighborhoods $U_i \subset \mathcal{T}$ around T_i, $i = 1, 2$, such that $T_1' \le T_2'$ for all $T_1' \in U_1$ and $T_2' \in U_2$. A set of tubes $S \subset \mathcal{T}$ is then called *hereditary* if $T \in S$ implies that $T' \in S$ for all $T' < T$. Note that the set of tubes $\text{Cr}(K)$ induced by a set of paths K is always hereditary. Therefore the state space for random compact sets of paths can be taken to be the space of *closed hereditary subsets* of $\mathcal{T}$, denoted by $\mathcal{H}$, equipped with the Fell topology. It can be shown that $\mathcal{H}$ is closed under the Fell topology, and hence compact. This gives the tube topology defined in [14, Section 2], and the Brownian web can be realized as a random variable taking values in $\mathcal{H}$.

The main advantages of the tube topology are: (i) tightness for a family of $\mathcal{H}$-valued random variables comes for free because the state space $\mathcal{H}$ is compact. For instance, under the tube topology, Theorem 6.6.9 on the convergence of general coalescing random walks to the Brownian web can be established under the optimal finite second moment assumption, since the higher moment assumption is only used to establish tightness under the path topology; (ii) the tube topology, being actually a weaker topology than the path topology of Fontes et al. [40] introduced in Section 6.2.1, makes it much easier to construct coalescing flows which do not satisfy the non-crossing property of the Arratia flow. In particular, the coalescing Brownian flow on the Sierpinski gasket was constructed in [14] using the tube topology, and an invariance principle was established.

To characterize a probability measure on $\mathcal{H}$, it turns out to be sufficient to determine the probability of the joint crossing of any finite collection of tubes chosen from a deterministic countable dense subset of $\mathcal{T}$. To prove the weak convergence of a sequence of probability measures on $\mathcal{H}$, it is sufficient to find a large enough set of tubes $\hat{\mathcal{T}} \subset \mathcal{T}$ such that the the probability of the joint crossing of any finite subset of $\hat{\mathcal{T}}$ converges to a limit (see [14, Proposition 2.12]). The strategy for verifying this convergence criterion for coalescing flows is similar to the proof of condition (E) in Theorem 6.6.3. First one approximates a given tube T with bottom time t_0 by tubes T_δ with bottom time $t_0 + \delta$ for $\delta > 0$, then one applies a coming down from infinity result to show that among paths which intersect the lower face of T, only finitely many remain in the lower face of T_δ at time $t_0 + \delta$, and lastly one uses condition (I) to control the joint distribution of the remaining finite collection of coalescing paths.

The tube topology is weaker than the paths topology introduced in Section 6.2.1, because the mapping from the space of compact sets of continuous paths (with the paths topology) to the space of closed hereditary sets of tubes $\mathcal{H}$ (with the Fell topology) is continuous as shown in [14, Lemma A.2]. In particular, given a set of paths K, restricting a path $\pi \in K$ to a later starting time and adding it to K has no effect under the tube topology. In a sense, the tube topology is also insensitive toward the behavior of the paths near their starting times. In particular, a sequence of paths π_n, which start at time 0 and converge to a path π uniformly on $[\delta, \infty)$ for any $\delta > 0$, will converge to π under the tube topology. But π_n may have wild oscillations in the interval $[0, 1/n)$ which prevent it from converging in the paths topology.

6.7.1.3 Marked Metric Measure Spaces

In [52], Greven, Sun and Winter treat the (dual) Brownian web as a stochastic process taking values in the space of spatially-marked metric measure spaces. The notion of a V-marked metric probability measure space was introduced

by Depperschmidt, Greven and Pfaelhuber in [30], which is simply a complete separable metric space (X,d), together with a Borel probability measure μ on the product space $X \times V$ for some complete separable metric space of marks V. The probability measure μ can be regarded as a sampling measure, and each V-marked metric measure (V-mmm) space can be uniquely determined (up to isomorphism) by the joint distribution of $(v_i)_{i\in\mathbb{N}}$ and the distance matrix $(d(x_i,x_j))_{i,j\in\mathbb{N}}$, where $(x_i,v_i)_{i\in\mathbb{N}}$ are i.i.d. samples drawn from $X \times V$ with common distribution μ. A sequence of V-mmm spaces is then said to converge if the associated random vector of marks $(v_i)_{i\in\mathbb{N}}$ and the random distance matrix $(d(x_i,x_j))_{i,j\in\mathbb{N}}$ converge in finite-dimensional distribution. By truncations in the mark space V, one can also extend the notion of V-mmm spaces to the case where μ is only required to be finite on bounded sets when projected from $X \times V$ to V, which was done in [52].

To see how the Brownian web fits in the framework of V-mmm spaces, note that the coalescing Brownian motions in the (dual) Brownian web can be regarded as the space-time genealogies of a family of individuals. More precisely, if $\widehat{\mathcal{W}}$ is the dual Brownian web introduced in Section 6.2.3, then for each $z = (x,t) \in \mathbb{R}^2$, interpreted as an individual at position x at time t, $\hat{\pi}_z \in \widehat{\mathcal{W}}$ determines its spatial genealogy, with $\hat{\pi}_z(s)$ (for $s \le t$) being the spatial location of its ancestor at time s. Coalescence of two paths $\hat{\pi}_{z_1}$ and $\hat{\pi}_{z_2}$ then signifies the merging of the two genealogy lines. If we consider all individuals indexed by $\mathbb{R}$ at a given time t, then we can measure the genealogical distance between individuals, i.e.,

$$d((x,t),(y,t)) := 2(t - \hat{\tau}_{(x,t),(y,t)}), \tag{6.68}$$

where $\hat{\tau}_{(x,t),(y,t)}$ is the time of coalescence between $\hat{\pi}_{(x,t)}$ and $\hat{\pi}_{(y,t)} \in \widehat{\mathcal{W}}$. It is easily seen that d is in fact an ultra-metric, i.e.,

$$d((x_1,t),(x_2,t)) \le d((x_1,t),(x_3,t)) \vee d((x_2,t),(x_3,t)) \qquad \text{for all } x_1,x_2,x_3 \in \mathbb{R}.$$

The collection of individuals indexed by $\mathbb{R}$ at time t then forms a metric space with metric d. The mark of an individual is just its spatial index, so that we can identify the metric space with the mark space $V := \mathbb{R}$, and a natural sampling measure is the Lebesgue measure on $\mathbb{R}$. There is one complication, namely that a.s. there exists $x \in \mathbb{R}$ such that multiple paths in $\widehat{\mathcal{W}}$ start from (x,t). The correct interpretation is that each such path encodes the genealogy of a distinct individual, which happens to occupy the same location. Therefore the space of population should be enriched from $\mathbb{R}$ to take into account these individuals. Such an enrichment does not affect the Lebesgue sampling measure, since it follows from Theorem 6.2.11 that a.s. there are only countably many points of multiplicity in $\widehat{\mathcal{W}}$ at a given time t. If we let (X_t,d_t,μ_t) denote the $\mathbb{R}$-mmm space induced by $\widehat{\mathcal{W}}$ on $\mathbb{R}$ at time t, then $\widehat{\mathcal{W}}$ a.s.

determines the stochastic process $(X_t, d_t, \mu_t)_{t\in\mathbb{R}}$, and conversely, it can be seen that $(X_t, d_t, \mu_t)_{t\in\mathbb{R}}$ a.s. determines $\widehat{\mathcal{W}}$, and hence the two can be identified.

We can also identify $\widehat{\mathcal{W}}$ with a single $\mathbb{R}^2$-mmm space instead of the $\mathbb{R}$-mmm space-valued process $(X_t, d_t, \mu_t)_{t\in\mathbb{R}}$. Namely, we can consider all points in $\mathbb{R}^2$ and extend the genealogical distance in (6.68) between individuals at the same time to individuals at different times:

$$d((x_1,t_1),(x_2,t_2)) := t_1 + t_2 - 2\hat{\tau}_{(x_1,t_1),(x_2,t_2)}, \tag{6.69}$$

where $\hat{\tau}_{(x_1,t_1),(x_2,t_2)}$ is the time of coalescence between $\hat{\pi}_{(x_1,t_1)}$ and $\hat{\pi}_{(x_2,t_2)} \in \widehat{\mathcal{W}}$. We then obtain a metric space whose elements can be identified with the mark space $V := \mathbb{R}^2$ (with suitable enrichment of $\mathbb{R}^2$ to take into account points of multiplicity in $\widehat{\mathcal{W}}$), on which we equip as a sampling measure the Lebesgue measure on $\mathbb{R}^2$. We remark that one can also enlarge the mark space and let the whole genealogy line $\hat{\pi}_z \in \widehat{\mathcal{W}}$ be the mark for a point $z \in \mathbb{R}^2$.

For models arising from population genetics, V-mmm space is a natural space for the spatial genealogies, and the dual Brownian web determines the genealogies of the so-called *continuum-sites stepping-stone model*. Because V-mmm spaces are characterized via sampling, proving convergence essentially reduces to condition (I) formulated in Section 6.6.1. However, proving tightness for V-mmm space-valued random variables is a nontrivial task and can be quite involved (see e.g., [52]). Also the Brownian net cannot be characterized using V-mmm spaces, since there is no natural analogue of the metric d in (6.69) when paths can branch.

6.7.2 Other Models which Converge to the Brownian Web and Net

We review here various models which have scaling limits that are connected to the Brownian web and net. These include population genetic models such as the voter model and the spatial Fleming-Viot processes, stochastic Potts models, the true self-avoiding walks, planar aggregation models, drainage networks, and supercritical oriented percolation in dimension $1 + 1$.

6.7.2.1 Voter Model and Spatial Fleming-Viot Processes

Voter model and spatial Fleming-Viot processes are prototypical population genetic models where the spatial genealogies of the population are coalescing random walks, which converge to the Brownian web under diffusive scaling.

Arratia [1, 2] first conceived the Brownian web in studying the scaling limit of the voter model on $\mathbb{Z}$, which is an interacting particle system $(\eta_t)_{t\geq 0}$ with $\eta_t \in \Omega := \{0,1\}^{\mathbb{Z}}$, modeling the opinions of a collection of individuals indexed by $\mathbb{Z}$. Independently for each $x \in \mathbb{Z}$, at exponential rate 1, a resampling event occurs where the voter at x picks one of its two neighbors with equal probability and

changes its opinion $\eta_t(x)$ to that of the chosen neighbor. The resampling events can be represented by a Poisson point process of arrows along the edges over time, with each arrow from x to y at time t signifying the voter at x changing its opinion at time t to that of y (see Figure 6.16(a)). This is known as Harris' graphical construction. To identify $\eta_t(x)$, one just needs to trace the genealogy of where the opinion $\eta_t(x)$ comes from backward in time, i.e., follow the arrows backward in time, until an ancestor y is reached at time 0 so that $\eta_t(x) = \eta_0(y)$. The genealogy line $\hat{\pi}_{(x,t)}$ for $\eta_t(x)$ is then a continuous time random walk running backward in time. Furthermore, for multiple space-time points $((x_i, t_i))_{1\le i\le k}$, their joint genealogy lines $(\hat{\pi}_{(x_i,t_i)})_{1\le i\le k}$ are a collection of coalescing random walks. This is known as the *duality* between the voter model and coalescing random walks (see [54, 61] for more details). Taking the diffusive scaling limit of the joint genealogies of all voter opinions at all possible times then leads to what we now call the (dual) Brownian web.

The spatial Fleming-Viot process $(\xi_t)_{t\ge 0}$ is a measure-valued process which extends the voter model by allowing a continuum of individuals at each site $x \in \mathbb{Z}$, represented by a probability measure on the type space $[0, 1]$. Individuals in the population migrate on $\mathbb{Z}$ as independent random walks, and resampling (one individual changes its type to that of another) takes place between every pair of individuals at the same site with exponential rate 1. This model has recently been studied in [52], and the joint spatial genealogies are a collection of coalescing random walks running backward in time, where each pair of walks at the same site coalesce with exponential rate 1. It is not difficult to see that the diffusive scaling limit of these genealogies also gives the (dual) Brownian web.

It is natural to ask whether the (dual) Brownian web also determines the genealogies of a population genetic model on $\mathbb{R}$. The answer is affirmative, and such a continuum model has been studied before and is known as the *continuum-sites stepping-stone model* (CSSM) (see e.g., Donnelly et al [29]). The Brownian web effectively gives a Harris' graphical construction of the CSSM.

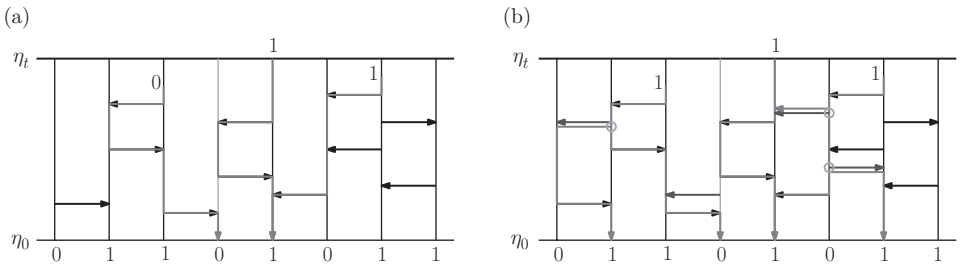

Figure 6.16. Harris' graphical construction of: (a) the voter model; (b) the biased voter model.

Convergence of the voter model as a measure-valued process on $\mathbb{R} \times \{0,1\}$ to the CSSM has been established in [6], and convergence of the genealogies of the spatial Fleming-Viot process to that of the CSSM as stochastic processes taking values in the space of $\mathbb{R}$-mmm spaces (see Section 6.7.1.3) has been established in [52]. In both results, the use of the sampling measure allows one to establish convergence under the optimal finite second moment assumption on the increments of the underlying coalescing random walks, in contrast to Theorem 6.6.9.

We remark that coalescing random walks on $\mathbb{Z}$ also model the evolution of boundaries between domains of different spins in the 0-temperature limit of the stochastic Potts model, and the Brownian web has been used to study aging in this model [37].

6.7.2.2 Biased Voter Model

We have seen that the voter model is dual to coalescing random walks. It turns out that branching-coalescing random walks is dual to the *biased voter-model*, also known as Williams-Bjerknes model [104, 83], which modifies the voter model by adding a selective bias so that type 1 is favoured over type 0. More precisely, each voter independently undergoes a second type of resampling event with rate ε, where the voter at x changes its type to that of the chosen neighbor y only if y is of type 1. The Harris' graphical construction can then be modified by adding a second independent Poisson point process of (selection) arrows along the edges over time, where a selection arrow is used only if it points to an individual of type 1 at that time (see Figure 6.16(b)). To determine $\eta_t(x)$, we then trace its genealogy backward in time by following the resampling arrows, and when a selection arrow is encountered, we follow both potential genealogies (by either ignoring or following the selection arrow). The potential genealogies then form a collection of branching-coalescing random walks with branching rate ε. It is not difficult to see that $\eta_t(x) = 1$ if and only if $\eta_0(y) = 1$ for at least one ancestor y that can be reached by the potential genealogies. This establishes the duality between the biased voter model and branching-coalescing random walks. If we consider a sequence of biased voter models with selection rate $\varepsilon \downarrow 0$, while rescaling space-time by $(\varepsilon, \varepsilon^2)$, then their genealogies converge to the Brownian net.

We remark that branching-coalescing random walks on $\mathbb{Z}$ also model evolving boundaries between domains of different spins in the low temperature limit of the stochastic Potts model, where a new domain of spins can nucleate at the boundary of two existing domains [68].

6.7.2.3 True Self-avoiding Walks and True Self-repelling Motion

Arratia first conceived what we now call the Brownian web in his unfinished manuscript [2], and the subject lay dormant until Tóth and Werner [99] discovered

a surprising connection between the Brownian web and the *true self-avoiding walk* on $\mathbb{Z}$ with bond repulsion [94]. A special case of the true self-avoiding walk is defined as follows. Let $l_0(\cdot)$ be an integer-valued function defined on the edges of $\mathbb{Z}$, which can be regarded as the initial condition for the edge local time of the walk, i.e., how many times each edge has been traversed. Given the walk's position $X_n = x$ and the edge local time $l_n(\cdot)$ at time n, if $l_n(\{x-1,x\}) < l_n(\{x,x+1\})$, then with probability 1, $X_{n+1} = x-1$; if $l_n(\{x,x+1\}) < l_n(\{x-1,x\})$, then with probability 1, $X_{n+1} = x+1$; and if $l_n(\{x-1,x\}) = l_n(\{x,x+1\})$, then $X_{n+1} = x \pm 1$ with probability $1/2$ each. The edge local time $l_n(\cdot)$ is then updated to $l_{n+1}(\cdot)$ by adding 1 to the local time at the newly traversed edge. Such a walk is called a true self-avoiding walk because it is repelled from the more visited regions, and the laws of $(X_i)_{1\le i\le n}$ are consistent as n varies, in contrast to the *self-avoiding walk*.

Interestingly, if we plot the evolution of the position of the walk together with its edge local time, then there is an almost sure coupling with a collection of forward/backward coalescing random walks. Figure 6.17 illustrates such a coupling, where the initial edge local time is given by $l_0(\{2x,2x+1\}) = 1$ for all $x \in \mathbb{N}_0$, $l_0(\{2x-1,2x\}) = 0$ for all $x \in \mathbb{N}$, and $l_0(\{-x-1,-x\}) = l_0(\{x,x+1\})$ for all $x \in \mathbb{Z}$. Such an l_0 corresponds to a special boundary condition for the coalescing random walks along the x-axis, as shown in Figure 6.17. The lattice-filling curve between the forward and backward coalescing random walks encodes the evolution of the position of the walk (the horizontal coordinate) and the edge local time (the vertical coordinate), and the area filled in by the curve is just the time the walk has spent. The collection of forward/backward coalescing random walks converge to the forward/backward Brownian web, albeit with the boundary condition that the path in the forward, respectively backward, web starting at $(0,0)$ is the constant

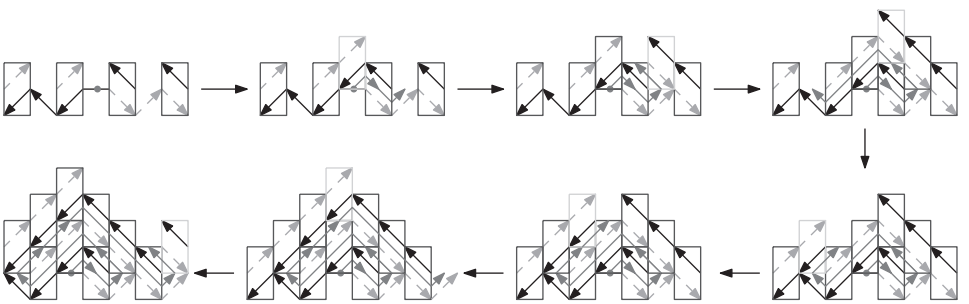

Figure 6.17. Coupling between a true self-avoiding walk (with a special initial edge local time) and a collection of forward/backward coalescing random walks (with a special boundary condition). The horizontal coordinate of the lattice-filling curve is the walk's position and the vertical coordinate is its edge local time.

path, and time for the webs now runs in the horizontal direction in Figure 6.17. It can then be shown (see [71]) that the lattice-filling curve squeezed between the coalescing random walks also converges to a space-filling curve $(X_t, L_t(X_t))_{t\geq 0}$ in $\mathbb{R}\times[0,\infty)$, squeezed between paths in the forward and backward Brownian webs. Here, X_t is the position of the so-called *true self-repelling motion*, and $(L_t(x))_{x\in\mathbb{R}}$ is its occupation time density at time t. Comparing with the discrete model, one sees that when a deterministic point $(x,y)\in\mathbb{R}\times[0,\infty)$ is first reached by the space-filling curve $(X_\tau, L_\tau(X_\tau))$ at some random time τ, $(L_\tau(a))_{a\leq x}$, respectively $(L_\tau(a))_{a\geq x}$, must equal the path in the forward, respectively backward, Brownian web starting at (x,y). Furthermore, the area under $(L_\tau(a))_{a\in\mathbb{R}}$ must be exactly equal to τ. In other words, from the realization of the forward/backward Brownian webs, one can determine when a deterministic point (x,y) is reached by the space-filling curve. By considering a deterministic countable dense set of $(x,y)\in\mathbb{R}\times[0,\infty)$, one can then determine the entire trajectory of $(X_t)_{t\geq 0}$. This is how the true self-repelling motion was constructed by Tóth and Werner in [99], which heuristically is a process with a drift given by the negative of the gradient of its occupation time density. It has the unusual scaling invariance of $(X_{at}/a^{2/3})_{t\geq 0}\overset{\text{dist}}{=}(X_t)_{t\geq 0}$ for any $a>0$, and it has finite variation of order $3/2$. For further details, see [99, 26].

6.7.2.4 Planar Aggregation Models

In [76], Norris and Turner discovered an interesting connection between the Hastings-Levitov planar aggregation model and the coalescing Brownian flow. We briefly explain the model and the connection here. At time 0, we start with a unit disc denoted by K_0. A small particle P_1, which we assume to be a disc of radius $\delta>0$ for simplicity, is attached to K_0 at a uniformly chosen random point on the boundary ∂K_0. This defines the new aggregate $K_1=K_0\cup P_1$. To define the next aggregate K_2, we apply a conformal map F_1 which maps K_1 back to the disc K_0 (more accurately, F_1 is a conformal map from $K_1^c\cup\{\infty\}$ to $K_0^c\cup\{\infty\}$, uniquely determined by the condition that $F_1(z)=Cz+O(1)$ for some $C>0$ as $|z|\to\infty$). A new particle P_2 of radius δ is then attached randomly at the boundary of $K_0=F_1(K_1)$. Reversing F_1 then defines the new aggregate $K_2=F_1^{-1}(F_1(K_1)\cup P_2)$. The dynamics can then be iterated as illustrated in Figure 6.18. The new particles added to the aggregates $(K_n)_{n\in\mathbb{N}}$ at each step are no longer discs due to distortion by the conformal maps, and it was shown in [76] that new particles tend to pile on top of each other and form protruding fingers. Interestingly, the image of ∂K_0 under F_n forms a coalescing flow on the circle, where for $x,y\in\partial K_0\subset\partial K_n$, the length of the arc between $F_n(x)$ and $F_n(y)$ on the unit circle is proportional to the probability that a new particle will be attached to the corresponding part of ∂K_n

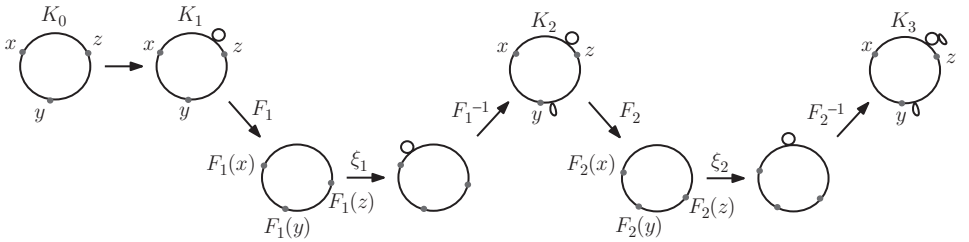

Figure 6.18. Hastings-Levitov planar aggregation model.

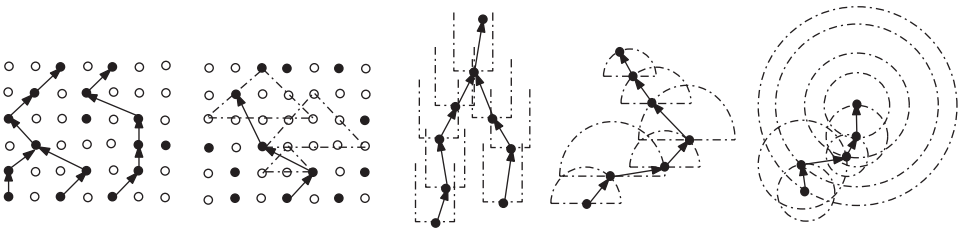

Figure 6.19. From left to right, illustration of paths in the drainage network of [50], the drainage network of [80], Poisson trees [43], the directed spanning forest [11], and the radial spanning tree [11].

between x and y. In the limit that the particle radius $\delta \downarrow 0$, while time is sped up by a factor of δ^{-3}, this coalescing flow can be seen as a localized disturbance flow, studied in [77], which converges to the coalescing Brownian flow on the circle w.r.t. the weak flow topology described in Section 6.7.1.1.

6.7.2.5 Drainage Networks and Directed Forests

Drainage networks are a class of models where coalescing paths arise naturally. First, a random subset of $\mathbb{Z}^{d+1}$ (or $\mathbb{R}^{d+1}$) is determined, which represents the water sources. Next, from each source (x,s), where $x \in \mathbb{R}^d$ and $s \in \mathbb{R}$, exactly one directed edge is drawn toward some other source (y,t) with $t > s$, representing the flow of water from (x,s) to (y,t). Examples include [50, 81], where the authors study a drainage network on $\mathbb{Z}^{d+1}$ with each vertex being a water source independently with probability $p \in (0,1)$. From each source (x,s), a directed edge is then drawn to the closest source in the next layer $\mathbb{Z}^d \times \{s+1\}$, and ties are broken by choosing each closest source with equal probability. There are also other variants such as in [5], where the directed edge from the source $(x,s) \in \mathbb{Z}^{d+1}$ connects to the closest source in the 45^o light cone rooted at (x,s), or as in [80] where the directed edge from (x,s) connects to the closest source in $\mathbb{Z}^d \times \{s, s+1, \ldots\}$ measured in ℓ_1 distance (see Figure 6.19).

There are also continuum space versions where the water sources form a homogeneous Poisson point process in $\mathbb{R}^{d+1}$. In the *Poisson trees* model

considered in [43], a directed edge is drawn from each source $(x,s) \in \mathbb{R}^{d+1}$ to the source in $\{(y,t) : |y-x| \le 1, t > s\}$ with the smallest t-coordinate. In the *directed spanning forest* model considered in [11], a directed edge is drawn from each source (x,s) to the closest source in $\mathbb{R}^d \times (s,\infty)$, measured in Euclidean distance.

Natural questions for such drainage networks include whether the directed edges form a single component, i.e., it is a tree rather than a forest. This has been shown to be the case for $d = 1$ and 2 for the drainage networks on $\mathbb{Z}^{d+1}$ described above [50, 5, 80] and for the Poisson trees [43], as well as for the directed spanning forest in $\mathbb{R}^{1+1}$ [21].

The collection of paths in the drainage network, obtained by following the directed edges, can be regarded as a collection of (dependent) coalescing random walk paths. Because the dependence is in some sense local in all the models described above, it is natural to conjecture that for $d = 1$, the collection of directed paths in the drainage network (after diffusive scaling) should converge to the Brownian web. This has indeed been verified for various drainage networks on $\mathbb{Z}^{1+1}$ [17, 22, 80], and for the Poisson trees on $\mathbb{R}^{1+1}$ [36]. The main difficulty in these studies lies in the dependence among the paths. The model considered in [80] is a discrete analogue of the directed spanning forest [11], and the dependence is handled using simultaneous regeneration times along multiple paths. Such arguments are not directly applicable to the directed spanning forest, and its convergence to the Brownian web remains open.

The directed spanning forest was introduced in [11] as a tool to study the *radial spanning tree*, where given a homogeneous Poisson point process $\Xi \subset \mathbb{R}^d$, a directed edge is drawn from each $x \in \Xi$ to the closest point $y \in \Xi \cup \{\mathrm{o}\}$ that lies in the ball with radius $|x|$ and centered at the origin o. If one considers directed paths in the radial spanning forest restricted to the region (in polar coordinates) $\{(r,\theta) : r \in [an,n], |\theta + \pi/2| \le n^{-b}\}$ for some $a \in (0,1)$ and $b \in (0,1/2)$, then after proper scaling, these paths are believed to converge to the so-called *Brownian bridge web* [46], which consists of coalescing Brownian bridges starting from $\mathbb{R} \times (-\infty, 0)$ and ending at $(0,0)$ (it can also be obtained from the Brownian web via a deterministic transformation of $\mathbb{R}^2$). For a couple of toy models, such a convergence has been established in [23, 46].

6.7.2.6 *Supercritical Oriented Percolation*

The oriented bond percolation model on $\mathbb{Z}^2_{\text{even}} := \{(x,t) \in \mathbb{Z}^2 : x + t \text{ is even}\}$ is defined by independently setting each oriented edge of the form $(x,t) \to (x \pm 1, t+1)$ to be either open with probability p, or closed with probability $1-p$, with retention parameter $p \in [0,1]$. The set of vertices in $\mathbb{Z}^2_{\text{even}}$ that can be reached from $z \in \mathbb{Z}^2_{\text{even}}$ by following open oriented edges is called the open cluster at z. It is known that there is a critical $p_c \in (0,1)$ such that for $p > p_c$, the open cluster at $(0,0)$

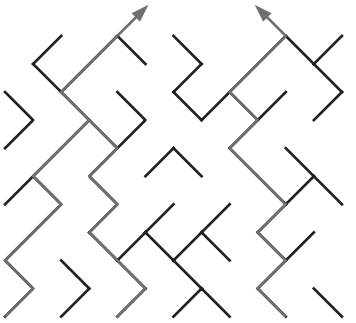

Figure 6.20. Rightmost infinite open paths in supercritical oriented percolation on $\mathbb{Z}^2_{\text{even}}$.

is infinite with positive probability, while for $p \leq p_c$, it is finite with probability one [24]. When the open cluster at $z \in \mathbb{Z}^2_{\text{even}}$ is infinite, z is called a percolation point.

In the supercritical regime $p > p_c$, the percolation points appear with a positive density, and starting from each percolation point $z = (x, t)$, we can find a rightmost path r_z among all the infinite open oriented paths starting from z (see Figure 6.20). It was shown by Durrett [24] that each $(r_z(n))_{n \geq t}$ satisfies a law of large numbers with drift $\alpha(p) > 0$, and subsequently Kuczek [58] showed that $r_z(n) - n\alpha$ satisfies a central limit theorem with variance $n\sigma(p)^2$ for some $\sigma(p) > 0$. Kuczek's central limit theorem can be easily extended to path level convergence to Brownian motion. It is then natural to ask what is the joint scaling limit of r_z for all percolation points z. Wu and Zhang [105] conjectured that the scaling limit should be the Brownian web. This may be surprising at first, because for different percolation points z_1 and z_2, r_{z_1} and r_{z_2} both depend on the infinite future, and hence could be strongly dependent on each other. However, it was shown by Sarkar and Sun in [89] that each r_z can be approximated by a thin percolation exploration cluster, such that different clusters evolve independently before they meet and quickly merge after they meet. It is then shown in [89] that after proper centering and scaling, the collection of rightmost infinity open paths from percolation points indeed converge to the Brownian web.

6.7.3 Brownian Web, Critical Percolation, and Noise

There are close parallels between the Brownian web and the scaling limit of critical planar percolation, starting from the topology. Indeed, as noted in Section 6.7.1.2, both the paths topology of [40] and the tube topology of [14] were inspired by similar topologies for planar percolation. The Brownian web is the scaling limit of discrete coalescing random walks on $\mathbb{Z}^2_{\text{even}}$ as shown in Figure 6.1, which can be

seen as a dependent planar percolation model and enjoys the same duality as bond percolation on $\mathbb{Z}^2_{\text{even}}$. The scaling limit of critical planar percolation is invariant under conformal maps, while the Brownian web is invariant under diffusive scaling of $\mathbb{R}^2$, a signature of a model at criticality. The Brownian web and the scaling limit of critical planar percolation are both (and the only known examples of) two-dimensional black noise in the language of Tsirelson [95, 96], which is intimately linked to the notion of noise sensitivity and the existence of near-critical scaling limits. Indeed, the Brownian net can be regarded as a near-critical scaling limit obtained by perturbing the Brownian web. As explained in Section 6.4, such a perturbation can be carried out by Poisson marking the set of "pivotal" points (i.e., the $(1,2)$-points of the Brownian web) according to a natural local time measure, and then turning these points into branching points. For planar percolation, exactly the same procedure has been proposed in [19] to construct a near-critical scaling limit from the critical scaling limit, and it has been rigorously carried out in [48] and [49], where defining the natural local time measure on pivotal points alone has been a very challenging task. The fact that changing the configuration at a countable subset of $\mathbb{R}^2$, i.e., the Poisson marked pivotal points, is sufficient to alter the scaling limit indicates noise sensitivity and that the scaling limit is a black noise. Allowing the intensity of the Poisson marking of the pivotal points to vary continuously leads to dynamical evolutions of the critical scaling limit, where unusual behaviour may emerge at random "dynamical times."

In what follows, we will briefly review Tsirelson's theory of noise [95, 96], and in what sense the Brownian web and the scaling limit of critical planar percolation are both black noise. This is the key common feature that lies behind the many parallels between the two models. Such parallels will likely extend to other examples of two- or higher-dimensional black noise, once they are found. We will also briefly review results on the dynamical Brownian web.

6.7.3.1 Brownian Web, Black Noise, and Noise Sensitivity

Intuitively, the Brownian web and the scaling limit of critical planar percolation should satisfy the property that the configurations are independent when restricted to disjoint domains in $\mathbb{R}^2$. This calls for the continuous analogue of the notion of product probability spaces (or independent random variables), which led Tsirelson to define (see [96, Definition 3c1]):

Definition 6.7.1 *A (one-dimensional)* continuous product of probability spaces *consists of a probability space* $(\Omega, \mathcal{F}, P)$ *and sub-*σ*-fields* $\mathcal{F}_{s,t} \subset \mathcal{F}$ *(for all* $s < t$*) such that* $\mathcal{F}$ *is generated by* $\bigcup_{s<t} \mathcal{F}_{s,t}$ *and for any* $r < s < t$*,*

$$\mathcal{F}_{r,t} = \mathcal{F}_{r,s} \otimes \mathcal{F}_{s,t}, \tag{6.70}$$

which means that $\mathcal{F}_{r,t}$ is generated by $\mathcal{F}_{r,s} \cup \mathcal{F}_{s,t}$, and $\mathcal{F}_{r,s}$ is independent of $\mathcal{F}_{s,t}$ (i.e., $P(A \cap B) = P(A)P(B)$ for all $A \in \mathcal{F}_{r,s}$ and $B \in \mathcal{F}_{s,t}$).

We can think of $\mathcal{F}_{s,t}$ as the σ-field generated by a family of observables that only depend on what happens in the interval (s,t).

Remark 6.7.2 In dimension two or higher, we should equip $(\Omega, \mathcal{F}, P)$ with a family of sub-σ-fields $\mathcal{F}_D$, indexed by a Boolean algebra of domains $D \subset \mathbb{R}^2$ (take for example the Boolean algebra of sets generated by open rectangles), while the factorization property (6.70) becomes

$$\mathcal{F}_D = \mathcal{F}_{D_1} \otimes \mathcal{F}_{D_2} \tag{6.71}$$

whenever $D_1 \cap D_2 = \emptyset$ and $\overline{D} = \overline{D_1 \cup D_2}$. It turns out that (6.71) cannot hold without some restrictions on the domains. The lack of a canonical choice of the family of domains in dimensions two and higher is still an issue to be resolved [98, Section 1.6].

Besides independence on disjoint domains, we also expect the Brownian web and the scaling limit of percolation to be translation invariant. This additional assumption leads to the notion of a noise (see [96, Definition 3d1]).

Definition 6.7.3 (Noise) *A noise is a continuous product of probability spaces $(\Omega, \mathcal{F}, P)$, equipped with a family of sub-σ-fields $\mathcal{F}_D$ indexed by an algebra of domains $D \in \mathcal{D}$ in $\mathbb{R}^d$, which is homogeneous in the following sense. There exists a group of isomorphisms $(T_h)_{h \in \mathbb{R}^d}$ on $(\Omega, \mathcal{F}, P)$ such that $T_{h+h'} = T_h \circ T_{h'}$ and T_h sends $\mathcal{F}_D$ to $\mathcal{F}_{D+h}$ for all $D \in \mathcal{D}$ and $h, h' \in \mathbb{R}^d$.*

In a tour de force, the scaling limit of critical planar percolation was shown to be a two-dimensional noise by Schramm and Smirnov in [87], where the sub-σ-field $\mathcal{F}_D$ (for D with piecewise smooth boundary) is generated by the indicator random variables for open crossing of quads (homeomorphic images of the unit square) contained in D. The Brownian web was shown to be a two-dimensional noise by Ellis and Feldheim in [33], where $\mathcal{F}_D$ (for open rectangles D) is generated by the Brownian web paths restricted to D. What is remarkable is that both noises are so-called *black noise* as defined by Tsirelson [95, 96], and they are the only known examples in dimension two or higher.

The notion of a noise being *black* turns out to be equivalent to the notion of a noise being *sensitive*, while a noise being *classical* (such as white noise or Poisson noise) turns out to be equivalent to a noise being *stable*. A *non-classical* noise is a noise that contains in some sense a sensitive part. To explain the underlying ideas, it is instructive to first discretize and then pass to the continuum, instead of giving directly the definition for the continuum noise.

Let us fix $\Lambda = (-1,1)^2$, and consider its discretization $\Lambda_\delta := \Lambda \cap S_\delta \mathbb{Z}^2_{\text{even}}$, where $S_\delta : (x,t) \to (\delta x, \delta^2 t)$. To each $z \in \Lambda_\delta$, we assign an i.i.d. symmetric random variable $\sigma_z \in \{+1,-1\}$, which can be defined via the coordinate map on the probability space $\Omega_\delta = \{+1,-1\}^{\Lambda_\delta}$, equipped with the discrete topology and the uniform probability measure P_δ. What noise we obtain in the continuum limit $\delta \downarrow 0$ depends crucially on the observables we choose, which will generate the σ-field $\mathcal{F}$ for the limiting noise.

If we choose our observables to be linear functions of $\sigma := (\sigma_z)_{z\in\Lambda_\delta}$, i.e., $f_\delta := \delta^{3/2} \sum_{z\in\Lambda_\delta} f(z)\sigma_z$ for continuous functions $f : \Lambda \to \mathbb{R}$, then their joint distributions converge to a nontrivial limit and the limiting noise is the classical white noise, where $\mathcal{F}$ is generated by the family of Gaussian random variables indexed by such continuous f, which we interpret as the integral of f w.r.t. the underlying white noise.

If we choose our observables to be nonlinear functions of $(\sigma_z)_{z\in\Lambda_\delta}$, then non-classical noise may appear in the limit. As illustrated in Figure 6.1, $\sigma := (\sigma_z)_{z\in\Lambda_\delta}$ uniquely determines the collection of coalescing random walks in Λ_δ. If we choose our observables to be the indicator random variables for whether there is a random walk path crossing a prescribed tube T in Λ (see Section 6.7.1.2), then their joint distributions converge to a nontrivial limit and the limiting noise is the Brownian web, where $\mathcal{F}$ is generated by the tube crossing events.

Roughly speaking, whether the limiting noise is classical or not depends on whether there are nontrivial observables that are noise sensitive as $\delta \downarrow 0$, defined as follows.

Definition 6.7.4 (Noise sensitivity) *For $\varepsilon \in (0,1)$, let $\sigma^\varepsilon := (\sigma^\varepsilon_z)_{z\in\Lambda_\delta}$ be obtained from $(\sigma_z)_{z\in\Lambda_\delta}$ by independently replacing each σ_z with an independent copy of σ_z with probability ε. A sequence of random variables $f_\delta : \Omega_\delta \to \mathbb{R}$ with $\sup_\delta E[f_\delta^2(\sigma)] < \infty$ is called noise sensitive if for each $\varepsilon > 0$,*

$$\lim_{\delta\downarrow 0} (E[f_\delta(\sigma) f_\delta(\sigma^\varepsilon)] - E[f_\delta(\sigma)]^2) = 0, \tag{6.72}$$

and the sequence is called noise stable if

$$\lim_{\varepsilon\downarrow 0} \limsup_{\delta\downarrow 0} |E[f_\delta(\sigma) f_\delta(\sigma^\varepsilon)] - E[f_\delta^2(\sigma)]| = 0. \tag{6.73}$$

There is a rich theory of noise sensitivity for functions on $(\Omega_\delta, P_\delta)$, and we refer to the lecture notes by Garban and Steif [51] for a detailed exposition and applications to percolation. For $\{0,1\}$-valued functions f_δ, noise sensitivity implies that $f_\delta(\sigma)$ and $f_\delta(\sigma^\varepsilon)$ become asymptotically independent if an arbitrarily small, but fixed, portion of σ is resampled, while noise stability implies that as $\varepsilon \downarrow 0$, the probability that $f_\delta(\sigma)$ and $f_\delta(\sigma^\varepsilon)$ coincide tends to 1 uniformly in δ close to 0.

There is a simple criterion for noise sensitivity/stability using Fourier analysis on $(\Omega_\delta, P_\delta)$. Observe that the set of functions $\chi_S := \prod_{i\in S}\sigma_i$, for all $S \subset \Lambda_\delta$, is an orthonormal basis for $L^2(\Omega_\delta, P_\delta)$. Therefore any $f_\delta \in L^2(\Omega_\delta, P_\delta)$ admits the orthogonal decomposition

$$f_\delta = \sum_{S\subset\Lambda_\delta} \hat{f}_\delta(S)\chi_S = E[f_\delta] + \sum_{k=1}^{\infty}\sum_{|S|=k} \hat{f}_\delta(S)\prod_{i\in S}\sigma_i. \tag{6.74}$$

Definition 6.7.5 (Spectral measure and energy spectrum) *The coefficients* $\hat{f}_\delta(S)$ *in* (6.74) *are called the* Fourier-Walsh *coefficients for* f_δ*, the measure* μ_{f_δ} *on* $\{S : S \subset \Lambda_\delta\}$ *with* $\mu_{f_\delta}(S) := \hat{f}_\delta^2(S)$ *is called the* spectral measure *of* f_δ*, and the measure* E_{f_δ} *on* $\mathbb{N}_0$ *with* $E_{f_\delta}(k) = \sum_{|S|=k}\hat{f}_\delta^2(S)$ *is called the* energy spectrum *of* f_δ *(see e.g., [51]).*

It is easily seen that

$$E[f_\delta(\sigma)f_\delta(\sigma^\varepsilon)] = E[f_\delta]^2 + \sum_{k=1}^{\infty}\varepsilon^k\sum_{|S|=k}\hat{f}_\delta^2(S) = E[f_\delta]^2 + \sum_{k=1}^{\infty}\varepsilon^k E_{f_\delta}(k), \tag{6.75}$$

and hence a sequence f_δ with $\sup_\delta E[f_\delta^2] < \infty$ is noise sensitive if and only if $E_{f_\delta}(\{1,\ldots,k\}) \to 0$ for every $k \in \mathbb{N}$, and f_δ is stable if and only if E_{f_δ} is a tight family of measures on $\mathbb{N}_0$.

As we take the continuum limit $\delta \downarrow 0$, what will happen to the spectral measure of the observables that generate the σ-field of the limiting noise? There are several possibilities: either all nontrivial square-integrable observables in the limit will have a spectral measure that is supported on $\mathcal{C}_{\text{fin}}(\Lambda) := \{S \in \Lambda : |S| < \infty\}$, in which case the limiting noise is called *classical* (and noise stable); or all nontrivial observables in the limit will have a spectral measure supported on $\mathcal{C}(\Lambda)\backslash\mathcal{C}_{\text{fin}}(\Lambda)$, with $\mathcal{C}(\Lambda)$ being the set of closed subsets of Λ, in which case the noise is called *black* (and noise sensitive); or there are nontrivial observables of both types, supported either on $\mathcal{C}_{\text{fin}}(\Lambda)$ or $\mathcal{C}(\Lambda)\backslash\mathcal{C}_{\text{fin}}(\Lambda)$, in which case the noise is called *non-classical*.

In practice, classifying a noise via discrete approximation as described above is subtle, because such approximations are not unique, and convergence of the spectral measure on $\mathcal{C}$ (equipped with Hausdorff topology) does not imply convergence of the energy spectrum. There could be observables f for the noise whose spectral measure is supported on $\mathcal{C}_{\text{fin}}$, and yet the approximating f_δ is noise sensitive (see [95, Section 5c] for the notion of *block sensitivity* which overcomes such issues). Instead, we can classify noise directly. The following definition extends [95, Definition 3d2] to general dimensions (see also [98] for an abstract formulation without associating the sub-σ-fields with domains in $\mathbb{R}^d$).

Definition 6.7.6 (Non-classical noise and black noise) *Let $(\Omega, \mathcal{F}, P)$ be a noise equipped with a family of sub-σ-fields $\mathcal{F}_D$ indexed by an algebra of domains $D \in \mathcal{D}$ in $\mathbb{R}^d$. For any $f \in L^2(\Omega, \mathcal{F}, P)$, its spectral measure μ_f is defined to be the unique positive measure on $\mathcal{C}$ (the space of compact subsets of $\mathbb{R}^d$) with the property that*

$$\mu_f(M \in \mathcal{C} : M \subset D) = E[E[f \mid \mathcal{F}_D]^2], \qquad D \in \mathcal{D}. \tag{6.76}$$

If for every $f \in L^2(\Omega, \mathcal{F}, P)$, its spectral measure μ_f is supported on $\mathcal{C}_{\text{fin}} := \{M \in \mathcal{C} : |M| < \infty\}$, then the noise is called classical*; otherwise the noise is called* non-classical. *If the spectral measure of every non-constant $f \in L^2(\Omega, \mathcal{F}, P)$ is supported on $\mathcal{C} \backslash \mathcal{C}_{\text{fin}}$, then the noise is called* black.

White noise and Poisson noise are classical, which can be seen from their chaos expansions. The scaling limit of critical planar percolation is a two-dimensional black noise as explained in [87], while the Brownian web was first shown to be a one-dimensional black noise in [95, Section 7] and [96, Section 7] (with sub-σ-fields $(\mathcal{F}_{s,t})_{s<t}$ associated with the strips $\mathbb{R} \times (s,t)$), and then shown to be a two-dimensional black noise in [33]. Interestingly, if a noise is black, then for all non-constant $f \in L^2(\Omega, \mathcal{F}, P)$, μ_f is in fact supported on $\mathcal{C}_{\text{perf}} := \{M \in \mathcal{C} : M \text{ is a perfect set}\}$, a collection of uncountable sets [95, Theorem 6d3].

The procedure of independently resampling each of a collection of ± 1-valued symmetric random variables with probability ε can in fact be carried out for noise. This allows one to define directly noise sensitivity vs stability for each $f \in L^2(\Omega, \mathcal{F}, P)$, where conditions (6.72)–(6.73) just need to be modified to remove the δ-dependence, and σ and σ^ε represent the noise and its ε-resampled version. A noise is then classical if and only if all $f \in L^2(\Omega, \mathcal{F}, P)$ are stable, and black if and only if all non-constant $f \in L^2(\Omega, \mathcal{F}, P)$ are sensitive (see [96, Section 4] and [95, Section 5] for further details). In particular, ε-resampling a black noise leads to an independent copy of the noise. To see nontrivial dependence, a different resampling procedure is needed, which is where dynamical evolution via Poisson marking of pivotal points comes in.

6.7.3.2 Dynamical Brownian Web

Because the Brownian web is a black noise, the correct way to resample the randomness underlying the Brownian web in order to see nontrivial dependence is to resample the "pivotal" points (i.e., the $(1,2)$-points of the Brownian web). This was carried out in [73] and is intimately linked to the authors' marking construction of the Brownian net in the same paper (see Section 6.4). Given a Poisson point process $M_\lambda \subset \mathbb{R}^2$ with intensity measure given by $\lambda(>0)$ times the intersection local time measure $\ell(\mathrm{d}z)$ on the $(1,2)$-points of the Brownian web (see Proposition 6.4.1), a resampled web is obtained by simply flipping the sign of

each $(1,2)$ point in M_λ as defined in (6.24), instead of turning each such point into a branching point as in the construction of the Brownian net in Theorem 6.4.2. Since there is a natural coupling between M_λ for all $\lambda \geq 0$, we obtain a stochastic process in the space of compact sets of paths $(\mathcal{H}, d_{\mathcal{H}})$, which is in fact reversible w.r.t. the law of the Brownian web. This defines the so-called *dynamical Brownian web*. The dynamical Brownian web can also be constructed as the scaling limit of the *dynamical discrete web* introduced in [56], where *discrete web* refers to the collection of coalescing simple symmetric random walks on $\mathbb{Z}^2_{\text{even}}$ illustrated in Figure 6.1, and dynamical evolution is introduced by independently flipping at rate one the sign of the arrow from each $z \in \mathbb{Z}^2_{\text{even}}$. If space-time is rescaled by $(\varepsilon, \varepsilon^2)$, and the dynamical time is slowed down by a factor of ε, it can then be shown that the dynamical discrete web converges to the dynamical Brownian web [73].

Similar questions have been studied for percolation models [47]. There has been much study about dynamical percolation (see [84] for a survey). Furthermore, dynamical critical site percolation on the triangular lattice has been shown to converge to a continuum limit [49], the proof of which follows the same line as in the Brownian web setting (via Poisson marking of the pivotal points), except that the implementation in the percolation setting has been much more difficult.

Key questions in the study of dynamical percolation, dynamical Brownian web, and "dynamical" processes in general, are the existence of exceptional times when a given property that holds almost surely at a deterministic time fails. For dynamical percolation, where the stationary law of the process is that of the percolation model at criticality, one such property is the non-existence of an infinite cluster [84]. For the dynamical discrete web, questions on exceptional times have also been studied in [45] and [57]. In particular, it was shown in [45] that there exists a.s. a random set of dynamical times with full Lebesgue measure, when the random walk in the discrete web starting from the origin does not satisfy the Law of the Iterated Logarithm (LIL); the authors also extended this result to the dynamical Brownian web.

This is in striking contrast with the dynamical simple symmetric random walk on $\mathbb{Z}$ introduced in [15], where the sign of each increment of the walk $(X_n - X_{n-1})_{n\in\mathbb{N}}$ is independently flipped at rate one in dynamical time. In contrast with the dynamical discrete web, there is no exceptional time at which the LIL fails in this setting. The difference between the two models can be accounted for by the effect of each individual switching. In [15], a switch affects only one increment of the walk, which induces an order one perturbation in sup-norm. In contrast, changing the direction of a single arrow in the discrete web can change the path of the walk by a "macroscopic" amount. Indeed, the switching arrows encountered by the walk are comparable to the Poisson marked $(1,2)$-points encountered by a Brownian motion in the dynamical Brownian web, where each switching at a

marked $(1,2)$-point leads the Brownian motion to embark on an excursion from the original path, and the size of the excursion has a heavy-tailed distribution with infinite mean. This is why the random walk is more sensitive to the dynamics on the discrete web than to the dynamics on the random walk considered in [15].

6.8 Open Questions

In this section, we discuss some open questions, including how the Brownian web and net could be useful in the study of more general one-dimensional interacting particle systems, conjectures on the geometry of the Brownian net, and some other miscellaneous questions.

6.8.1 Voter Model Perturbations and Brownian Net with Killing

One further direction of research is to use the Brownian web and net to study one-dimensional interacting particle systems with migration, birth/branching, coalescence, and death. Such a particle system often arises as the spatial genealogies of a dual particle system. In suitable parameter regimes, such a particle system may converge to the Brownian net, which incorporates diffusive particle motion, branching and coalescence, and killing effect can be added by defining the so-called *Brownian net with killing*, which we will explain below. Using the convergence of the discrete particle systems to a continuum model which is amenable to analysis, one hopes to draw conclusions for the discrete particle systems, such as local survival vs extinction, as well as convergence to equilibrium.

In dimensions three and higher, a study in this spirit has been carried out by Cox, Durrett and Perkins in [18], where they studied a class of interacting particle systems on $\mathbb{Z}^d$ whose transition rates can be written as perturbations of the transition rates of a (possibly non-nearest neighbor) voter model. Their interest stemmed from the observation that several models in ecology and evolution, as well as in statistical mechanics – the spatial Lotka-Volterra model [70], the evolution of cooperation [78], nonlinear voter models [67] – can be seen as perturbations of the voter model in a certain range of their parameter space. In general, these models are difficult to study either because of a lack of monotonicity (Lotka-Volterra model) or because of the intrinsic complexity of the model (certain models of evolution of cooperation). However, by considering them as small perturbations of the voter model, it was shown in [18] that for $d \geq 3$, the properly rescaled local density of particles converges to the solution of a reaction diffusion equation, and that properties of the underlying particle system (such as coexistence of species in the Lotka-Volterra model) can be derived from the behavior of this Partial Differential Equation (PDE). Implementing an analogous program in

dimensions one and two is still an open problem, although in dimension two, it is believed that the behavior of voter model perturbations should be similar to the case $d = 3$.

The genealogies of the voter model are given by coalescing random walks [61, 54], and the type of perturbations of the voter model in [18] lead to branching of the genealogies. In dimensions $d \geq 3$, because of transience of the random walk, such a collection of branching-coalescing genealogies was shown in [18] to converge (under suitable scaling and moment assumptions on the random walk and the branching mechanism) to a system of *independent* branching Brownian motions. The branching mechanism is given explicitly in terms of the "microscopic" descriptions of the system and also takes into account coalescence at the mesoscopic level. In dimension one, such a collection of branching-coalescing spatial genealogies should converge to the Brownian net under suitable scaling of parameters and space-time. One is then hopeful that in dimension one, one can also draw conclusions for the particle systems from such a convergence. One model that falls in this class is the biased annihilating branching process, introduced by Neuhauser and Sudbury in [75]. Another class is a family of one-dimensional models of competition with selection and mutation (see e.g., [7]). However, proving convergence of the particle systems to the Brownian net is a serious challenge.

6.8.1.1 Brownian Net with Killing

We briefly recall here the Brownian net with killing, which allows the modeling of the death of a particle. A discrete space analogue has been introduced in [68]. Let $b, \kappa >= 0$ be such that $b + \kappa \leq 1$, corresponding respectively to the branching and killing parameters of the system. From each point in $\mathbb{Z}^2_{\text{even}}$, an arrow is drawn from (x,t) to either $(x-1,t+1)$ or $(x+1,t+1)$ with equal probability $\frac{1}{2}(1-b-\kappa)$; two arrows are drawn with probability b; no arrow is drawn with probability κ. Similar to the discrete net (see Figure 6.9), we can consider the set of all paths starting in $\mathbb{Z}^2_{\text{even}}$ by following arrows until the path terminates. This defines an infinite collection of coalescing random walks which branch with probability b and are killed with probability κ. This model encompasses several classical models from statistical mechanics. For $p \in (0,1)$ and $b = p^2$, $\kappa = (1-p)^2$, one recovers the standard one-dimensional oriented percolation model. When $b = \kappa = 0$, the trajectories are distributed as coalescing random walks.

In [74], it is shown that when the branching and killing parameters depend on some small scaling parameter ε such that $\lim_{\varepsilon \to 0} b_\varepsilon / \varepsilon = b_\infty$ and $\lim_{\varepsilon \to 0} \kappa_\varepsilon / \varepsilon^2 = \kappa_\infty$, then the collection of branching-coalescing random walks with killing introduced above converges (after proper rescaling) to a continuum object called the Brownian net with killing with parameter $(b_\infty, \kappa_\infty)$. The latter scaling of the

parameters was motivated by the stochastic q-color Potts model on $\mathbb{Z}$. Indeed, in [74], it was conjectured that this model at large inverse temperature β, is asymptotically dual to a system of (continuous time) branching-coalescing random walks with killing, with branching and killing parameters

$$b_\beta = \frac{q}{2} e^{-\beta} \quad \text{and} \quad \kappa_\beta = q e^{-2\beta},$$

where q is the number of colors in the system. Thus the Brownian net could be used to construct a natural scaling limit for the stochastic Potts model at low temperature in dimension one.

The Brownian net with killing can also be directly constructed as follows. Let $\mathcal{N}_b$ be a Brownian net with branching parameter b. For every realization of the Brownian net, one can introduce a natural time-length measure on the set of points of type (p,p) in $\mathbb{R}^2$ w.r.t. the Brownian net (see Theorem 6.3.24). More precisely, for every Borel set $E \subset \mathbb{R}^2$, define the time-length measure as

$$\mathcal{L}(E) = \int_{\mathbb{R}} |\{x : (x,t) \in E \text{ and is of type (p,p)}\}| \, \mathrm{d}t. \tag{6.77}$$

The Brownian net with killing with parameter (b,κ) is then obtained by killing the Brownian net paths at the points of a Poisson point process on $\mathbb{R}^2$ with intensity measure $\kappa\mathcal{L}$.

It was shown in [74] that the Brownian net with killing undergoes a phase transition as the killing parameter κ varies, similar to the percolation transition in oriented percolation. There is reason to believe [53] that the Brownian net with killing belongs to the same universality class as oriented percolation in $\mathbb{Z}^{1+1}$, in which case it may serve as a simpler model to study this universality class.

6.8.2 Fractal Structure of the Brownian Net

We explain here some open questions concerning the geometry of the Brownian net. Recall from Remark 6.3.20 the set $\mathcal{N}(*,-\infty)$ of paths started at time $-\infty$, called the *backbone* of the Brownian net $\mathcal{N}$. Let $N := \{(x,t) \in \mathbb{R}^2 : x = \pi(t) \text{ for some } \pi \in \mathcal{N}(*,-\infty)\}$ be the graph of the backbone. By the image set property (Proposition 6.3.17), every continuous function $\pi : \mathbb{R} \to \mathbb{R}$ that lies entirely inside N is a path in the backbone. In view of this, much information can be read off from the random closed set $N \subset \mathbb{R}^2$.

If we set $\xi_t := \{x \in \mathbb{R} : (x,t) \in N\}$, then $(\xi_t)_{t\in\mathbb{R}}$ is just the branching-coalescing point set in stationarity (see Section 6.3.5), which is in fact a Markov process [86, Section 1.9–1.10]. At each deterministic time t, ξ_t is distributed a Poisson point process with intensity two, which is a reversible law for this Markov process. As

noted in Remark 6.3.20, by reversibility, the law of $N \subset \mathbb{R}^2$ (and likewise the collection of paths $\mathcal{N}(*,-\infty)$) is symmetric under time reversal.

It is clear that N has some sort of fractal structure. The main reason for this is the "infinite" branching rate of the Brownian particles that ξ_t consists of. Indeed, in each open time interval, a path in $\mathcal{N}(*,-\infty)$ will split infinitely often into two paths, which usually coalesce again after a very short time. But during the time that there are two paths, these too will for short moments split into two further paths, and so on, ad infinitum (see Figure 6.21).

Natural questions are: 1) is this the only way in which N has a fractal structure? 2) how can we formulate this rigorously? and 3) can we prove this? Here is an attempt to answer question 2. Let S and M be the set of all *separation points* and *meeting points* (i.e., points of types $(\mathrm{p},\mathrm{pp})_\mathrm{s}$ and (pp,p), respectively, as defined in Section 6.2.5) that lie on N. Set $S_* := S \cup \{(*,-\infty)\}$ and $M_* := M \cup \{(*,\infty)\}$. Define $\phi : S_* \to M_*$ by

$$\phi(z) := \text{ the first meeting point of the left-most and right-most paths starting at } z, \tag{6.78}$$

where we set $\phi(z) := (*,\infty)$ if the left-most and right-most paths starting at z never meet, and for definitiveness we also define $\phi(*,-\infty) := (*,\infty)$.

By the symmetry of the backbone with respect to time reversal, we define $\phi : M_* \to S_*$ analogously, following paths in $\mathcal{N}(*,-\infty)$ backwards in time. (Note that this is different from our usual way of looking backwards. Instead of following dual paths, here we follow *forward* paths in the backbone backwards in time.)

A lot about the structure of the backbone can be understood in terms of the map ϕ. For example, separation points z such that $\phi(z) = (*,\infty)$ are $(-\infty,\infty)$-relevant separation points as defined in Section 6.4.1. By Proposition 6.4.7, these form a locally finite subset of $\mathbb{R}^2$. If we have a pair of points $z \in S$ and $z' \in M$ such that $\phi(z) = z'$ and $\phi(z') = z$, then this corresponds to a path that splits at z into two paths, which meet again at z' (see Figure 6.21). If $z = (x,t)$ and $z' = (x',t')$ form such a pair, then the left-most and right-most paths starting at z, up to their first meeting point z', enclose a compact set that we will call a *bubble*.

We conjecture that bubbles are the *only* source of the fractal structure of N. The following two conjectures make this idea precise.

Conjecture 6.8.1 (Bubble hypothesis)

(a) *Almost surely, there does not exist an infinite sequence of points* $(z_k)_{k\geq 0}$ *in* $S_* \cup M_*$*, all different from each other, such that* $\phi(z_{k-1}) = z_k$ *for all* $k \geq 1$.

(b) *Almost surely, there does not exist an infinite sequence of points* $(z_k)_{k\geq 0}$ *in* $S_* \cup M_*$*, all different from each other, such that* $\phi(z_k) = z_{k-1}$ *for all* $k \geq 1$.

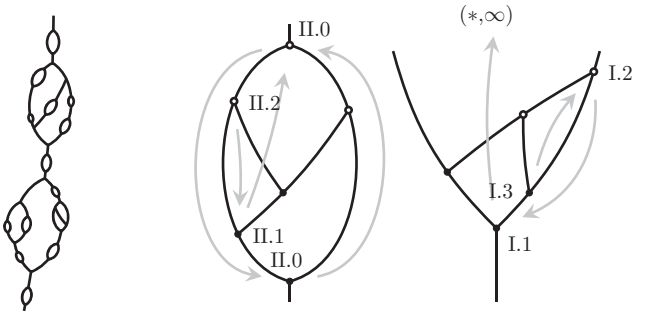

Figure 6.21. Left: schematic depiction of the fractal structure of the backbone. Middle and right: illustration of the classification of separation and meeting points on the backbone with the map ϕ (gray arrows). The leftmost picture contains two separation and two meeting points of type II. 1; all others are of type II. 0.

If Conjecture 6.8.1 (a) is correct, then for each separation or meeting point z that lies on the backbone, after some moment in the sequence $(z, \phi(z), \phi^2(z), \dots)$, the same points start to repeat. Simple geometric considerations show that the limiting cycle must have length 2, i.e., there is an integer $n \geq 0$ such that $\phi^{n+2}(z) = \phi^n(z)$. Letting n denote the smallest such integer, we say that z is of *class* I. n if $\phi^n(z) = (*, \infty)$ or $(*, -\infty)$ and of *class* II. n if $\phi^n(z) \in \mathbb{R}^2$. Note that separation and meeting points of type II. 0 are the bottom and top points of bubbles, as we have just defined them. Points of type II. n all lie inside bubbles, with larger n leading to a more complex left-right crossover pattern in the containing bubble (see Figure 6.21). The following conjecture says that most bubbles have a very simple internal structure.

Conjecture 6.8.2 (Bubble complexity) *Let* $\mathcal{C}_{\mathrm{I,n}}$, *respectively* $\mathcal{C}_{\mathrm{II,n}}$, *denote the set of meeting and separation points on the backbone* N *which are of class* I. n, *respectively* II. n. *Then*

(a) *The sets* $\mathcal{C}_{\mathrm{II},0}$ *and* $\mathcal{C}_{\mathrm{II},1}$ *are dense in* N.
(b) *The sets* $\mathcal{C}_{\mathrm{II,n}}$ *with* $n \geq 2$ *and* $\mathcal{C}_{\mathrm{I,n}}$ *with* $n \geq 1$ *are locally finite subsets of* $\mathbb{R}^2$.

As motivation for these conjectures, we state a somewhat more applied problem. Modulo a time reversal, the set $\mathcal{W}(\{0\} \times \mathbb{R})$ of all paths in the Brownian web can be interpreted as the spatial genealogies of a population living in one-dimensional space (see Section 6.7.2.1). Let M be the set of all meeting points of paths in $\mathcal{W}(\{0\} \times \mathbb{R})$, i.e.,

$$M := \{(x,t) \in \mathbb{R}^2 : t > 0, \exists \pi, \pi' \in \mathcal{W}(\{0\} \times \mathbb{R}) \text{ s.t. } \pi(t) = x = \pi'(t),\ \pi < \pi' \text{ on } (0,t)\}. \tag{6.79}$$

Biologically, we can interpret a point $z = (x,t)$ that is the first meeting point of paths $\pi_{(x,0)}$ and $\pi_{(x',0)}$ as the *most recent common ancestor* (MRCA) of x and x',

that lived a time t in the past. It is easy to see that M is a locally finite subset of $\mathbb{R} \times (0, \infty)$.

Moving away from neutral evolution, one can interpret paths in the Brownian net as *potential* genealogies (cf. the biased voter model in Section 6.7.2.2), where in order to determine the true genealogy of an individual, one has to have information about what happens at selection events, which correspond to separation points. In this case, replacing $\mathcal{W}(\{0\} \times \mathbb{R})$ in (6.79) by $\mathcal{N}(\{0\} \times \mathbb{R})$, we can interpret the resulting set M as the set of *potential most recent common ancestors* (PMRCAs) of individuals living at time zero.

Conjecture 6.8.3 (Potential most recent common ancestors) *If we replace* $\mathcal{W}(\{0\} \times \mathbb{R})$ *in* (6.79) *by* $\mathcal{N}(\{0\} \times \mathbb{R})$*, then the set* M *is a locally finite subset of* $\mathbb{R} \times (0, \infty)$.

The definition of a potential most recent common ancestor (PMRCA) is somewhat reminiscent of the definition of a $(0, \infty)$-relevant separation point, except that in order to determine whether a meeting point is a PMRCA, we have to follow forward paths backwards in time. This is similar to the definition of the map $\phi : M_* \to S_*$ and indeed, Conjecture 6.8.3 seems to be closely related to our previous two conjectures. In particular, meeting points that lie inside a bubble (properly defined w.r.t. the finite time horizon 0) can never be PMRCAs.

6.8.3 Miscellaneous Open Questions

We collect below some other interesting questions.

(1) Is the Brownian net integrable in the sense that the branching-coalescing point set defined in (6.20) admits an explicit characterization, similar to the Pfaffian point process characterization of the coalescing point set discussed in Remark 6.2.10?
(2) Find effective criteria for general branching-coalescing particle systems, where paths can cross, to converge to the Brownian net. This would extend Theorem 6.6.11 for the non-crossing case, and pave the way for the study of voter model perturbations discussed in Section 6.8.1.
(3) Can one formulate a well-posed martingale problem for the branching-coalescing point set $(\xi_t)_{t \geq 0}$ defined in (6.20)? This may offer an alternative route to prove convergence to the Brownian net, which is based on generator convergence and avoids the paths topology.
(4) For the Howitt-Warren flow $(K^{\uparrow}_{s,t}(x, \mathrm{d}y))_{s<t, x \in \mathbb{R}}$ introduced in Section 6.5.1, which gives the transition probability kernels of a random motion in a continuum space-time random environment constructed from the Brownian

web and net, can one show that for any $x \in \mathbb{R}$, $-\log K^{\uparrow}_{0,t}(0,[tx,tx+1])$ has Tracy-Widom GUE fluctuations on the scale $t^{1/3}$ as $t \to \infty$, similar to recent results in [12] for random walks on $\mathbb{Z}$ in special i.i.d. Beta-distributed space-time random environments (cf. Remark 6.5.15)?

(5) In [35], Evans, Morris and Sen showed that coalescing stable Lévy processes on $\mathbb{R}$ with stable index $\alpha \in (1,2)$ come down from infinity, i.e., starting from everywhere on $\mathbb{R}$, the coalescing Lévy processes become locally finite on $\mathbb{R}$ at any time $t > 0$. This suggests the existence of a family of *Lévy webs*. In fact, using the tube topology reviewed in Section 6.7.1.2 and the property of coming down from infinity, it should be straightforward to construct the Lévy web w.r.t. the tube topology. For coalescing Brownian flow on the Sierpinski gasket, this was carried out in [14]. Can one construct the Lévy web in the paths topology, and what type of special points may arise in the spirit of Theorem 6.2.11?

By scaling invariance, the density of the coalescing stable Lévy flow on $\mathbb{R}$ with index $\alpha \in (1,2)$ should be $\rho_t = \frac{C}{t^{1/\alpha}}$. Identifying C would allow one to determine the sharp asymptotic rate of decay for the density of coalescing random walks on $\mathbb{Z}$ in the domain of attraction of a stable Lévy process, similar to [72, Corollary 7.1] and [88] in the Brownian case.

(6) The Brownian web has been shown by Ellis and Feldheim [33] to be a two-dimensional black noise, equipped with a family of sub-σ-fields indexed by finite unions of open rectangles in $\mathbb{R}^2$. Is there a maximal extension of the family of sub-σ-fields, indexed by a Boolean algebra of domains in $\mathbb{R}^2$ that include the rectangles, such that the Brownian web remains a noise (cf. Remark 6.7.2)? A similar question is also open for the scaling limit of critical planar percolation [98, Section 1.6].

Recently Tsirelson [97] showed by general arguments that there exists some continuous path f starting at time $-\infty$, such that $\mathcal{F}_{f_-}$ and $\mathcal{F}_{f_+}$, the σ-fields generated by Brownian web paths restricted respectively to the left or the right of the graph of f, do not jointly generate the full σ-field $\mathcal{F}$. Can we characterize the set of f for which the factorization property $\mathcal{F} = \mathcal{F}_{f_-} \vee \mathcal{F}_{f_+}$ holds? (The case $f \equiv 0$ has been treated in [33]).

(7) We expect the Brownian net with killing (see Section 6.8.1.1) to belong to the same universality class as oriented percolation (OP) on $\mathbb{Z}^{1+1}$, which corresponds to the so-called Reggeon field theory [20, 65] (see also [53] for an extensive survey on the OP universality class). The Brownian net without killing belongs to a different universality class. But is there also a field theory corresponding to the Brownian net? A positive indication is that the Brownian net has competing effects of instantaneous coalescence vs infinite effective

rate of branching, which is very much in the spirit of renormalizations in field theories.

(8) The Brownian web appears in the scaling limit of super-critical oriented percolation, as discussed in Section 6.7.2.6, and the Brownian net is expected to arise if the percolation parameter p is allowed to vary in a small interval. However, the most interesting question is: what is the scaling limit of critical (and near-critical) oriented percolation (OP) on $\mathbb{Z}^{1+1}$? The Brownian net with killing provides a simpler model to study the OP universality class. We expect that in the near-critical scaling limit, one would obtain a family of models which interpolates between the Brownian net and the scaling limit of critical oriented percolation. However, such a goal appears far out of reach at the moment, because there are no conjectures at all on what the critical scaling limit might be, and neither have critical exponents been shown to exist.

Acknowledgments

This article is based on lectures given by R. Sun in the trimester program *Disordered Systems, Random Spatial Processes and Some Applications*, at the Institut Henri Poincaré in the spring of 2015. We thank the program and the organizers for the opportunity to lecture on this topic. We thank all our collaborators and colleagues with whom we have had valuable discussions. In particular, R. Sun and E. Schertzer would like to thank Chuck Newman and Krishnamurthi Ravishankar for introducing them to this topic. R. Sun also thanks University of Warwick, Leiden University, and Academia Sinica for hospitality, where part of these notes were written. R. Sun is supported by AcRF Tier 1 grant R-146-000-185-112. J.M. Swart is sponsored by GACR grant 15-08819S.

References

[1] R. Arratia, *Coalescing Brownian motions on the line*. Ph.D. Thesis, University of Wisconsin, Madison, 1979.

[2] R. Arratia, *Coalescing Brownian motions and the voter model on* $\mathbb{Z}$. Unpublished partial manuscript.

[3] M. Aizenman, *Scaling limit for the incipient spanning clusters*, in *Mathematics of Multiscale Materials: Percolation and Composites* (K. Golden et al, eds.). New York: Springer, IMA Vol. Math. Appl. 99, (1998).

[4] M. Aizenman and A. Burchard, *Hölder regularity and dimension bounds for random curves*. Duke Math. J. 99, 419–453 (1999).

[5] S. Athreya, R. Roy and A. Sarkar, *Random directed trees and forest – drainage networks with dependence*. Electron. J. Probab. 13, 2160–2189 (2008).

[6] S. Athreya and R. Sun, *One-dimensional voter model interface revisited*. Electron. Commun. Probab. 16, 792–800 (2011).

[7] S. Athreya and J. M. Swart. *Systems of branching, annihilating, and coalescing particles*. Electron. J. Probab. 17, 1–32 (2012).
[8] D. ben-Avraham, M. A. Burschka and C. R. Doering, *Statics and dynamics of a diffusion-limited reaction: anomalous kinetics, nonequilibrium self-ordering, and a dynamic transition*. J. Stat. Phys. 60, 695–728 (1990).
[9] P. Billingsley, *Convergence of probability measures, 2nd edition*. New York: John Wiley & Sons, (1999).
[10] N. Berestycki, *Recent progress in coalescent theory*. Rio de Janeiro: Sociedade Brasileira de Matemática, Ensaios Matemáticos 16, (2009).
[11] F. Bacelli and C. Bordenave. *The radial spanning tree of a Poisson point process*. Ann. Appl. Probab. 17, 305–359 (2007).
[12] G. Barraquand and I. Corwin, *Random-walk in Beta-distributed random environment*. (2015). arXiv:1503.04117.
[13] C. Bezuidenhout and G. Grimmett, *The critical contact process dies out*. Ann. Probab. 18, 1462–1482 (1990).
[14] N. Berestycki, C. Garban and A. Sen, *Coalescing Brownian flows: a new approach*. (2013). arXiv:1307.4313. Ann. Probab. 43, 3177–3215 (2015).
[15] I. Benjamini, O. Háaggstrom, Y. Peres and J. E. Steif, *Which properties of a random sequence are dynamically sensitive?* Ann. Probab. 31, 1–34 (2003).
[16] S. Belhaouari, T. Mountford, R. Sun and G. Valle, *Convergence results and sharp estimates for the voter model interfaces*. Electron. J. Probab. 11, 768–801 (2006).
[17] C. F. Coletti, E. S. Dias and L. R. G. Fontes, *Scaling limit for a drainage network model*. J. Appl. Probab. 46, 1184–1197 (2009).
[18] J. T. Cox, R. Durrett and E. Perkins, *Voter model perturbations and reaction diffusion equations*. Astérisque 349 (2013).
[19] F. Camia, L. R. G. Fontes and C. M. Newman, *Two-dimensional scaling limits via marked nonsimple loops*. Bull. Braz. Math. Soc. (N.S.) 37, 537–559 (2006).
[20] J. L. Cardy and R. L. Sugar, *Directed percolation and Reggeon field theory*. J. Phys. A: Math. Gen. 13, L423–L427 (1980).
[21] D. Coupier and V. C. Tran, *The 2D-directed spanning forest is almost surely a tree*. Random Structures & Algorithms. 42, 59–72 (2013).
[22] C. F. Coletti and G. Valle, *Convergence to the Brownian web for a generalization of the drainage network model*. Ann. Inst. H. Poincaré Probab. Statist. 50, 899–919 (2014).
[23] C. F. Coletti and L. A. Valencia, *Scaling limit for a family of random paths with radial behavior*. (2014). arXiv:1310.6929.
[24] R. Durrett, *Oriented percolation in two dimensions*. Ann. Probab. 12, 999–1040 (1984).
[25] D. Dhar, *Theoretical studies of self-organized criticality*. Physica A: Statistical Mechanics and its Applications. 369, 29–70 (2006).
[26] L. Dumaz, *A clever (self-repelling) burglar*, Electron. J. Probab. 17, 1–17 (2012).
[27] D. Dhar, *Fragmentation of a sheet by propagating, branching and merging cracks*. Journal of Physics A: Mathematical and Theoretical. 48, 17 (2015).
[28] C. R. Doering and D. ben-Avraham, *Interparticle distribution functions and rate equations for diffusion-limited reactions*. Phys. Rev. A. 38, 3035 (1988).
[29] P. Donnelly, S. N. Evans, K. Fleischmann, T. G. Kurtz and X. Zhou, *Continuum-sites stepping-stone models, coalescing exchangeable partitions and random trees* Ann. Probab. 28, 1063–1110 (2000).
[30] A. Depperschmidt, A. Greven and P. Pfaelhuber, *Marked metric measure spaces*. Electron. Commun. Probab. 16, 174–188 (2011).

[31] M. Damron and C. L. Winter, *A non-Markovian model of rill erosion*, Networks and Heterogeneous Media. 4, 731–753 (2009).
[32] T. Ellis, *Coalescing stochastic flows driven by Poisson random measure and convergence to the Brownian web*. Ph.D. thesis, University of Cambridge, 2010.
[33] T. Ellis and O. N. Feldheim. *The Brownian web is a two-dimensional black noise*. (2012). arXiv:1203.3585. Ann. Inst. H. Poincar Probab. Statist. 52, 162–172 (2016).
[34] A. Etheridge, N. Freeman and D. Straulino, *The Brownian net and selection in the spatial Lambda-Fleming-Viot process*. (2015). arXiv:1506.01158.
[35] S. N. Evans, B. Morris and A. Sen, *Coalescing systems of non-Brownian particles*, Probab. Theory Related Fields. 156, 307–342 (2013).
[36] P. A. Ferrari, L. R. G. Fontes and X.-Y. Wu, *Two-dimensional Poisson Trees converge to the Brownian web*. Ann. Inst. H. Poincaré Probab. Statist. 41, 851–858 (2005).
[37] L. R. G. Fontes, M. Isopi, C. M. Newman and D. Stein, *Aging in 1D discrete spin models and equivalent systems*. Physical Review Letters. 87, 110201 (2001).
[38] L. R. G. Fontes, M. Isopi, C. M. Newman and K. Ravishankar, *The Brownian web*. Proc. Nat. Acad. Sciences. 99, 15888–15893 (2002).
[39] L. R. G. Fontes, M. Isopi, C. M. Newman and K. Ravishankar, *The Brownian web: characterization and convergence*. (2003). math/0304119.
[40] L. R. G. Fontes, M. Isopi, C. M. Newman and K. Ravishankar, *The Brownian web: characterization and convergence*. Ann. Probab. 32(4), 2857–2883 (2004).
[41] L. R. G. Fontes, M. Isopi, C. M. Newman and K. Ravishankar, *Coarsening, nucleation, and the marked Brownian web*. Ann. Inst. H. Poincaré Probab. Statist. 42, 37–60 (2006).
[42] C. M. Fortuin, P. Kasteleyn and J. Ginibre, *Correlation inequalities on some partially ordered sets*. Commun. Math. Phys. 22, 89–103 (1970).
[43] P. A. Ferrari, C. Landim and H. Thorisson, *Poisson trees, succession lines and coalescing random walks*. Ann. Inst. H. Poincaré Probab. Statist. 40, 141–152 (2004).
[44] L. R. G. Fontes and C. M. Newman, *The full Brownian web as scaling limit of stochastic flows*. Stoch. Dyn. 6, 213–228 (2006).
[45] L. R. G. Fontes, C. M. Newman, K. Ravishankar and E. Schertzer, *Exceptional times for the dynamical discrete web*. Stoch. Proc. Appl. 119, 2832–2858 (2009).
[46] L. R. G. Fontes, L. A. Valencia and G. Valle, *Scaling limit of the radial Poissonian web*. Electron. J. Probab. 20, paper no. 31 (2015).
[47] G. Grimmett, *Percolation*. New York: Springer-Verlag, (1999).
[48] C. Garban, G. Pete and O. Schramm, *Pivotal, cluster and interface measures for critical planar percolation*. J. Amer. Math. Soc. 26, 939–1024 (2013).
[49] C. Garban, G. Pete and O. Schramm, *The scaling limits of near-critical and dynamical percolation*. (2013). arXiv:1305.5526.
[50] S. Gangopadhyay, R. Roy and A. Sarkar, *Random Oriented Trees: A Model of Drainage Networks*. Ann. App. Probab. 14, 1242–1266 (2004).
[51] C. Garban and J. Steiff, *Noise sensitivity and percolation*, in *Probability and statistical physics in two and more dimensions* (D. Ellwood et al. eds.). Providence, RI: Amer. Math. Soc. 49–154 (2012).
[52] A. Greven, R. Sun and A. Winter, *Continuum space limit of the genealogies of interacting Fleming-Viot processes on* $\mathbb{Z}$. (2015). arXiv:1508.07169.
[53] H. Hinrichsen, *Non-equilibrium critical phenomena and phase transitions into absorbing states*. Advances in Physics. 49, 815–958 (2000).
[54] R. Holley and T. Liggett. *Ergodic theorems for weakly interacting infinite systems and the voter model*. Annals of Probab. 3, 643–663 (1975).

[55] C. Howitt and J. Warren, *Consistent families of Brownian motions and stochastic flows of kernels*. Ann. Probab. 37, 1237–1272 (2009).
[56] C. Howitt and J. Warren, *Dynamics for the Brownian web and the erosion flow*. Stochastic Processes Appl. 119, 2028–2051 (2009).
[57] D. Jenkins, *Superdiffusive and subdiffusive exceptional times in the dynamical discrete web*. Stochastic Processes Appl. 125, 3373–3400 (2015).
[58] T. Kuczek, *The central limit theorem for the right edge of supercritical oriented percolation*. Ann. Probab. 17, 1322–1332 (1989).
[59] I. Karatzas and S.E. Shreve, *Brownian motion and stochastic calculus, 2nd edition.* New York: Springer-Verlag, (1991).
[60] K. Krebs, M.P. Pfannmüller, B. Wehefritz and H. Hinrichsen, *Finite-size scaling studies of one-dimensional reaction-diffusion systems. Part I. Analytical results*. J. Stat. Phys. 78, 1429–1470 (1995).
[61] T. Liggett, *Interacting particle systems*. New York: Springer, (1975).
[62] Y. Le Jan and S. Lemaire, *Products of beta matrices and sticky flows*. Probab. Th. Relat. Fields. 130, 109–134 (2004).
[63] Y. Le Jan and O. Raimond, *Flows, Coalecence and Noise*. Annals of Probab. 32, 1247–1315 (2004).
[64] Y. Le Jan and O. Raimond, *Sticky flows on the circle and their noises*. Probab. Th. Relat. Fields. 129, 63–82 (2004).
[65] M. Moshe, *Recent developments in Reggeon field theory*. Physics Reports. 37, 255–345 (1978).
[66] I. Molchanov, *Theory of random sets*. Berlin: Springer (2005).
[67] J. Molofsky, R. Durrett, J. Dushoff, D. Griffeath and S. Levin, *Local frequency dependence and global coexistence*. Theor. Pop. Biol. 55, 270–282 (1999).
[68] Y. Mohylevskyy, C. M. Newman and K. Ravishankar, *Ergodicity and percolation for variants of one-dimensional voter models*. ALEA. 10, 485–504 (2013).
[69] R. Munasinghe, R. Rajesh, R. Tribe and O. Zaboronski. *Multi-scaling of the n-point density function for coalescing Brownian motions*. Commun. Math. Phys. 268, 717–725 (2006).
[70] C. Neuhauser and S. W. Pacala, *An explicitly spatial version of the Lotka-Volterra model with interspecific competition*. Ann. Appl. Probab. 9, 1226–1259 (2000).
[71] C. M. Newman and K. Ravishankar, *Convergence of the Tóth lattice filling curve to the Tóth-Werner plane filling curve*. ALEA Lat. Am. J. Probab. Math. Stat. 1, 333–345 (2006).
[72] C. M. Newman, K. Ravishankar and R. Sun, *Convergence of coalescing nonsimple random walks to the Brownian web*. Electron. J. Prob. 10, 21–60 (2005).
[73] C. M. Newman, K. Ravishankar and E. Schertzer, *Marking* $(1,2)$ *points of the Brownian web and applications*. Ann. Inst. Henri Poincaré Probab. Statist. 46, 537–574 (2010).
[74] C. M. Newman, K. Ravishankar and E. Schertzer, *The Brownian net with killing*. Stoch. Proc. and App. 125, 1148–1194 (2015).
[75] C. Neuhauser and A. Sudbury, *The biased annihilating branching process*. Adv. Appl. Prob. 25, 24–38 (1993).
[76] J. Norris and A. Turner, *Hastings-Levitov aggregation in the small-particle limit*. Comm. Math. Phys. 316, 809–841 (2012).
[77] J. Norris and A. Turner, *Weak convergence of the localized disturbance flow to the coalescing Brownian flow*. Ann. Probab. 43, 935–970 (2015).

[78] H. Ohtsuki, H. C. Hauert, E. Lieberman and M. A. Nowak, *A simple rule for the evolution of cooperation on graphs and social networks*. Nature. 441, 502–505 (2006).
[79] V. Privman, *Nonequilibrium statistical mechanics in one dimension*. Cambridge: Cambridge University Press, (1997).
[80] R. Roy, K. Saha and A. Sarkar. *Random directed forest and the Brownian web*. Ann. Inst. H. Poincaré Probab. Statist. 52, 1106–1143 (2016).
[81] R. Roy, K. Saha and A. Sarkar, *Hack's law in a drainage network model: a Brownian web approach*. Ann. Appl. Probab. 26, 1807–1836 (2016).
[82] F. Spitzer, *Principles of random walk, 2nd edition*. New York: Springer-Verlag, (1976).
[83] D. Schwartz, *Applications of duality to a class of Markov processes*. Ann. Probab. 5, 522–532 (1977).
[84] J. Steif, *A survey of dynamical percolation. Fractal geometry and stochastics IV*. Progress in Probability. 61, 145–174 (2009).
[85] T. Seppäläinen, *Current fluctuations for stochastic particle systems with drift in one spatial dimension*, in *Ensaios Matemáticos [Mathematical Surveys]*, 18. Rio de Janeiro: Sociedade Brasileira de Matemática, (2010).
[86] R. Sun and J. M. Swart, *The Brownian net*. Ann. Probab. 36, 1153–1208 (2008).
[87] O. Schramm and S. Smirnov, *On the scaling limits of planar percolation*. Ann. Probab. 39, 1768–1814 (2011).
[88] A. Sarkar and R. Sun, *Brownian web and oriented percolation: density bounds*. RIMS Kokyuroku, No. 1805, Applications of Renormalization Group Methods in Mathematical Sciences, 90–101 (2012).
[89] A. Sarkar and R. Sun, *Brownian web in the scaling limit of supercritical oriented percolation in dimension* $1+1$. Electron. J. Probab. 18, Paper 21 (2013).
[90] K. Saha and A. Sarkar, *Convergence of drainage networks to the Brownian web*. (2015). arXiv:1508.06919.
[91] E. Schertzer, R. Sun and J.M. Swart, *Special points of the Brownian net*. Electron. J. Prob. 14, Paper 30, 805–864 (2009).
[92] E. Schertzer, R. Sun and J.M. Swart, *Stochastic flows in the Brownian web and net*. Mem. Amer. Math. Soc. 227, 1065 (2014).
[93] F. Soucaliuc, B. Tóth and W. Werner, *Reflection and coalescence between one-dimensional Brownian paths*. Ann. Inst. Henri Poincaré Probab. Statist. 36, 509–536 (2000).
[94] B. Tóth, *The "true" self-avoiding walk with bond repulsion on* $\mathbb{Z}$*: Limit theorems*. Ann. Probab. 23, 1523–1556 (1995).
[95] B. Tsirelson, *Scaling limit, noise, stability*. Lecture Notes in Mathematics. 1840, 1–106, (2004).
[96] B. Tsirelson, *Nonclassical stochastic flows and continuous products*. Probability Surveys. 1, 173–298 (2004).
[97] B. Tsirelson, *Random compact set meets the graph of nonrandom continuous function*. (2013). arXiv:1308.5112.
[98] B. Tsirelson, *Noise as a Boolean algebra of -fields*. Ann. Probab. 42, 311–353 (2014).
[99] B. Tóth and W. Werner, *The true self-repelling motion*. Probab. Theory Related Fields. 111, 375–452 (1998).
[100] R. Tribe, S.K. Yip and O. Zaboronski, *One dimensional annihilating and coalescing particle systems as extended Pfaan point processes*. Electron. Commun. Probab. 17, No. 40 (2012).

[101] R. Tribe and O. Zaboronski, *Pfaan formulae for one dimensional coalescing and annihilating systems*. Electron. J. Probab. 16, No. 76, 2080–2103 (2011).
[102] J. Warren, *The noise made by a Poisson snake*. Electron. J. Probab. 7, No. 21, 1–21 (2002).
[103] J. Warren, *Dyson's Brownian motions, intertwining and interlacing*. Electron. J. Probab. 12, No. 19, 573–590 (2007).
[104] T. Williams and R. Bjerknes, *A stochastic model for the spread of an abnormal clone through the basal layer of the epithelium*, in Symp. Tobacco Research Council, London, (1971).
[105] X.-Y. Wu and Y. Zhang, *A geometrical structure for an infinite oriented cluster and its uniqueness*. Ann. Probab. 36, 862–875 (2008).
[106] Jinjiong Yu, *Edwards-Wilkinson fluctuations in the Howitt-Warren flows*. Stochastic Processes Appl. 126, 948–982 (2016).