

Abelian Sandpiles: an Overview and Results on Certain Transitive Graphs

A.A. Járai

Department of Mathematical Sciences, University of Bath, Claverton Down, Bath BA2 7AY, United Kingdom. E-mail: A.Jarai@bath.ac.uk

Received October 20, 2011

Abstract. We review the Majumdar–Dhar bijection between recurrent states of the Abelian sandpile model and spanning trees. We generalize earlier results of Athreya and Járai on the infinite volume limit of the stationary distribution of the sandpile model on \mathbb{Z}^d , $d \geq 2$, to a large class of graphs. This includes: (i) graphs on which the wired spanning forest is connected and has one end; (ii) transitive graphs with volume growth at least cn^5 on which all bounded harmonic functions are constant. We also extend a result of Maes, Redig and Saada on the stationary distribution of sandpiles on infinite regular trees, to arbitrary exhaustions.

KEYWORDS: Abelian sandpile, burning algorithm, uniform spanning forest, infinite volume limit, transitive graph, bounded harmonic function

AMS SUBJECT CLASSIFICATION: Primary 60K35, Secondary 82B20

1. Introduction

This paper is based on a talk given at an IRS meeting in Paris¹, and contains most of the results discussed in the talk, with proofs. We give an overview of the Abelian sandpile model, with particular emphasis on the Majumdar–Dhar bijection. Then we discuss recent results on the infinite volume limit of the model on certain transitive graphs.

The Abelian sandpile model and close variants were discovered independently in various contexts. Our focus here will be the context of probability models on graphs; see the references in [12] for surprising connections with other

¹Inhomogeneous Random Systems, *Stochastic Geometry and Statistical Mechanics*, Institut Henri Poincaré, Paris, 27 January 2010.

fields of mathematics. “Sandpile” models were introduced by Bak, Tang and Wiesenfeld [2] as simple toy examples, in an attempt to explain the physical mechanisms underlying the widespread occurrence of power-law distributions and fractals in nature. They introduced the idea of self-organized criticality (SOC) as a possible mechanism, and studied “sandpile” models numerically to support their claims. The importance of the model as a theoretical tool to study SOC was recognized by Dhar [5], who generalized it and discovered some of its fundamental properties, including the Abelian property. Dhar coined the name “Abelian sandpile”. The definition of the model is recalled in Section 2, and the key results needed are summarized in Section 3. For further background on the mathematical results, see the survey by Redig [27]. See also the paper by Holroyd, Levine, Mészáros, Peres, Propp and Wilson [12] for a rigorous and self-contained introduction as well as an account of the connection of sandpiles with the rotor-router model. The paper [12] also contains extensions to directed graphs of some of the results discussed in Sections 2 and 3. See the survey by Dhar [6] for the theoretical physics context.

Our main focus will be the following type of question. Let $G = (V, E)$ be an infinite, locally finite graph, for example \mathbb{Z}^d , or the Cayley graph of a finitely generated discrete group. Let $V_1 \subset V_2 \subset \dots \subset V$ be a sequence of finite subsets such that $\cup_{n=1}^{\infty} V_n = V$. Do the Abelian sandpile models on the V_n 's converge to a limiting model on V ?

The above question was first addressed in the case of \mathbb{Z} by Maes, Redig, Saada and Van Moffaert [21]. Here the limiting model has a trivial stationary distribution, nevertheless the question of convergence to this distribution is non-trivial [21].

Maes, Redig and Saada [22] considered sandpiles on infinite regular trees. A stationary Markov process was constructed, obtained as the limit of sandpile Markov chains on finite subgraphs. For the most part, the construction given in [22] is very general, and applies to a general infinite graph. There were two key steps, however, that were specific to the tree: (i) to show that the stationary measures of sandpiles on a suitable sequence of finite subgraphs converge weakly to a unique (automorphism invariant) limit; (ii) to show that avalanches are almost surely finite in the limit. These steps were carried out making use of results of Dhar and Majumdar [7].

Maes, Redig and Saada [23] studied a so-called *dissipative* version of the \mathbb{Z}^d model, where particles are removed on each toppling (not only at the boundary). The presence of dissipation introduces fast decay of correlations. Making use of this, the steps (i)–(ii) above could be carried out, and the infinite volume process was constructed. A nice feature of the limiting process is that it is shown to live on a compact Abelian group, extending the finite volume formalism.

For the usual (non-dissipative) model on \mathbb{Z}^d , the step (i) above for $d \geq 2$ was solved by Athreya and Járai [1], and step (ii) for $d \geq 3$ was solved by Járai and Redig [16]. The main new ingredient in these papers was to exploit a result

of Majumdar and Dhar [25] that gives a bijection between the recurrent states of the sandpile model and wired spanning trees of the underlying graph. This made it possible to use techniques from the theory of uniform spanning trees, in particular Pemantle’s theorem [26] on the existence of the wired uniform spanning forest. There is a difference between the cases $2 \leq d \leq 4$ and $d \geq 5$, that are a reflection of Pemantle’s result that in the former case the spanning forest is a.s. connected, while in the latter case it is not. Another essential ingredient is that the each tree in the spanning forest has one end. This was first proved by Pemantle [26] and Benjamini, Lyons, Peres and Schramm [3]. (In what follows, we will abbreviate the latter authors to BLPS.)

The surveys [24] and [13] discuss the above developments.

More recently, a continuous height dissipative model was studied on \mathbb{Z}^d by Járai, Redig and Saada [17]. This extends the discrete dissipative model considered in [23] by allowing the amount dissipated per toppling to be any non-negative real value, rather than an integer. This has the advantage that the limit of zero dissipation can be formulated precisely. In this limit, the discrete non-dissipative model is recovered. This work is also based on an adaptation of the Majumdar–Dhar bijection.

In the present paper, we extend some of the \mathbb{Z}^d results to more general graphs. Part of our motivation is the well-known fact that Pemantle’s argument [26] (made explicit by Häggström [9]) shows that the wired spanning forest measure exists on any infinite locally finite graph, as a limit from uniform spanning trees on finite graphs (see Theorem 4.1 below). Also, the alternative distinguishing a connected spanning forest from the disconnected case can be vastly generalized [3, Theorems 9.2, 9.4]. At first sight, the Majumdar–Dhar bijection may suggest that a general convergence statement should also exist for sandpiles and could be derived from the bijection. However, the situation is more subtle. Although the wired spanning forest measure is always unique, the limits of sandpile measures may be non-unique. Results of Járai and Lyons [15] (see Theorem 6.2 below) show that this is the case for a class of graphs with two ends, on which the wired spanning forest has two ends a.s. After making appropriate assumptions to exclude the above phenomenon, a general convergence statement can be proved for certain “low-dimensional” graphs. One can follow essentially the same argument as the one made in [1] for \mathbb{Z}^d , $2 \leq d \leq 4$. Nevertheless, we decided to include a new proof in the present paper, that follows a somewhat different route. The argument we present is based on coupling, and hence gives more than weak convergence; see Theorem 6.1 below.

Considerably more work is needed to extend the line of argument made in [1] for \mathbb{Z}^d , $d \geq 5$, to a more general class of transitive graphs. Here we need to make more restrictive assumptions on the graph to make the proof work, see Assumption 7.1 and Theorem 7.2. Nevertheless, parts of our argument for this case are still quite general, and are potentially useful beyond the validity of Assumption 7.1; see Lemmas 7.3 and 7.4 and Proposition 7.1. As discussed in

Section 7, results of BLPS [3] and Lyons, Morris and Schramm [19] imply that there are many graphs on which our assumptions are satisfied.

Our results make it possible to apply the general machinery developed by Maes, Redig and Saada [22] to a large class of graphs. This is outlined in Section 9.

The bijection is also useful on infinite regular trees. We discuss some interesting symmetry properties of the bijection on trees, and in Theorem 8.1 we use them to extend the convergence result of [22] to a general exhaustion.

The outline of the paper is as follows. In Section 2 we review the definition and basic properties of the Abelian sandpile on a finite undirected multigraph. Section 3 is devoted to a discussion of the Majumdar–Dhar bijection that establishes a one-to-one mapping between recurrent sandpile configurations and spanning trees. In Section 4 we recall the wired spanning forest measure and Pemantle’s alternative distinguishing the cases $2 \leq d \leq 4$ and $d \geq 5$ for \mathbb{Z}^d . In Section 5 we recall Wilson’s method and its extensions to infinite graphs. In Section 6, we state and prove the general convergence theorem for “low-dimensional” graphs. In Section 7 we state and prove a convergence theorem for certain “high-dimensional” transitive graphs. In Section 8, we discuss the results for regular trees. Finally, in Section 9 we make the connection with the results of Maes, Redig and Saada [22].

2. The Abelian sandpile model

We first define the model on a finite, connected multigraph $G = (V^+, E)$ that has a distinguished vertex s , called the *sink*. We write $V = V^+ \setminus \{s\}$. We allow G to have loop-edges, as this has no major consequence. We write $a_{xy} = a_{yx}$ for the number of edges between x and y in G , where $x, y \in V^+$. Sometimes we will consider the directed graph $\vec{G} = (V^+, \vec{E})$ that is obtained from G by replacing each edge by two directed edges, one in each direction. A directed edge $e \in \vec{E}$, points from the *tail* of e , denoted e_- , to the *head* of e , denoted e_+ . When there is no ambiguity, we also write $e = [e_-, e_+]$, to specify an oriented edge by its tail and head. For a set of vertices $A \subset V^+$, we denote by $G \setminus A$ the graph obtained by removing all vertices in A from G , as well as all edges incident with vertices in A .

We define the set of *stable configurations* of particles:

$$\Omega_G := \prod_{x \in V} \{0, \dots, \deg_G(x) - 1\},$$

and the set of all particle configurations:

$$\mathcal{X}_G := \prod_{x \in V} \{0, 1, \dots\}.$$

The dynamics of the model is defined in terms of the *toppling matrix*, that is the graph Laplacian:

$$\Delta_{xy} = (\Delta_G)_{xy} = \begin{cases} \deg_G(x) - a_{xx}, & \text{if } x = y; \\ -a_{xy}, & \text{if } x \neq y; \\ 0 & \text{otherwise;} \end{cases} \quad x, y \in V.$$

We define the basic operation of *toppling*. If $\eta \in \mathcal{X}_G$ and $\eta_x \geq \deg_G(x)$, then x is allowed to *topple*, which means that it sends one particle along each edge incident with x . This can be written:

$$\eta_y \longrightarrow \eta_y - \Delta_{xy}, \quad y \in V,$$

that is, the row of Δ_G corresponding to x is subtracted from the configuration η . Note that if x was allowed to topple, the new configuration is also in \mathcal{X}_G . When x is a neighbour of s , a_{xs} particles are lost as the result of toppling, otherwise the number of particles is conserved.

We define the *stabilization map* $\mathcal{S}: \mathcal{X}_G \rightarrow \Omega_G$, by applying topplings as long as possible. It can be shown that any configuration stabilizes in finitely many steps, and the resulting stable configuration is independent of the sequence of topplings used. This is summarized in the following lemma.

Lemma 2.1 (Dhar [5]; see also [12, Lemmas 2.2, 2.4]). *\mathcal{S} is well-defined.*

We define the *addition operators* $a_x: \Omega_G \rightarrow \Omega_G$ by $a_x \eta := \mathcal{S}(\eta + \delta_x)$, $x \in V$, where $\delta_{x,y} = 1$ if $y = x$ and $= 0$ when $y \neq x$. The addition operators satisfy the *Abelian property*: $a_x a_y = a_y a_x$, $x, y \in V$ [5]; see also [12, Lemma 2.5]. Let $\{p(x)\}_{x \in V}$ be a distribution on V satisfying $p(x) > 0$, $x \in V$. We define a Markov chain with statespace Ω_G , where a single step consists of picking a vertex $x \in V$ at random, according to the distribution p , and applying a_x to the configuration. The set of *recurrent configurations* \mathcal{R}_G is the set of recurrent states of the Markov chain. The *Sandpile Group* of G is defined as

$$K_G := \mathbb{Z}^V / \mathbb{Z}^V \Delta_G,$$

that is, \mathbb{Z}^V factored by the integer row span of Δ_G .

Theorem 2.1 (Dhar [5]; see also [12, Corrolary 2.16]).

- (i) *The restriction of the map a_x to \mathcal{R}_G is a one-to-one transformation of \mathcal{R}_G onto itself, for each $x \in V$. These restricted maps generate an Abelian group isomorphic to K_G .*
- (ii) $|\mathcal{R}_G| = |K_G| = \det(\Delta_G)$.

- (iii) *The Markov chain has a unique stationary distribution ν_G and this is the uniform distribution on \mathcal{R}_G .*

By the Matrix-Tree Theorem [4, Theorem II.12], $\det(\Delta_G)$ also equals the number of spanning trees of G . Let us write \mathcal{T}_G for the set of all spanning trees of G . It is natural to ask for an explicit bijection between \mathcal{R}_G and \mathcal{T}_G , and such a bijection is discussed in Section 3. See [12] for a different class of bijections, based on the rotor-router walk.

3. The Majumdar – Dhar bijection

In this section we describe our main tool for studying infinite volume limits of sandpiles. Let $G = (V^+, E)$ be a finite, connected multigraph, and s the distinguished vertex (the sink). Recall that ν_G is the stationary distribution, \mathcal{R}_G is the set of recurrent configurations, and \mathcal{T}_G is the set of spanning trees of G . We describe a bijection between \mathcal{R}_G and \mathcal{T}_G that was introduced by Majumdar and Dhar [25].

3.1. Allowed configurations

For a subset $F \subset V$ and $x \in F$, we write $\deg_F(x) = \sum_{y \in F} a_{yx}$, which is the degree of x in the subgraph induced by F . We write η_F for the restriction of the configuration η to the subset F . We say that η_F is a *forbidden subconfiguration* (FSC) if for all $x \in F$, $\eta_x < \deg_F(x)$. We say that $\eta \in \Omega_G$ is *allowed*, if there is no $F \subset V$, $F \neq \emptyset$, such that η_F is a FSC. Let us write \mathcal{A}_G for the set of allowed configurations. In Section 3.2 we review Dhar’s Burning Algorithm that decides if a given configuration is allowed or not.

It was proved in [5] that $\mathcal{R}_G \subset \mathcal{A}_G$. It was proved in [25], with the introduction of the bijection in Section 3.3, that $|\mathcal{A}_G| = |\mathcal{T}_G|$. Hence it follows that $|\mathcal{R}_G| = \det(\Delta_G) = |\mathcal{T}_G| = |\mathcal{A}_G|$, and therefore $\mathcal{R}_G = \mathcal{A}_G$. See [12, Lemma 4.2] for a different proof of the latter fact, that is still based on the Burning Algorithm, but does not require the bijection.

Lemma 3.1. (Dhar [5]; Majumdar, Dhar [25]; [12, Lemma 4.2].) *Suppose that G is a connected multigraph with a sink s specified. Then $\mathcal{R}_G = \mathcal{A}_G$.*

3.2. The Burning Algorithm

The following algorithm, introduced by Dhar [5], checks if a configuration is allowed. Let $\eta \in \Omega_G$. Set $B(0) := \{s\}$, and $U(0) = V$. For $i \geq 1$, we inductively define

$$\begin{aligned} B(i) &:= \{x \in U(i-1) : \eta_x \geq \deg_{U(i-1)}(x)\}, \\ U(i) &:= U(i-1) \setminus B(i) = V \setminus \left(\cup_{j=0}^i B(j)\right). \end{aligned}$$

We call $B(i)$ the *set of vertices burning at time i* , and $U(i)$ the *set of unburnt vertices at time i* . We say that the algorithm *terminates*, if for some $i \geq 1$ we have $U(i) = \emptyset$. It is easy to prove by induction on i that for all $i \geq 1$, no vertex in $B(i)$ can be part of any FSC. It also follows from the definition of $B(i)$ that if for some $i \geq 1$ we have $B(i) = \emptyset$ and $U(i-1) \neq \emptyset$ (i.e. the algorithm does not terminate), then $\eta_{U(i-1)}$ is an FSC. Hence the algorithm terminates if and only if η is allowed.

This algorithm can be generalized to Eulerian digraphs with a sink; see [12, Lemma 4.1] The algorithm does not work on general directed graphs. An extension to that case, called the *script algorithm*, was given by Speer [28].

3.3. The bijection

Based on the Burning Algorithm, we now give the bijection between \mathcal{A}_G and \mathcal{T}_G . The bijection is not canonical, in the sense that some choices can be made how to set it up. Suppose that for every $x \in V$, every non-empty $P \subset \{e \in \vec{E} : e_- = x\}$ and every finite $K \subset \{0, 1, 2, \dots, \deg_G(x) - 1\}$ of the form $K = \{j, j+1, \dots, j+|P|-1\}$, an arbitrary bijection $\alpha_{P,K}: P \rightarrow K$ is fixed. Then the bijection between \mathcal{A}_G and \mathcal{T}_G will be uniquely defined in terms of the $\alpha_{P,K}$'s.

We define the map $\sigma_G : \mathcal{A}_G \rightarrow \mathcal{T}_G$. Let $\eta \in \mathcal{A}_G$, and consider the sets $\{B(i)\}_{i \geq 0}$ defined in Section 3.2. By the definition of the Burning Algorithm we have $V = \cup_{i \geq 1} B(i)$, and this is a disjoint union. We build the tree $t = \sigma_G(\eta)$ by connecting a vertex $x \in B(i)$, $i \geq 1$ to some vertex in $B(i-1)$. This ensures that there are no loops, and since $V = \cup_{i \geq 1} B(i)$, t will be a spanning tree of G . Suppose then that $x \in B(i)$ for some $i \geq 1$. Let

$$\begin{aligned} n_x &:= \sum_{y \in \cup_{j=0}^{i-1} B(j)} a_{yx} \\ P_x &:= \{e \in \vec{E} : e_+ \in B(i-1), e_- = x\} \\ K_x &= \{\deg_G(x) - n_x, \dots, \deg_G(x) - n_x + |P_x| - 1\}. \end{aligned} \tag{3.1}$$

We claim that $\eta_x \in K_x$. For this, note that due to $x \in B(i)$ we have

$$\begin{aligned} \eta_x &\geq \deg_{U(i-1)}(x) = \sum_{y \in U(i-1)} a_{yx} \\ &= \deg_G(x) - \sum_{y \in \cup_{j=0}^{i-1} B(j)} a_{yx} = \deg_G(x) - n_x. \end{aligned}$$

On the other hand, we have

$$|P_x| = \sum_{y \in B(i-1)} a_{xy} = \sum_{y \in B(i-1)} a_{yx},$$

and since $x \notin B(i-1)$, for $i \geq 2$ we have

$$\eta_x < \deg_{U(i-2)}(x) = \deg_G(x) - \sum_{y \in \cup_{j=0}^{i-2} B(j)} a_{yx} = \deg_G(x) - n_x + |P_x|. \quad (3.2)$$

When $i = 1$, we have $n_x = |P_x|$, so we still have $\eta_x < \deg_G(x) - n_x + |P_x|$. This shows that indeed $\eta_x \in K_x$. It follows that the edge $e_x := \alpha_{P_x, K_x}^{-1}(\eta_x)$ is an edge pointing from x to a vertex in $B(i-1)$. If we define

$$\sigma_G(\eta) := t := \{e_x : x \in V\},$$

then t is a spanning tree of G with each edge directed towards s , or equivalently, disregarding the orientedness, a spanning tree of G .

Lemma 3.2 (Majumdar, Dhar [25]). *The map $\sigma_G : \mathcal{A}_G \rightarrow \mathcal{T}_G$ is a bijection between these sets.*

Proof. We first show that σ_G is injective. Let $\eta^1, \eta^2 \in \mathcal{A}_G$, $\eta^1 \neq \eta^2$, and let $t^1 := \sigma_G(\eta^1)$, $t^2 := \sigma_G(\eta^2)$. Let $i \geq 1$ be the smallest index such that either $B(i, \eta^1) \neq B(i, \eta^2)$ or there exists $x \in B(i, \eta^1) = B(i, \eta^2)$ with $\eta_x^1 \neq \eta_x^2$. If such index did not exist, we would get by induction on i that $\eta^1 = \eta^2$ on $\cup_{i \geq 1} B(i, \eta^1) = \cup_{i \geq 1} B(i, \eta^2) = V$, a contradiction. By the choice of i , we have

$$B(j, \eta^1) = B(j, \eta^2) \quad \text{for } 1 \leq j \leq i-1. \quad (3.3)$$

If $B(i, \eta^1) \neq B(i, \eta^2)$, then pick a vertex x in the symmetric difference. Then by the construction of σ_G , in one of t^1 and t^2 there is an edge from x to $B(i-1, \eta^1) = B(i-1, \eta^2)$ and there is no such edge in the other, so $t^1 \neq t^2$. Suppose therefore that $B(i, \eta^1) = B(i, \eta^2)$, but there exists $x \in B(i, \eta^1) = B(i, \eta^2)$ such that $\eta_x^1 \neq \eta_x^2$. By the equality (3.3), we have $n_x(\eta^1) = n_x(\eta^2)$, $P_x(\eta^1) = P_x(\eta^2)$, and hence also $K_x(\eta^1) = K_x(\eta^2)$. However, since $\eta_x^1 \neq \eta_x^2$ we have $\alpha_{P_x, K_x}^{-1}(\eta_x^1) \neq \alpha_{P_x, K_x}^{-1}(\eta_x^2)$, and therefore the edge between x and $B(i-1)$ is different in t^1 and t^2 . This completes the proof of injectivity.

We now show that σ_G is surjective. In the course of doing so, we find the inverse map $\sigma_G^{-1} =: \phi_G : \mathcal{T}_G \rightarrow \mathcal{A}_G$. First we note that for any $\eta \in \mathcal{A}_G$, the sets $B(0), B(1), \dots$ and the data in (3.1) can be easily expressed in terms of $t = \sigma_G(\eta)$ as well. Namely, let $d_t(\cdot, \cdot)$ denote graph distance in the tree t . Then due to the construction of t , we have

$$\begin{aligned} B(0) &= \{s\}; \\ B(i) &= \{x \in V : d_t(s, x) = i\}, \quad i \geq 1. \end{aligned} \quad (3.4)$$

Since this expresses $B(0), B(1), \dots$ in terms of t , the formulas (3.1) show that n_x, P_x and K_x are also expressed in terms of t . Also, by the definition of σ_G , the unique edge of t in P_x is e_x , hence we have $\eta_x = \alpha_{P_x, K_x}(e_x)$.

The above makes it clear what the inverse $\phi_G = \sigma_G^{-1}$ has to be. Suppose that $t \in \mathcal{T}_G$ is given. We use (3.4) to *define* the $B(i)$'s and for $x \in B_i$, $i \geq 1$, we use (3.1) as the definition of n_x , P_x and K_x . It is immediate from these definitions that P_x is non-empty, and t has a unique edge in P_x . Therefore, for $x \in B(i)$, $i \geq 1$ we let e_x be the unique edge of t in P_x , and we set $\eta_x = \alpha_{P_x, K_x}(e_x)$. We define $\phi_G(t) := \eta$. It is clear that if $\eta \in \mathcal{A}_G$, then $\sigma_G(\phi_G(t)) = t$. What is left to show is that we always have $\eta \in \mathcal{A}_G$.

We prove that for every $t \in \mathcal{T}_G$ we have $\eta = \phi_G(t) \in \mathcal{A}_G$, by applying the Burning Test to η . By definition, $B(0) = \{s\}$. We also set $U(0) = V$, and recursively, $U(i) := U(i-1) \setminus B(i)$ for $i \geq 1$. We show by induction on i that at time $i \geq 0$ precisely $B(i)$ burns.

We know that at time 0, $B(0)$ and $U(0)$ are the set of burning and unburnt sites. Suppose inductively that $i \geq 1$ and we already know that at time $0 \leq j \leq i-1$ exactly $B(j)$ burns, and hence $U(i-1)$ is the set of unburnt sites at time $i-1$. We show that at time i , precisely $B(i)$ burns.

Let $x \in B(i)$. Then due to the inductive hypothesis and the definition of n_x , we have

$$\deg_{U(i-1)}(x) = \sum_{y \in U(i-1)} a_{yx} = \sum_{y \in V^+ \setminus \bigcup_{j=0}^{i-1} B(j)} a_{yx} = \deg_G(x) - n_x. \quad (3.5)$$

Since $\eta_x \in K_x$ (by the definition of $\eta = \phi_G(t)$), we have $\eta_x \geq \deg_G(x) - n_x$. Hence due to (3.5), x burns at time i .

Let now $x \in B(j)$ with $j \geq i+1$. Then by the induction hypothesis, $B(j-1), B(j), \dots$ are unburnt at time $i-1$, and hence

$$\begin{aligned} \deg_{U(i-1)}(x) &\geq \sum_{y \in \bigcup_{k \geq j-1} B(k)} a_{yx} \\ &= \deg_G(x) - \sum_{y \in \bigcup_{0 \leq k \leq j-2} B(k)} a_{yx} = \deg_G(x) - n_x + |P_x|. \end{aligned}$$

Since $\eta_x \in K_x$, we have $\eta_x < \deg_G(x) - n_x + |P_x|$, and therefore x does not burn at time i . This shows that at time i precisely the set $B(i)$ burns, and completes the induction. Therefore η is allowed, and we have shown that σ_G is a bijection between \mathcal{A}_G and \mathcal{T}_G . \square

We define the *uniform spanning tree measure* μ_G as the probability measure on \mathcal{T}_G that assigns each $t \in \mathcal{T}_G$ equal weight. Lemma 3.2 has the following important corollary.

Corollary 3.1. *The stationary measure ν_G of the Abelian sandpile on G is the image under ϕ_G of the uniform spanning tree measure μ_G .*

The next lemma summarizes an observation about the nature of the inverse map ϕ_G that will be important for infinite volume limits. For $t \in \mathcal{T}_G$, write $d_t(\cdot, \cdot)$ for graph distance in the tree t . For $x \in V$, write π_x for the unique self-avoiding path from x to s in t . Let us write $x \sim y$ if there exists an edge in G between x and y . Let

$$\mathcal{N}_x = \{y \in V^+ : y \sim x \text{ or } y = x\}.$$

Let $v_x \in V^+$ be the unique vertex such that $v_x \in \pi_y$ for all $y \in \mathcal{N}_x$, and $d_t(s, v_x)$ is maximal. (Informally, this is the ‘‘first meeting point’’ of the paths $\{\pi_y\}_{y \in \mathcal{N}_x}$.) Let us write \vec{F}_x for the following directed subtree of t :

$$\vec{F}_x := \{e \in t : e_- \in \cup_{y \in \mathcal{N}_x} \pi_y, d_t(s, e_+) \geq d_t(s, v_x)\}.$$

Each edge in \vec{F}_x is directed towards v_x , so specifying \vec{F}_x is equivalent to specifying the undirected, rooted tree (F_x, v_x) . Recall that e_x is the unique edge of t satisfying $(e_x)_- = x$ and $d_t(s, (e_x)_+) = d_t(s, x) - 1$. Write $\eta = \phi_G(t)$.

Lemma 3.3.

- (i) *The value of η_x only depends on t through the rooted subtree (F_x, v_x) .*
- (ii) *The value of η_x only depends on t through the differences $\{d_t(s, x) - d_t(s, y) : y \sim x\}$ and the edge e_x .*
- (iii) *In fact, η_x only depends on the cardinality of the set $\{y \sim x : d_t(s, x) - d_t(s, y) \geq 1\}$, the set $\{y \sim x : d_t(s, x) - d_t(s, y) = 1\}$ and the edge e_x .*

Proof.

(i) By the definition of ϕ_G , η_x only depends on n_x , P_x (which determine K_x) and e_x . Due to the characterization of the $B(i)$'s in terms of graph distance (3.4) we have

$$n_x = \sum_{\substack{y: y \sim x \\ d_t(s, y) < d_t(s, x)}} a_{yx} = \sum_{\substack{y: y \sim x \\ d_t(v_x, y) < d_t(v_x, x)}} a_{yx} = \sum_{\substack{y: y \sim x \\ d_{F_x}(v_x, y) < d_{F_x}(v_x, x)}} a_{yx},$$

and the last expression only depends on (F_x, v_x) . Similarly,

$$\begin{aligned} P_x &= \{e \in \vec{E} : e_- = x, d_t(s, e_+) = d_t(s, x) - 1\} \\ &= \{e \in \vec{E} : e_- = x, d_{F_x}(v_x, e_+) = d_{F_x}(v_x, x) - 1\}, \end{aligned}$$

and the last expression only depends on (F_x, v_x) . Finally, since e_x is the unique edge of t incident with x that is directed away from x , we have

$$e_x \text{ is the unique edge } e \text{ in } \vec{F}_x \text{ such that } e_- = x \text{ and } d_{F_x}(v_x, e_+) = d_{F_x}(v_x, x) - 1.$$

(ii) This is similar to part (i). We have

$$n_x = \sum_{\substack{y: y \sim x \\ d_t(s,x) - d_t(s,y) > 0}} a_{yx},$$

$$P_x = \{e \in \vec{E} : e_- = x, d_t(s,x) - d_t(s,e_+) = 1\}.$$

This proves the claim. (iii) also follows from the above expressions. \square

4. The Wired Spanning Forest

Let now $G = (V, E)$ be an infinite locally finite graph. For simplicity, from now on we restrict our attention to simple graphs (no multiple edges or loops), but note that it is possible to extend all our results in Sections 6 and 7 to multigraphs, with essentially the same arguments.

An *exhaustion* of V is a sequence $V_1 \subset V_2 \subset \dots \subset V$ such that $\cup_{n=1}^{\infty} V_n = V$. Let $G_n = (V_n^+, E_n)$ denote the graph obtained from G by identifying all vertices in $V \setminus V_n$ to a single vertex s , so that $V_n^+ = V_n \cup \{s\}$, and removing loops at s . Sometimes G_n is called the *wired graph* associated to V_n , where “wired” refers to the fact that all connections outside V_n have been forced to occur. Recall that μ_{G_n} is the uniform probability measure on the set of spanning trees \mathcal{T}_{G_n} . We will write \Rightarrow to denote weak convergence of measures.

The usefulness for infinite volume limits of the bijection in Section 3.3 lies in the well-known theorem stated below. This theorem is implicit in the work of Pemantle [26], and was made explicit by Häggström [9], in the case of \mathbb{Z}^d . The \mathbb{Z}^d proof immediately applies in the generality stated.

Theorem 4.1 (Pemantle [26]; see also [9]). *Let $G = (V, E)$ be an infinite locally finite graph. There exists a measure μ on $\{0, 1\}^E$ such that $\mu_{G_n} \Rightarrow \mu$ independently of the exhaustion. The measure μ concentrates on spanning forests of G all of whose components are infinite.*

The measure μ is also called the *Wired Spanning Forest (WSF) measure*. Theorem 4.1 naturally leads to the following question.

Open question 1. Assume the same conditions as in Theorem 4.1. Under what extra conditions does ν_{G_n} have a unique weak limit ν on the space $\prod_{x \in V} \{0, \dots, \deg_G(x) - 1\}$, independently of the exhaustion?

It is not possible to deduce a general convergence statement only from Theorem 4.1. On certain graphs with two ends the limit is not unique; see Theorem 6.2 in Section 6. However, as we will see in Section 6, there is a general convergence theorem on certain “low-dimensional” graphs. We will need to consider the number of components of the WSF, and the *ends* of the components.

We say that an infinite tree *has one end*, if any two infinite self-avoiding paths in the tree have infinitely many vertices in common. In the theorem below, statement (i) and the first part of statement (ii) are due to Pemantle [26]. The statement on one end in part (ii) was first proved by BLPS [3] and in much greater generality. Lyons, Morris and Schramm [19] gave a simpler and even more general proof with quantitative estimates.

Theorem 4.2 (Pemantle [26]; BLPS [3]). *Let G be the \mathbb{Z}^d lattice.*

- (i) *Suppose $2 \leq d \leq 4$. The Wired Spanning Forest is μ -a.s. connected, and has one end.*
- (ii) *Suppose $d \geq 5$. The Wired Spanning Forest μ -a.s. consists of infinitely many trees, and each tree has one end.*

5. Wilson's method

In this section we recall some facts about Wilson's method, that is an indispensable tool in studying uniform spanning trees.

Let $\pi = [\pi_0, \pi_1, \dots, \pi_M]$ be a finite path in some graph. The *loop-erasure* of π is defined by chronologically removing loops from the path as they are created. That is, we set $\sigma = \text{LE}(\pi) := [\sigma_0, \dots, \sigma_K]$, where we inductively define

$$\begin{aligned} s_0 &:= 0, & s_j &:= \max\{k \geq s_{j-1} : \pi(k) = \sigma(j-1)\}, & j &\geq 1, \\ \sigma(0) &:= \pi(0), & \sigma(j) &:= \pi(s_j + 1), & j &\geq 1. \end{aligned}$$

Note that loop-erasure also makes sense for an infinite path that visits any vertex only finitely often.

Suppose now that $G = (V, E)$ is a finite graph, and $w : E \rightarrow (0, \infty)$ is a function. We call $w(e)$ the *weight* of the edge e . The pair (G, w) is called a *network*. Most of the time no weights will be specified, and then it is assumed that $w(e) = 1$ for all $e \in E$. The weight of a spanning tree $t \in \mathcal{T}_G$ is defined by $w(t) := \prod_{e \in t} w(e)$. We extend the definition of μ_G to networks by requiring that each element of \mathcal{T}_G receives probability proportional to its weight.

A *network random walk* on (G, w) is a Markov chain $\{S(n)\}_{n \geq 0}$ with state space V and transition probabilities:

$$\mathbf{P}[S(k+1) = v \mid S(k) = u] = \frac{w(u, v)}{\sum_{v' \sim u} w(u, v')}.$$

When the weights are constant, we call this *simple random walk* on G . The definition of network random walk immediately extends to infinite networks as long as for each vertex $u \in V$ we have $\sum_{v' \sim u} w(u, v') < \infty$.

Let v_1, \dots, v_N be an enumeration of V , and let r be a fixed vertex of G . Let $\{S_k^j\}_{k \geq 0}$, $1 \leq j \leq N$ be independent network random walks on G , with

$S^j(0) = v_j$. We define a sequence of subtrees $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \dots \subset \mathcal{F}_N$ of G . Put $\mathcal{F}_0 = \{r\}$, and inductively define for $j \geq 1$:

$$\begin{aligned} T_j &:= \inf\{k \geq 0 : S^j(k) \in \mathcal{F}_{j-1}\}, \\ \mathcal{F}_j &:= \mathcal{F}_{j-1} \cup \text{LE}(S^j[0, T^j]). \end{aligned} \quad (5.1)$$

It is clear from the construction that \mathcal{F}_N is a spanning tree of G .

Theorem 5.1 (Wilson [30]). *On any finite network, regardless of what enumeration was chosen, \mathcal{F}_N is distributed according to μ_G .*

Suppose now that $G = (V, E)$ is a locally finite infinite recurrent graph. Essentially the same method can be applied as in the finite case. Let v_1, v_2, \dots be an enumeration of V , and let $r \in V$ be fixed. Define \mathcal{F}_j , $j \geq 0$ as in the finite case, and set $\mathcal{F} := \cup_{j \geq 0} \mathcal{F}_j$. Then \mathcal{F} is a.s. a spanning tree of G .

Theorem 5.2 (BLPS [3, Theorem 5.6]). *On any recurrent infinite graph, regardless of the enumeration chosen, \mathcal{F} is distributed according to μ .*

Suppose now that $G = (V, E)$ is a locally finite infinite transient graph. Wilson's method can be applied to this case as well, by letting the root r be "at infinity". That is, let v_1, v_2, \dots be an enumeration of V , set $\mathcal{F}_0 := \emptyset$, and define T_j and \mathcal{F}_j as in (5.1). Now some of the T_j 's will be infinite, but as noted earlier, loop-erasure still makes sense due to transience. We set $\mathcal{F} := \cup_{j \geq 1} \mathcal{F}_j$. Then \mathcal{F} is a.s. a spanning forest of G .

Theorem 5.3 (BLPS [3, Theorem 5.1]). *On any transient infinite graph, regardless of the enumeration chosen, \mathcal{F} is distributed according to μ .*

6. Infinite volume limits – single tree

Let $G = (V, E)$ be an infinite locally finite graph as in Section 4. Let $V_1 \subset V_2 \subset \dots \subset V$ be an exhaustion, and recall the wired graph $G_n = (V_n^+, E_n)$. In this section we assume that the WSF of G is μ -a.s. connected and has one end. By Theorem 4.2(i), this includes \mathbb{Z}^d with $2 \leq d \leq 4$. The theorem below was proved in [1]. There it was stated in the case of \mathbb{Z}^d , $2 \leq d \leq 4$, however, the proof there directly applies to the more general setting. Nevertheless, below we present a somewhat different proof, based on coupling. Let us write

$$\Omega_G := \prod_{x \in V} \{0, \dots, \deg_G(x) - 1\}.$$

Theorem 6.1 (Athreya, Járai [1]). *Let $G = (V, E)$ be an infinite locally finite graph. Suppose that the WSF of G is μ -a.s. connected and has one end. There exists a measure ν on Ω_G such that $\nu_{G_n} \Rightarrow \nu$, independently of the exhaustion.*

Before proving Theorem 6.1, let us comment on when the assumptions are satisfied. If $G = (V, E)$ is any graph, an *automorphism* of G is a bijection $\varphi : V \rightarrow V$, such that $\{x, y\} \in E$ if and only if $\{\varphi(x), \varphi(y)\} \in E$. We say that G is *vertex-transitive*, if for any $x, y \in V$ there exists an automorphism that takes x to y .

Suppose that $G = (V, E)$ is a locally finite (vertex)-transitive graph. Let $o \in V$ be a fixed vertex of G , and let v_n be the number of vertices of G with distance at most n from o . It was shown by Lyons, Peres and Schramm [20, Corollary 5.3] and BLPS [3, Corollary 9.6] that if $v_n \leq cn^4$, then the WSF is a.s. connected. Regarding the number of ends, it was shown by BLPS [3, Theorem 10.3] that if G is transitive and transient and the WSF has a single tree a.s., then that tree has one end a.s. This was further generalized by Lyons, Morris and Schramm [19, Theorem 7.1] who gave a sufficient condition in terms of the isoperimetric profile of the graph, without assuming transitivity. Regarding the recurrent case, it was shown in [3, Theorem 10.6, Proposition 10.10] that in a recurrent transitive graph G , the WSF has one end a.s. unless G is roughly isometric to \mathbb{Z} .

Proof of Theorem 6.1. Fix $x \in V$. We will use a subscript n for objects associated with the graph G_n . In particular, we write $F_{n,x}$, $v_{n,x}$, etc. for the data associated to a $t_n \in \mathcal{T}_{G_n}$ appearing in Lemma 3.3.

Write \mathcal{T}_G for the set of spanning trees of G with one end, and let $t \in \mathcal{T}_G$. Due to the one end property, we can think of each edge of t being directed towards infinity. For $u, v \in V$ let us write $u \preceq v$ if there is a directed path (possibly of length 0) from u to v in t , and write $u \prec v$ in the case when $u \neq v$. For $y \in \mathcal{N}_x$ let π_y denote the unique infinite directed path in t starting at y . Let $v_x \in V$ be the unique vertex such that $v_x \in \pi_y$ for all $y \in \mathcal{N}_x$, and v_x is minimal with respect to the relation \preceq (the ‘‘first meeting point’’). Let us write \vec{F}_x for the following directed subtree of t :

$$\vec{F}_x := \{e \in t : e_- \in \cup_{y \in \mathcal{N}_x} \pi_y, e_+ \preceq v_x\}.$$

Each edge in \vec{F}_x is directed towards v_x , so specifying \vec{F}_x is equivalent to specifying the undirected, rooted tree (F_x, v_x) .

We now define a mapping $\phi_G : \mathcal{T}_G \rightarrow \Omega_G$. Let

$$\begin{aligned} n_x &= |\{y : y \sim x, d_{F_x}(v_x, y) < d_{F_x}(v_x, x)\}|, \\ P_x &= \{e \in \vec{E} : e_- = x, d_{F_x}(v_x, e_+) = d_{F_x}(v_x, x) - 1\}, \\ K_x &= \{\deg_G(x) - n_x, \dots, \deg_G(x) - n_x + |P_x| - 1\}. \end{aligned}$$

Let e_x be the unique edge of t satisfying $(e_x)_- = x$. Set $\eta_x := \alpha_{P_x, K_x}(e_x)$, $x \in V$. Define $\phi_G(t) := \eta$, and let ν be the image of μ under the map ϕ_G .

We show that $\nu_n \Rightarrow \nu$. In fact, we consider a coupling between the measures μ_n and μ , with the following property. With $\eta_n = \phi_{G_n}(t_n)$ and $\eta = \phi_G(t)$,

for all finite $A \subset V$ we have

$$\lim_{n \rightarrow \infty} \mathbf{P}[\eta_{n,x} = \eta_x, x \in A] = 1. \quad (6.1)$$

This clearly implies weak convergence.

We first consider the case when G is recurrent. For any $x \in V$ let

$$D_x := \{e \in E : \exists u \in V \text{ incident with } e \text{ such that } u \preceq v_x\}.$$

Due to the assumption on one end, D_x is μ -a.s. finite.

Lemma 6.1. *Let $K \subset E$ be a fixed finite set of edges. On the event $D_x \subset K$, the value of (F_x, v_x) , and hence of $\eta_x = (\phi_G(t))_x$ is determined by the status of the edges in K , that is by the pair $(K \cap t, K \setminus t)$. Similarly, on the event $D_{n,x} \subset K$, the value of $(F_{n,x}, v_{n,x})$, and hence the value of $\eta_{n,x} = (\phi_{G_n}(t_n))_x$ is determined by the status of the edges in K .*

Proof of Lemma 6.1. All edges of F_x belong to $D_x \cap t$ and hence to $K \cap t$. Therefore, F_x is determined as the smallest connected set of edges in $K \cap t$ containing all vertices of \mathcal{N}_x . It remains to show that v_x is also determined.

We claim that v_x is the unique vertex belonging to F_x such that there exists a path in $K \cap t$ from v_x to the vertex boundary of K that is edge-disjoint from F_x . First note that v_x satisfies the requirement, by virtue of the path in t from v_x to infinity. Suppose $v \neq v_x$ was another such vertex, and let $f_1, \dots, f_L \in K \cap t$ be a path from v to the vertex boundary of K that is disjoint from F_x . By the definition of D_x and induction, we have $f_j \in D_x \cap t$, $j = 1, \dots, L$. Let $f \notin K$ be an edge that shares an endvertex with f_L . The common endvertex of f_L and f , call it u , satisfies $u \prec v_x$. Hence we get $f \in D_x \subset K$, a contradiction. \square

We continue the proof of (6.1) (in the case when G is recurrent). Fix $\varepsilon > 0$. Choose $B \subset E$ a large enough finite set, so that

$$\mathbf{P}[\cup_{x \in A} D_x \subset B] > 1 - \varepsilon. \quad (6.2)$$

Assume n is large enough so that $E_n \supset B$. The following coupling between μ_n and μ is due to BLPS [3, Proposition 5.6]. Let u_1, \dots, u_K be an enumeration of all vertices incident with the edges in B . We use Wilson's method to generate samples t_n (resp. t) from μ_n (resp. μ), where the enumeration of vertices starts with u_1, \dots, u_K , and the root is some fixed vertex $r \in A$. The same random walks are used in the case of G_n and G , up to the first time τ_n^j when the walk crosses an edge between V_n and the sink s . After time τ_n^j , the construction on G_n is continued using an independent simple random walk on G_n started at s . Due to recurrence, for large enough n ,

$$\mathbf{P}[\tau_n^j > T^j \text{ for } j = 1, \dots, K] > 1 - \varepsilon. \quad (6.3)$$

If the event in (6.3) occurs, the status of all edges in B are the same for t_n and t . When the event in (6.2) also occurs, Lemma 6.1, Lemma 3.3(i), and the definitions of ϕ_G and ϕ_{G_n} imply that $(\phi_G(t))_x = (\phi_{G_n}(t_n))_x$ for all $x \in A$. Since ε was arbitrary, this proves (6.1) in the recurrent case.

When G is transient, the proof is fairly similar, and somewhat simpler. This time, we let u_1, \dots, u_K be an enumeration of $\cup_{x \in A} \mathcal{N}_x$. On G we use Wilson's method rooted at infinity, and on G_n we use it with root equal to the sink s . The constructions use the same random walks S^j , up to the first exit time τ_n^j from V_n for $j = 1, \dots, K$. Given $\varepsilon > 0$, let $C \subset E$ be a large enough finite set so that

$$\mathbf{P}[\cup_{x \in A} F_x \subset C] > 1 - \varepsilon. \quad (6.4)$$

Let $\hat{\tau}_C^j$ be the time of the last visit to C by S^j . Due to transience, $\hat{\tau}_C^j < \infty$ a.s. It follows, using transience again, that if n is large enough

$$\mathbf{P}[S^j[\tau_n^j, \infty) \cap S^j[0, \hat{\tau}_C^j] = \emptyset, j = 1, \dots, K] > 1 - \varepsilon. \quad (6.5)$$

Note that on the event in (6.5), using the notation from Section 5, we have $\mathcal{F}_K \cap C = \mathcal{F}_K^n \cap C$. When the event in (6.4) also occurs, we have $(F_x, v_x) = (F_{n,x}, v_{n,x})$, $x \in A$. This implies, due to Lemma 3.3(i) and the definitions of $\phi_G(t)$ and $\phi_{G_n}(t_n)$, that $\eta_x = \eta_{n,x}$, $x \in A$. Since ε was arbitrary, we obtain (6.1) in the transient case. \square

Uniqueness of the limit in Theorem 6.1 can fail, if the assumption on one end is dropped. The following theorem was proved in [15]. Let G_0 be a connected finite graph. Let G be the product $\mathbb{Z} \times G_0$, that is, (n_1, u_1) and (n_2, u_2) are connected by an edge, if either $n_1 = n_2$ and $u_1 \sim u_2$ in G_0 , or if $u_1 = u_2$ and $|n_1 - n_2| = 1$. Write $G_{n,m}$ for the wired graph associated to $\{n, n+1, \dots, m-1, m\} \times G_0$.

Theorem 6.2 (Járai, Lyons [15]). *If G_0 has at least two vertices, then $\{\nu_{G_{n,m}} : n < 0, m > 0\}$ has precisely two ergodic weak limit points.*

Remark 6.1. Here the WSF on G has two ends a.s., as can be seen by using Wilson's method, Theorem 5.2. Therefore, the conditions of Theorem 6.1 are not satisfied. It is a natural question whether this is the only thing that can go wrong with the existence of a unique limit. If the answer is yes, this would solve Open question 1.

7. Infinite volume limits — multiple trees

7.1. Statement of result

In this section we will be interested in graphs where the WSF is not a single tree. Theorem 4.2(ii) states that this is the case when G is the \mathbb{Z}^d lattice

for $d \geq 5$. The method of proof of Theorem 6.1 breaks down in this case, because with probability bounded away from 0, $v_{n,x}$ equals the sink, and hence $\{(F_{n,x}, v_{n,x})\}_{n \geq 1}$ is not tight. The following theorem was proved in [1] in the case when the exhaustion satisfies a regularity property. The restriction on the exhaustion was removed in [16, Appendix], using the result of [14].

Theorem 7.1 (Athreya and Járai [1]; Járai and Redig [16]). *Consider the \mathbb{Z}^d lattice with $d \geq 5$, and let $V_1 \subset V_2 \subset \dots \subset \mathbb{Z}^d$ be any exhaustion. There exists a measure ν on $\Omega_{\mathbb{Z}^d}$ such that $\nu_n \Rightarrow \nu$, independently of the exhaustion.*

The goal of this section is to generalize Theorem 7.1 to other graphs under certain conditions.

Let $G = (V, E)$ be an infinite locally finite graph. We denote by $\text{AUT}(G)$ the group of graph automorphisms of G . With the topology of pointwise convergence, $\text{AUT}(G)$ is a locally compact group [31, Lemma 1.27].

A function $h : V \rightarrow \mathbb{R}$ is called *harmonic*, if for every $x \in V$ we have

$$\frac{1}{\deg_G(x)} \sum_{y: y \sim x} h(y) = h(x).$$

We make the following assumptions on G .

Assumption 7.1.

- (i) G is vertex-transitive.
- (ii) The probability that two independent simple random walks on G started at some vertex intersect infinitely often is 0.
- (iii) Each component of the WSF of G has one end a.s.
- (iv) Every bounded harmonic function on G is constant.

We are going to prove the following theorem.

Theorem 7.2. *Let $G = (V, E)$ be an infinite, locally finite graph, satisfying Assumption 7.1 (i)–(iv). There exists a measure ν on Ω_G such that for any exhaustion $V_1 \subset V_2 \subset \dots \subset V$ we have $\nu_{G_n} \Rightarrow \nu$.*

Before setting out to prove Theorem 7.2, let us discuss examples where the conditions are satisfied.

Condition (i). Suppose that Γ is a finitely generated group, and let S be a fixed finite generating set with the property that if $s \in S$ then also $s^{-1} \in S$. The (right-)Cayley graph of (Γ, S) is the graph with vertex set $V = \Gamma$ and edge

set $E := \{\{x, xs\} : x \in \Gamma, s \in S\}$. Any Cayley graph is vertex-transitive, as shown by left-multiplication by elements of Γ .

Condition (ii). Suppose that G is a vertex-transitive graph, and let o be a fixed vertex of G . Write $d(\cdot, \cdot)$ for graph distance in G . Let

$$V(n) := |\{x \in V : d(o, x) \leq n\}|.$$

Suppose that there exists a constant $c > 0$ such that $V(n) \geq cn^5$. Let $\{S_n\}_{n \geq 0}$ be simple random walk on G . Due to [31, Corollary 14.5], the return probability of S satisfies $\mathbf{P}[S_{2n} = o \mid S_0 = o] \leq Cn^{-5/2}$. As explained in [20, Section 5], this implies that the expected number of intersections (with multiplicity) between two independent simple random walks starting at o is finite. Hence (ii) is satisfied in this case. Note that by [3, Theorem 9.4], the WSF has infinitely many trees a.s., whenever (i) and (ii) are satisfied.

Condition (iii). Suppose that G is a vertex-transitive graph satisfying $V(n) \geq cn^3$. It follows from results of Lyons, Morris and Schramm [19, Theorem 7.1, Corollary 7.3], that every tree of the WSF has one end a.s. In the cases when the WSF is a single tree, and when the WSF is disconnected with $\text{AUT}(G)$ unimodular, this was earlier proved by BLPS [3, Theorems 10.3, 10.4], Hence (iii) is satisfied for a large class of graphs.

Condition (iv). Let $G = (V, E)$ be a graph on which the group $\Gamma \subseteq \text{AUT}(G)$ acts transitively, i.e., for any $x, y \in V$ there exists $\varphi \in \Gamma$ such that $\varphi(x) = y$. Examples where Assumption 7.1 (iv) is satisfied are given by *nilpotent* groups Γ . Recall that for $a, b \in \Gamma$, their commutator is defined as $[a, b] := a^{-1}b^{-1}ab$. Let $\Gamma_1 := \Gamma$, and for $k \geq 2$ let Γ_k be the subgroup of Γ generated by all elements of the form $\dots [[a_1, a_2], a_3], \dots, a_k$. Then $\Gamma = \Gamma_1 \supseteq \Gamma_2 \supseteq \Gamma_3 \supseteq \dots$ is called the *lower central series* of Γ . If there exists an r such that Γ_{r+1} is the trivial group, Γ is called *nilpotent* [11, Chapter 10]. It was shown in [8] that if Γ is nilpotent then any bounded harmonic function on G is constant.

Remark 7.1. We believe that the technical Assumption 7.1(iv) is not necessary. However, at present the only example where we know the existence of ν without this assumption is the case of a regular tree, discussed in Section 8.

7.2. Notation and coupling

We prepare for the proof of Theorem 7.2 by defining the appropriate analogue of (F_x, v_x) . This is done in the same way as for the case of \mathbb{Z}^d , $d \geq 5$ in [1]. In order to be self-contained, we give the details. The idea behind the definitions is that there is probability bounded away from zero, as $n \rightarrow \infty$, that two given vertices $y_1, y_2 \in \mathcal{N}_x$ will be connected through the sink s . This means that $(F_{n,x}, v_{n,x})_{n \geq 1}$ is not tight. We want to replace it with an object that is tight, by removing the connections through s .

We first give the finite volume definitions. We use notation similar to Section 6, that is, lower indices n refer to the graph G_n . Fix $x \in V$, and assume that $\mathcal{N}_x \subset V_n$. Let $t_n \in \mathcal{T}_{G_n}$, and define the forest $\mathbf{t}_n := t_n \setminus \{s\}$. Let

$$\begin{aligned} K_{n,x}(\mathbf{t}_n) &:= \text{number of connected components of } \mathbf{t}_n \text{ intersecting } \mathcal{N}_x, \\ t_{n,x}^{(1)}, \dots, t_{n,x}^{(K_{n,x})} &:= \text{the components of } \mathbf{t}_n \text{ that intersect } \mathcal{N}_x, \\ A_{n,x}^{(i)} &:= t_{n,x}^{(i)} \cap \mathcal{N}_x, \quad 1 \leq i \leq K_{n,x}. \end{aligned} \quad (7.1)$$

Here the indexing of the $t^{(i)}$'s and the $A^{(i)}$'s is determined as follows. We fix an ordering of \mathcal{N}_x , let us say $\mathcal{N}_x = \{y_0 = x, y_1, \dots, y_{\deg_G(x)}\}$. We let $t_{n,x}^{(1)}$ be the component of \mathbf{t}_n containing y_0 , let $t_{n,x}^{(2)}$ be the component containing the earliest y_i not in $t_{n,x}^{(1)}$, etc.

Let $v_{n,x}^{(i)}$ be the unique vertex v of $t_{n,x}^{(i)}$ such that $v \in \pi_{n,y}$ for all $y \in A_{n,x}^{(i)}$, and $d_t(s, v)$ is maximal. We define

$$\vec{F}_{n,x}^{(i)} := \left\{ e \in t_{n,x}^{(i)} : e_- \in \bigcup_{y \in A_{n,x}^{(i)}} \pi_{n,y}, d_t(s, e_+) \geq d_t(s, v_{n,x}^{(i)}) \right\}. \quad (7.2)$$

Specifying the $\vec{F}_{n,x}^{(i)}$'s is equivalent to specifying the undirected rooted trees $(F_{n,x}^{(i)}, v_{n,x}^{(i)})_{i=1}^{K_{n,x}}$.

We introduce the set of relative distances:

$$d_{n,x}^{(i,j)} = d_{n,x}^{i,j}(t_n) = d_n(v_{n,x}^{(i)}, s) - d_n(v_{n,x}^{(j)}, s), \quad 1 \leq i < j \leq K_{n,x}.$$

Due to Lemma 3.3 (ii), $\eta_{n,x} = (\phi_n(t_n))_x$ only depends on the data:

$$K_{n,x}(t_n), \quad (F_{n,x}^{(i)}, v_{n,x}^{(i)})_{i=1}^{K_{n,x}}, \quad \{d_{n,x}^{(i,j)}(t_n)\}_{1 \leq i < j \leq K_{n,x}}.$$

When each tree in the WSF on G has one end, we can expect that the joint law of

$$K_{n,x}, \quad (F_{n,x}^{(i)}, v_{n,x}^{(i)})_{i=1}^{K_{n,x}}$$

converges as $n \rightarrow \infty$. The candidate for the limit is given by the natural analogues in the graph G , that we now define.

Let $\mathcal{T}_G \subset \{0, 1\}^E$ denote the set of all spanning forests of G such that each component is infinite and has one end. Let $\mathbf{t} \in \mathcal{T}_G$. Due to the one end property, we can direct each edge of \mathbf{t} towards the end of the component containing it. Again, we write $u \preceq v$, if there is a directed path from u to v . As in Section 6, for $y \in V$ we denote by π_y the unique infinite directed path in \mathbf{t} starting at y . Fix $x \in V$, and let

$$\begin{aligned} K_x(\mathbf{t}) &:= \text{number of connected components of } \mathbf{t} \text{ intersecting } \mathcal{N}_x, \\ t_x^{(1)}, \dots, t_x^{(K_x)} &:= \text{the components of } \mathbf{t} \text{ that intersect } \mathcal{N}_x, \\ A_x^{(i)} &:= t_x^{(i)} \cap \mathcal{N}_x, \quad 1 \leq i \leq K_x. \end{aligned} \quad (7.3)$$

Here the indexing of the $t^{(i)}$'s and $A^{(i)}$'s follows the same rule as in the case of G_n . Let $v_x^{(i)}$ be a vertex of $t_x^{(i)}$ minimal with respect to the relation \preceq among all vertices v with the property that $y \preceq v$ for all $y \in A_x^{(i)}$. Such a vertex exists, due to the one end property, and there is a unique minimal one. Let

$$\vec{F}_x^{(i)} := \left\{ e \in t_x^{(i)} : e_- \in \bigcup_{y \in A_x^{(i)}} \pi_y, e_+ \preceq v_x^{(i)} \right\}, \quad 1 \leq i \leq K_x.$$

Specifying the $\vec{F}_x^{(i)}$'s is equivalent to specifying the undirected, rooted trees $(F_x^{(i)}, v_x^{(i)})$.

Lemma 7.1. *Suppose that $G = (V, E)$ is a transient graph that satisfies Assumption 7.1 (ii)–(iii). For any finite $A \subset V$ there is a coupling of μ_n , $n \geq 1$, and μ such that in this coupling*

$$\lim_{n \rightarrow \infty} \mathbf{P}[K_{n,x} = K_x, (F_{n,x}^{(i)}, v_{n,x}^{(i)}) = (F_x^{(i)}, v_x^{(i)}), 1 \leq i \leq K_{n,x}] = 1. \quad (7.4)$$

Proof. Let u_1, \dots, u_L be an enumeration of $\cup_{x \in A} \mathcal{N}_x$. On G we use Wilson's method rooted at infinity with random walks S^j , started at u_j for $j = 1, \dots, L$. On G_n , we use Wilson's method with root equal to the sink, and with the same random walks S^j , up to their first exit time τ_n^j from V_n . Recall the notation from Section 5: $(\mathcal{F}_i)_{i \geq 0}$ and $(\mathcal{F}_{n,i})_{i \geq 0}$ are the growing forests constructed by Wilson's method, and T^j is the hitting time of \mathcal{F}_{j-1} by S^j . For any $C \subset E$, let

$$\hat{\tau}_C^j := \sup\{k \geq 0 : S^j(k) \in C\}.$$

Let $J \subset \{1, \dots, L\}$ be the (random) set of indices such that $T^j = \infty$.

Given $\varepsilon > 0$, let $C \subset E$ be a large enough finite set such that

$$\mathbf{P}\left[\bigcup_{j \notin J} S^j[0, T^j] \subset C\right] > 1 - \varepsilon. \quad (7.5)$$

By transience, we can find n_1 , such that for all $n \geq n_1$ we have

$$\mathbf{P}[\text{for all } j \in J \text{ we have } S^j[\tau_n^j, \infty) \cap S^j[0, \hat{\tau}_C^j] = \emptyset] > 1 - \varepsilon. \quad (7.6)$$

The significance of the event in (7.6) is that on this event, the loop-erasing procedure on $S^j[0, \infty)$ after time τ_n^j has no effect on the configuration in C , so the configurations in C will be the same when the algorithm is run in G_n and G .

Observe that if $i, j \in J$, $i < j$, then $\text{LE}(S^i[0, \infty)) \cap S^j[0, \infty) = \emptyset$. Assumption 7.1 (ii) and transitivity implies that almost surely $|S^j[0, \infty) \cap S^i[0, \infty)| < \infty$. Since the points in this intersection are not present in $\text{LE}(S^i[0, \infty))$, we can find n_2 large enough such that for all $n \geq n_2$ we have

$$\mathbf{P}[\text{for all } i, j \in J, i < j \text{ we have } \text{LE}(S^i[0, \tau_n^i]) \cap S^j[0, \tau_n^j] = \emptyset] > 1 - \varepsilon. \quad (7.7)$$

Assume now the intersection of the events in (7.5), (7.6) and (7.7). Let $n \geq \max\{n_1, n_2\}$. We prove that the event in (7.4) then must occur, implying the lemma. We show that for all $1 \leq i \leq L$ the following holds:

- (i) if $T^i = \infty$ then $T_n^i = \tau_n^i$ and $\text{LE}(S^i[0, \infty))$ and $\text{LE}(S^i[0, \tau_n^i])$ agree up to their last visit to C ;
- (ii) if $T^i < \infty$ then $T^i = T_n^i < \tau_n^i$ and $\text{LE}(S^i[0, T^i]) = \text{LE}(S^i[0, T_n^i]) \subset C$;
- (iii) $\mathcal{F}_{n,i} \cap C = \mathcal{F}_i \cap C$.

The proof is by induction on i . For $i = 1$ we have $T^1 = \infty$ and $T_n^1 = \tau_n^1$ always, so we are in case (i). Let γ be the initial segment of $\text{LE}(S^1[0, \tau_n^1])$ up to the last visit to C . Due to the event in (7.6), S^1 makes no further visit to γ after time τ_n^1 , and therefore the initial segment of $\text{LE}(S^1[0, \infty))$ up to the last exit from C coincides with γ , as required.

Consider now $2 \leq i \leq L$. We prove (i). The event in (7.7) implies that $S^i[0, \tau_n^i)$ does not intersect any of the paths $\text{LE}(S^j[0, \tau_n^j])$ with $j < i$, $j \in J$. It also does not intersect $\text{LE}(S^j[0, T_n^j])$ for $j < i$, $j \notin J$, since by the induction hypothesis for j , case (ii), we have $\text{LE}(S^j[0, T_n^j]) = \text{LE}(S^j[0, T^j])$ and $T^j = \infty$.

We now prove (ii). By virtue of the event in (7.5), $S^i[0, T^i]$ does not leave C . By the induction hypothesis (iii) we have $\mathcal{F}_{n,i-1} \cap C = \mathcal{F}_{i-1} \cap C$, and the claim in (ii) follows immediately.

Statement (iii) follows from (i) and (ii).

It follows immediately from (i) and (ii) that that $K_{n,x} = K_x$. It also follows from (i) and (ii) that for each $x \in A$ and $1 \leq i \leq K_x$ we have $(F_{n,x}^{(i)}, v_{n,x}^{(i)}) = (F_x^{(i)}, v_x^{(i)}) \subset C$. This completes the proof. \square

7.3. Permutation of components

As in [1], the key difficulty to overcome is to analyze the behaviour of the $d_{n,x}^{(i,j)}$'s. Lemma 3.3(ii) implies that when $|d_{n,x}^{(i,j)}|$ is large for all $1 \leq i < j \leq K_{n,x}$ then their exact value is irrelevant. More precisely, if

$$\min \{ |d_{n,x}^{(i,j)}| : 1 \leq i < j \leq K_{n,x} \} > \max \{ \text{diam}(F_{n,x}^{(i)}) : 1 \leq i \leq K_{n,x} \}, \quad (7.8)$$

then all that matter for the value of $\eta_{n,x} = \varphi_{G_n}(t_n)_x$ are the signs of $d_{n,x}^{(i,j)}$. Therefore, we introduce the permutation $\sigma_{n,x}$ of $\{1, \dots, K_{n,x}\}$ by requiring

$$d_{t_n}(v_{n,x}^{(\sigma(1))}, s) \leq d_{t_n}(v_{n,x}^{(\sigma(2))}, s) \leq \dots \leq d_{t_n}(v_{n,x}^{(\sigma(K_{n,x}))}, s).$$

In case of ties, we break them according to an arbitrary fixed rule. We write Σ_k for the set of permutations of $\{1, \dots, k\}$, so that on the event $\{K_{n,x} = k\}$, we have $\sigma_{n,x} \in \Sigma_k$. We summarize the above observations on the dependence on $\sigma_{n,x}$ in the following lemma.

Lemma 7.2. *Let $G = (V, E)$ be an infinite graph. For every $x \in V$, $1 \leq k \leq \deg_G(x)$ and $V_n \supset \mathcal{N}_x$ there exist functions $f_{k,x}(F^{(1)}, v^{(1)}, \dots, F^{(k)}, v^{(k)}, s)$ (where $s \in \Sigma_k$) with values in $\{0, \dots, \deg_G(x) - 1\}$ such that whenever (7.8) holds, we have*

$$\eta_{n,x} = f_{K_{n,x},x}(F_{n,x}^{(1)}, v_{n,x}^{(1)}, \dots, F_{n,x}^{(K_{n,x})}, v_{n,x}^{(K_{n,x})}, \sigma_{n,x}).$$

We will show that if all bounded harmonic functions are constant, then $\sigma_{n,x}$ is asymptotically uniform. More precisely, conditioned on $K_{n,x} = k$ and $(F_{n,x}^{(i)}, v_{n,x}^{(i)})_{i=1}^{K_{n,x}}, \sigma_{n,x}$ converges in distribution, as $n \rightarrow \infty$, to a uniform random element of Σ_k . Assuming this (and a certain consistency property between the permutations corresponding to different $x_1, x_2 \in V$), we can define the measure ν that is the candidate for the limit.

Let \mathbf{t} be a sample from the WSF on G . Consider a random linear ordering of the components of \mathbf{t} that has the property that it induces the uniform permutation on any finite subset of components. This can be realized for example by assigning i.i.d. $\text{Unif}(0, 1)$ variables to the components, and considering the ranking induced by these. Given components $t_1 \neq t_2$ of \mathbf{t} , we write $t_1 < t_2$, if t_1 precedes t_2 in the ordering. For any $x \in V$, define $\sigma_x \in \Sigma_{K_x}$ by requiring:

$$\sigma_x(i) < \sigma_x(j) \quad \text{if and only if} \quad t_x^{(i)} < t_x^{(j)} \quad \text{for all } 1 \leq i < j \leq K_x.$$

Define the configuration $\eta \in \Omega_G$ by setting

$$\eta_x := f_{K_x,x}(F_x^{(1)}, v_x^{(1)}, \dots, F_x^{(K_x)}, v_x^{(K_x)}, \sigma_x), \quad (7.9)$$

that is defined μ -a.s., under Assumption 7.1 (iii). Let ν be the image of the measure μ under the map $\mathbf{t} \mapsto \eta$.

In order to prove Theorem 7.2, we need to show that for any $A \subset V$ finite, the joint distribution of $\{\eta_{n,x}\}_{x \in A}$ converges to the joint distribution of $\{\eta_x\}_{x \in A}$. We extend to this situation some of the definitions made for single points. Namely, let

$$\begin{aligned} K_{n,A}(\mathbf{t}_n) &:= \text{number of connected components of } \mathbf{t}_n \text{ intersecting } \bigcup_{x \in A} \mathcal{N}_x, \\ t_{n,A}^{(1)}, \dots, t_{n,A}^{(K_{n,A})} &:= \text{the components of } \mathbf{t}_n \text{ that intersect } \bigcup_{x \in A} \mathcal{N}_x. \end{aligned} \quad (7.10)$$

We define $(F_{n,A}^{(i)}, v_{n,A}^{(i)})$, $d_{n,A}^{(i,j)}$ and $\sigma_{n,A}$ completely analogously to the single point case. We also introduce K_A , $t_A^{(1)}, \dots, t_A^{(K_A)}$, $(F_A^{(i)}, v_A^{(i)})$ in the infinite graph G .

The following two propositions make precise the intuition about fluctuations of $d_{n,A}^{(i,j)}$ and the uniformity of $\sigma_{n,A}$.

Proposition 7.1. *Suppose $G = (V, E)$ satisfies Assumption 7.1 (i)–(iii). For any finite $A \subset V$ we have*

$$\lim_{M \rightarrow \infty} \limsup_{n \rightarrow \infty} \mathbf{P} \left[\min_{1 \leq i < j \leq K_{n,A}} |d_{n,A}^{(i,j)}| \leq M \right] = 0. \quad (7.11)$$

Proposition 7.2. *Suppose $G = (V, E)$ satisfies Assumption 7.1 (i)–(iv). For any finite $A \subset V$, any $k \geq 1$, $s \in \Sigma_k$ and any sequence of finite rooted trees $(F^{(i)}, v^{(i)})_{i=1}^k$ we have*

$$\begin{aligned} & \lim_{n \rightarrow \infty} \mathbf{P} [K_{n,A} = k, \sigma_{n,A} = s, (F_{n,A}^{(i)}, v_{n,A}^{(i)}) = (F^{(i)}, v^{(i)}), 1 \leq i \leq k] \\ &= \frac{1}{k!} \mathbf{P} [K_A = k, (F_A^{(i)}, v_A^{(i)}) = (F^{(i)}, v^{(i)}), 1 \leq i \leq k]. \end{aligned} \quad (7.12)$$

We prove Proposition 7.1 in the next section, and Proposition 7.2 in Section 7.5. The short proof of Theorem 7.2, using the two Propositions, is at the end of Section 7.5.

7.4. Lower bound on fluctuations

We start with some preparations for the proof of Proposition 7.1. Let $k \geq 1$ and let $v^{(1)}, \dots, v^{(k)} \in V$ be fixed vertices. We analyze the event

$$\{K_x = k, v_x^{(1)} = v^{(1)}, \dots, v_x^{(k)} = v^{(k)}\} \quad (7.13)$$

using Wilson's method. Let u_1, \dots, u_L be an enumeration of $\{v^{(1)}, \dots, v^{(k)}\} \cup (\cup_{x \in A} \mathcal{N}_x)$ such that $u_i = v^{(i)}$ for $i = 1, \dots, k$. Similarly to the proof of Lemma 7.1, we couple the algorithms in G and in G_n by using the same random walks in Wilson's method, with the walk S^j starting at u_j . (But note that this time the enumeration is different.) On the event (7.13), for fixed $1 \leq i < j \leq k$, the occurrence of $|d_{n,A}^{(i,j)}| \leq M$ implies that the lengths of the two (independent) paths $\gamma^i = \text{LE}(S^i[0, \tau_n^i])$ and $\gamma^j = \text{LE}(S^j[0, \tau_n^j])$ differ by at most M . This will be unlikely, if there is any fluctuation in the length of the paths, and we show that this is the case whenever G is not a tree. The proof will be based on some lemmas that follow. The first lemma makes a deterministic statement about the existence of a cycle with two infinite paths satisfying some requirements. The significance of the statement is that it will allow us to construct two finite random walk paths of equal number of steps between two vertices such that the loop-erasures of the paths have different lengths. This will be sufficient to establish non-trivial fluctuations in the length of the loop-erased walk.

In what follows, o will denote a fixed vertex of G .

Lemma 7.3. *Let $G = (V, E)$ be an infinite vertex-transitive graph that is not a tree. There exists a cycle $C = \{t_1, \dots, t_L\}$ in G , vertices t_I, t_J in C and infinite paths $\pi = \{\pi_0, \pi_1, \dots\}$ and $\rho = \{\rho_0, \rho_1, \dots\}$ in G such that:*

- (i) t_I and t_J are not antipodal in C , i.e. $I \neq J + L/2 \pmod L$;
- (ii) $\pi(0) = t_I$ and $\rho(0) = t_J$;
- (iii) C , $\pi[1, \infty)$ and $\rho[1, \infty)$ are disjoint.

Proof. The conditions imply that the vertex degree is ≥ 3 . Let C be a cycle of length ≥ 3 in G that passes through o . We assume that C has minimal length. Using transitivity, we see that there exists a bi-infinite path $\dots, s_{-2}, s_{-1}, s_0 = o, s_1, s_2, \dots$ in G . Let

$$\begin{aligned} v &:= s_{-k}, \quad \text{where } k = \max\{r \geq 1 : s_{-r} \in C\}; \\ w &:= s_l, \quad \text{where } l = \max\{r \geq 1 : s_r \in C\}. \end{aligned}$$

If C has odd length, we can set $t_i = v$, $t_j = w$, $\pi = \{s_{-k}, s_{-(k+1)}, \dots\}$, $\rho = \{s_l, s_{l+1}, \dots\}$. Henceforth assume that $|C| \geq 4$ and $|C|$ even. If v and w are not antipodal, there is nothing further to prove. Assume that v and w are antipodal, and let $u \in C$ be a neighbour of v .

Case (a): $|C| = 4$. Let u' be the other neighbour of v in C . Since u has at least three distinct neighbours, u has a neighbour $u_1 \neq v, w$.

If $u_1 = u'$, then $\{v, u, u_1\}$ is a cycle of length 3, contradicting the minimality of C .

If $u_1 = s_{-k_1}$ for some $k_1 > k$, then the triple $C' := \{u_1 = s_{-k_1}, s_{-k_1+1}, \dots, s_{-k} = v, u\}$, $\pi := s[-k_1, -(k_1+1), \dots)$, $\rho := \{u, w = s_l, s_{l+1}, \dots\}$ satisfies the requirements of the lemma. Similarly, if $u_1 = s_{l_1}$ for some $l_1 > l$, we are done.

If none of the above holds, select an infinite self-avoiding path

$$u, u_1, u_2, u_3 \dots;$$

such a path is easily seen to exist, using transitivity. If this path is disjoint from

$$B := C \cup \{s_{-k}, s_{-(k+1)}, \dots\} \cup \{s_l, s_{l+1}, \dots\}, \quad (7.14)$$

then the triple C , $\pi = \{s_{-k}, s_{-(k+1)}, \dots\}$ and $\rho = \{u, u_1, u_2, \dots\}$ satisfies the requirements of the lemma. Therefore, suppose that for some $m \geq 1$ we have $u_m \in B$, and let m be the smallest index with this property.

If $u_m = u'$, then the triple $C' = \{v, u, u_1, \dots, u_m = u'\}$, $\pi = \{v = s_{-k}, s_{-(k+1)}, \dots\}$, $\rho = \{u, w = s_l, s_{l+1}, \dots\}$ satisfies the requirements. If $u_m = v$, then necessarily $m \geq 2$ and the triple $C' = \{v, u, u_1, \dots, u_{m-1}\}$, $\pi = \{s_{-k}, s_{-(k+1)}, \dots\}$, $\rho = \{u, w = s_l, s_{l+1}, \dots\}$ works. If $u_m = s_{-k_1}$ for some $k_1 > k$, then consider the cycle $C' = \{v, u, u_1, \dots, u_m = s_{-k_1}, s_{-k_1+1}, \dots, s_{-k-1}\}$. If u_m and u are not antipodal in C' , then we can set $\pi = \{s_{-k_1}, s_{-(k_1+1)}, \dots\}$ and $\rho = \{u, w = s_l, s_{l+1}, \dots\}$. If they are antipodal, then u_m and v are not antipodal in C' , so we can replace ρ by $\rho' = \{v, u', w = s_l, s_{l+1}, \dots\}$. Similarly, we are done if $u_m = s_{l_1}$ for some $l_1 \geq l$. This complete Case (a).

Case (b): $|C| \geq 6$, $|C|$ even. Again we start by letting u be a neighbour of v in C , and let \bar{u} be the neighbour of u in C different from v . Let u_1 be a neighbour of u different from v, \bar{u} . Let b denote the path in C from u to w passing through \bar{u} , and let a denote the path in C from v to w not passing through u .

Select a self-avoiding path u, u_1, u_2, \dots . If this path does not intersect B (as defined in (7.14)), we are done similarly to Case (a). If there is an intersection, let the first one be u_m ($m \geq 1$).

If $u_m = s_{-k_1}$ for some $k_1 > k$, we consider $C' = \{s_{-k_1}, s_{-k_1+1}, s_{-k} = v, u, u_1, \dots, u_{m-1}\}$. If u_m and u are not antipodal in C' , we set $\pi = \{s_{-k_1}, s_{-(k_1+1)}, \dots\}$ and $\rho = b \cup \{s_{l+1}, s_{l+2}, \dots\}$. If they are antipodal, then u_m and v are not antipodal, and we can replace ρ by $\rho' = a \cup \{s_{l+1}, s_{l+2}, \dots\}$. If $u_m = v$ (and necessarily $m \geq 2$), then we set $C' = \{v, u, u_1, \dots, u_{m-1}\}$, $\pi = \{s_{-k}, s_{-(k+1)}, \dots\}$ and $\rho = b \cup \{s_{l+1}, s_{l+2}, \dots\}$.

If $u_m = s_{l_1}$ for some $l_1 > l$, then the triple $C, \pi = \{v = s_{-k}, s_{-(k+1)}, \dots\}$, $\rho = \{u, u_1, \dots, u_m = s_{l_1}, s_{l_1+1}, \dots\}$ works. If $u_m = w$, we note that the path $\{u, u_1, \dots, u_m = w\}$ has to be longer than b , otherwise their union gives a cycle shorter than C . In particular in the cycle $C' = \{u, u_1, \dots, u_m\} \cup b$ the vertices u and w are not antipodal. Hence the choice $\pi = \{u, v = s_{-k}, s_{-(k+1)}, \dots\}$ and $\rho = \{w = s_{l_1}, s_{l_1+1}, \dots\}$ works.

If $u_m \in a \setminus \{v, w\}$, we can find a cycle containing u, v , and part of a , and use $\pi = \{v = s_{-k}, s_{-(k+1)}, \dots\}$ and $\rho = b \cup \{s_{l+1}, s_{l+2}, \dots\}$. Finally, assume that $u_m \in b \setminus \{u, w\}$, and let $c = \{u, u_1, \dots, u_m\}$. Let b' be the subpath of b from u to u_m . If $|c| = |b'|$, then $c \cup b'$ yields a cycle shorter than C , a contradiction (note that the case $m = 1$ is excluded here, since $u_1 \neq \bar{u}$). Therefore, $|c| \neq |b'|$, and hence u and u_m are not antipodal in the cycle $c \cup b'$. Therefore, we can set $\pi = \{u, v = s_{-k}, s_{-(k+1)}, \dots\}$ and $\rho = (b \setminus b') \cup \{s_{l_1}, s_{l_1+1}, \dots\}$. \square

Consider the configuration constructed in Lemma 7.3. We assume the labeling is such that $I = 1$. Shifting by an automorphism we may assume that $\pi(1) = o$. Let G_0 denote the finite graph consisting of the cycle $\{t_1, t_2, \dots, t_L\}$ together with the edges $\{\pi(1), t_1\}$ and $\{\rho(1), t_J\}$. We define two nearest neighbour paths in G_0 such that:

- (i) they both start at $\pi(1)$ and end at $\rho(1)$;
- (ii) they both visit each edge of G_0 ;
- (iii) they have the same number of steps $2L + J + 1$;
- (iv) their loop-erasures have different lengths.

Let

$$\begin{aligned} \beta_1 &:= [\pi(1), t_1, t_2, \dots, t_L, t_1, t_L, t_{L-1}, \dots, t_J, t_{J-1}, \dots, t_2, t_1, t_2, \dots, t_J, \rho(1)], \\ \beta_2 &:= [\pi(1), t_1, t_2, \dots, t_L, t_1, t_L, t_{L-1}, \dots, t_J, \rho(1), t_J, \dots, \rho(1), t_J, \rho(1)]. \end{aligned}$$

Here $\rho(1), t_J$ is repeated as many times as necessary so that the length of β_2 is $2L + J + 1$. The loop-erasure of β_1 has length $L - J + 3$, while the loop-erasure of β_2 has length $J + 1 \neq L - J + 3$.

We want to show that a long loop-erased random walk in G will contain copies of $\text{LE}(\beta_1)$ and $\text{LE}(\beta_2)$ with positive densities. We can do this by an adaptation of an argument of Lawler [18, Theorem 7.7.2]. For this it will be convenient to define a bi-infinite simple random walk by letting $\{S(m)\}_{m \geq 0}$ and $\{S(-m)\}_{m \geq 0}$ be independent realizations of simple random walk on G starting at o .

We set $M = 2L + J + 1$, and we consider the blocks

$$B_k = [S(Mk), S(Mk + 1), \dots, S(M(k + 1))].$$

Let AUT_o denote the stabilizer of o in $\text{AUT}(G)$. By [31, Lemma (1.27)], AUT_o is compact, and hence it carries a right-invariant Haar measure λ of total mass 1. For each $x \in V$ we fix an automorphism ϕ_x that takes o to x .

Definition 7.1. We say that an index $j \geq 0$ is *good*, if the following conditions are satisfied:

- (a) For some $\psi \in \text{AUT}_o$ we have $\psi\phi_{S(Mj)}^{-1}B_j = \beta_1$ or $\psi\phi_{S(Mj)}^{-1}B_j = \beta_2$;
- (b) $S(-\infty, Mj) \cap B_j = \emptyset$ and $S(M(j + 1), \infty) \cap B_j = \emptyset$;
- (c) $S(-\infty, Mj] \cap S[M(j + 1), \infty) = \emptyset$.

Let $b := \mathbf{P}[0 \text{ is good}]$. In what follows, we write

$$B(x, k) := \{y \in V : \text{dist}(x, y) \leq k\}.$$

We also introduce the notation $\xi_B = \inf\{n \geq 0 : S(n) \in B\}$ for the hitting time of B by S . The following lemma shows that good indices occur with positive frequency.

Lemma 7.4. *Suppose the graph $G = (V, E)$ satisfies Assumption 7.1 (i)–(ii).*

- (1) *We have $b > 0$.*
- (2) *For any $\varepsilon > 0$, we have*

$$\mathbf{P} \left[\exists K_0 \forall K \geq K_0 \text{ there are at least } (b - \varepsilon)K \text{ good} \right. \\ \left. \text{indices among } 0, \dots, K - 1 \right] = 1.$$

Proof of (1). Assumption 7.1 (ii) implies that G is transient, in particular. Transience and reversibility of the simple random walk imply that for any finite $B \subset V$ we have

$$\lim_{K \rightarrow \infty} \sup_{z \notin B(o, K)} \mathbf{P}[\xi_B < \infty \mid S(0) = z] = 0. \quad (7.15)$$

Let $\{S^{(1)}(n)\}_{n \geq 0}$ and $\{S^{(2)}(n)\}_{n \geq 0}$ be independent simple random walks on G , with possibly different initial states. It is easy to see that Assumption 7.1 (ii) and transitivity implies

$$\mathbf{P}[S^{(1)}(0, \infty) \cap S^{(2)}(0, \infty) = \emptyset \mid S^{(1)}(0) = o = S^{(2)}(0)] =: \delta_0 > 0.$$

We show that we also have

$$\inf_{x \neq y} \mathbf{P}[S^{(1)}[0, \infty) \cap S^{(2)}[0, \infty) = \emptyset \mid S^{(1)}(0) = x, S^{(2)}(0) = y] =: \delta > 0. \quad (7.16)$$

By transitivity, we may assume $x = o$. By (7.15), we can find K_0 such that $\text{dist}(o, y) \geq K_0$ implies $\mathbf{P}[\xi_o^{(2)} < \infty \mid S^{(2)}(0) = y] < \delta_0/2$, where $\xi_o^{(2)}$ denotes the hitting time of o by $S^{(2)}$. Let $\{S^{(3)}(n)\}_{n \geq 0}$ and $\{S(n)\}_{n \geq 0}$ be a third and a fourth independent simple random walk, both starting at o . Assume the event

$$A_{1,3} = \{S^{(1)}(0, \infty) \cap S^{(3)}(0, \infty) = \emptyset\}.$$

By Lévy's 0–1 law, a.s. on the event $A_{1,3}$ we have

$$\lim_{n \rightarrow \infty} \mathbf{P}[S^{(3)}[0, \infty) \cap S^{(1)}[0, \infty) = \emptyset \mid S^{(3)}[0, n], S^{(1)}] = 1.$$

In particular, a.s. on $A_{1,3}$ the random variables

$$\begin{aligned} X_{1,3} &= \inf_{z \in S^{(3)}(0, \infty)} \mathbf{P}[S[0, \infty) \cap S^{(1)}[0, \infty) = \emptyset \mid S(0) = z, S^{(3)}] \\ X_{3,1} &= \inf_{z \in S^{(1)}(0, \infty)} \mathbf{P}[S[0, \infty) \cap S^{(3)}[0, \infty) = \emptyset \mid S(0) = x, S^{(1)}] \end{aligned}$$

are positive. Hence we can find $c > 0$ and $0 < \delta' < \delta_0/4$ such that

$$\mathbf{P}[A_{1,3}, X_{1,3} \geq c, X_{3,1} \geq c, \xi_o^{(2)} = \infty] \geq \delta'. \quad (7.17)$$

On the event in (7.17), either $S^{(2)}$ never hits $S^{(1)}[0, \infty) \cup S^{(3)}[0, \infty)$, or if it hits one of the paths, then with conditional probability at least c , its continuation from the first hitting point does not hit the other path. By symmetry of the roles of $S^{(1)}$ and $S^{(3)}$, we get $\mathbf{P}[S^{(1)}[0, \infty) \cap S^{(2)}[0, \infty) = \emptyset] \geq c\delta'$. The statement involving all $y \neq o$ now follows.

We continue with the proof of statement (1). The probability that $S[0, M]$ traces out exactly β_1 or β_2 is positive. Assume that this occurs. Consider some large K , and let A_K be the event that $\{S(M+k)\}_{k \geq 0}$ and $\{S(-k)\}_{k \geq 0}$ follow the paths ρ and π , respectively, constructed in Lemma 7.3 until they both leave $B(o, K)$, at vertices x and y . By (7.15), we can choose K large enough so that

$$\sup_{z \notin B(o, K)} \mathbf{P}[\xi_{G_0} < \infty \mid S(0) = z] \leq \delta/4. \quad (7.18)$$

Due to (7.18) and (7.16), the conditional probability given A_K , that the walks satisfy the requirements (b) and (c) is at least $\delta/2$. This proves part (1). \square

Proof of (2). We want to apply the ergodic theorem. A technical difficulty is that there may be no canonical way to “translate” a vertex $x \in V$ back to o . Hence after shifting the path by ϕ_x^{-1} , we average over AUT_o , which is possible, since AUT_o is compact. This way we can define a certain path-valued stationary process. Let Ψ_0, Ψ_1, \dots be an i.i.d. sequence, independent of the random walk, with each element distributed according to λ . Put

$$X_k(m) := \Psi_k \phi_{S(Mk)}^{-1} S(Mk + m), \quad -\infty < m < \infty, \quad k \geq 0.$$

We claim that the path-valued sequence $\{X_k(\cdot)\}_{k \geq 0}$ is stationary and mixing (on the space of paths we consider the topology of pointwise convergence and the induced Borel σ -field). The somewhat tedious proof of this intuitive claim is deferred to the Appendix.

The proof of (2) is now straightforward from part (1) and the ergodic theorem, noting that j is good if and only if 0 is good relative to the path X_j . \square

A time $-\infty < j < \infty$ is called *loop-free* for S if $S(-\infty, j] \cap S(j, \infty) = \emptyset$. The significance of loop-free points is that loop-erasure on the two sides of a loop-free point do not influence each other. Note that if $k \geq 0$ is good, then kM and $(k+1)M$ are loop-free. This observation and Lemma 7.4 immediately implies the following lemma.

Lemma 7.5. *Suppose the graph $G = (V, E)$ satisfies Assumption 7.1 (i)–(ii). There exists $b' > 0$ such that*

$$\mathbf{P} \left[\exists K_0 \forall K \geq K_0 \text{ there are at least } b'K \text{ loop-free points among } 0, \dots, K-1 \right] = 1.$$

The lower bound on the fluctuations can now be achieved by conditioning on “all information outside the good blocks”. In order to make this precise, for each good index $k \geq 0$, we choose $\psi_k \in \text{AUT}_o$ such that $\psi_k \phi_{S(kM)}^{-1} B_k \in \{\beta_1, \beta_2\}$. Note that since β_1 and β_2 both traverse G_0 , if ψ'_k is another such automorphism, then $\psi'_k \psi_k^{-1}$ fixes G_0 pointwise. In particular, $\psi_k^{-1}|_{G_0} \equiv \psi'_k|_{G_0}$, where $|_{G_0}$ denotes restriction to G_0 . We define the σ -algebra \mathcal{G} generated by the following random objects:

$$\begin{aligned} & S(kM), \quad k \geq 0; \\ & Y_k := I[k \text{ is good}], \quad k \geq 0; \\ & \psi_k^{-1}|_{G_0} \text{ for } k \geq 0 \text{ such that } Y_k = 1; \\ & \text{the paths } B_{k'} = [S(k'M), S(k'M+1), \dots, S((k'+1)M)], \\ & \quad \text{for } k' \geq 0 \text{ such that } Y_{k'} = 0; \\ & S(-\infty, 0]. \end{aligned} \tag{7.19}$$

Lemma 7.6. *Given \mathcal{G} , the good blocks are conditionally independent, and conditional on \mathcal{G} , such a block B_k takes the values $\Phi_k\beta_1$ and $\Phi_k\beta_2$ with probabilities $1/2$ each, for some \mathcal{G} -measurable automorphisms $\{\Phi_k\}$ that take o to $S(kM)$, respectively.*

Proof. We know that almost surely there are infinitely many good indices. Fix $N \geq 2$. Consider the class \mathcal{P}_N of events of the form:

$$\begin{aligned} E = \{ & S(-j) = y_j, j = 1, \dots, J; \\ & S(kM) = z_k, k = 0, \dots, K; \\ & k \in I \text{ are good;} \\ & k' \in \{0, \dots, K-1\} \setminus I \text{ are not good} \\ & \psi_k \phi_{S(kM)}^{-1} B_k \in \{\beta_1, \beta_2\} \text{ for } k \in I; \\ & \psi_k^{-1}|_{G_0} = \alpha; \\ & B_{k'} = b_{k'} \text{ for } k' \in \{0, \dots, K-1\} \setminus I\}, \end{aligned}$$

where $J, K \geq N$, $z_k, I \subset \{0, \dots, K-1\}$, $|I| = N$, α , and $b_{k'}$ are fixed. Let $A_k^{\varepsilon_k}$ be the event $\{\psi_k \phi_{z_k}^{-1} B_k = \beta_{\varepsilon_k}\}$, where $\varepsilon_k \in \{1, 2\}$ for $k = 0, \dots, K-1$. By decomposing the path of S into $\{S(-n)\}_{n \geq 0}$, the blocks B_k , $k = 0, \dots, K-1$, and $\{S(n)\}_{n \geq KM}$, we see that

$$\mathbf{P}\left[\left(\bigcap_{k \in I} A_k^{\varepsilon_k}\right) \cap E\right] = \left(\prod_{k \in I} \frac{1}{2}\right) \mathbf{P}[E].$$

Since \mathcal{P}_N is closed under intersection, and generates \mathcal{G} , this implies conditional independence of the first N good blocks. The lemma follows. \square

Proof of Proposition 7.1. Let $\varepsilon > 0$ be given. Due to Lemma 7.1, there exists a finite $B \subset V$ such that for all large enough n , with probability at least $1 - \varepsilon$, we have $v_{n,x}^{(i)} \in B$ for $i = 1, \dots, K_{n,x}$. Hence the Proposition will follow, once we show that for any $v, w \in B$ we have

$$\lim_{M \rightarrow \infty} \limsup_{n \rightarrow \infty} \mathbf{P}[\pi_{n,v} \cap \pi_{n,w} = \emptyset, |d_{t_n}(v, s) - d_{t_n}(w, s)| \leq M] = 0. \quad (7.20)$$

Let S^1 and S^2 be independent simple random walks starting at v and w , respectively. Let $\gamma^i := \text{LE}(S^i[0, \tau_n^i])$. We apply Lemma 7.6 to the random walk $S = S^1$. Let

$$T_n := \sup\{j \geq 0 : B_j \subset V_n \text{ and } j \text{ is good}\}$$

be the index of the last good block completed before τ_n^1 . Let the set of good indices be: $\{g_1, g_2, \dots\}$. We define the stretches between good blocks: we let

$$\rho_0 := [S^1(0), S^1(1), \dots, S^1(g_1 M)],$$

and for $k \geq 1$ we let

$$\rho_k := [S^1((g_k + 1)M), S^1((g_k + 1)M + 1), \dots, S^1(g_{k+1}M)].$$

Observe that loop-erasure of the ρ_k 's do not interfere with each other, due to item (c) of Definition 7.1. Hence $\gamma^1 = \text{LE}(S^1[0, \tau_n^1])$ is the concatenation of:

$$\begin{aligned} & \text{LE}(\rho_0), \text{LE}(B_{g_1}), \text{LE}(\rho_1), \text{LE}(B_{g_2}), \dots, \\ & \text{LE}(\rho_{T_n-1}), \text{LE}(B_{T_n}), \text{LE}(S^1[(T_n + 1)M, \tau_n^1]). \end{aligned}$$

Due to Lemma 7.4 (ii), for any K we have

$$\lim_{n \rightarrow \infty} \mathbf{P}[T_n \geq K] = 1.$$

Now condition on the random walk S^2 , condition on the set of good indices and the bad blocks of S^1 up to exit from V_n , and condition on the event $\{T_n \geq K\}$. Then

$$\text{length}(\gamma^1) - \text{length}(\gamma^2) = (Y_0 - \text{length}(\gamma^2)) + \sum_{j=1}^K Y_j,$$

where the value of the first term is fixed by the conditioning, and the Y_j are conditionally i.i.d. with positive variance. By the local central limit theorem [29], we get

$$\mathbf{P}[|\text{length}(\gamma^1) - \text{length}(\gamma^2)| \leq M] \leq \frac{cM}{\sqrt{K}}.$$

Letting $K \rightarrow \infty$ implies the claim in (7.20), and hence the Proposition follows. \square

7.5. Asymptotic uniformity of the permutation

In this section we prove Proposition 7.2, and complete the proof of Theorem 7.2.

Let $k \geq 1$ be fixed and let $(F^{(i)}, v^{(i)})$, $1 \leq i \leq k$ be a fixed sequence of finite rooted trees in G . We will use Wilson's method to generate samples \mathbf{t} and t_n from the measures μ and μ_n . The set-up is the same as in Section 7.4, that is, we use the same random walks S^j started at the vertices

$$u_1 = v^{(1)}, \dots, u_k = v^{(k)}, u_{k+1}, \dots, u_L,$$

where u_{k+1}, \dots, u_L is an enumeration of $\cup_{x \in A} \mathcal{N}_x$. Recall that T^i and T_n^i denote the hitting times of \mathcal{F}_{i-1} and $\mathcal{F}_{n,i-1}$, respectively, by S^i .

Let B_1 denote the set of vertices in $\cup_{i=1}^k F^{(i)}$. Let

$$C := \left(\bigcap_{i=1}^k \{T^i = \infty\} \right) \cap \{\mathcal{F}_k \cap B_1 = \{v^{(1)}, \dots, v^{(k)}\}\}.$$

Let

$$C' := \left\{ \bigcup_{j=k+1}^L \text{LE}(S^j[0, T^j]) = \bigcup_{i=1}^k F^{(i)} \right\}.$$

Observe that as long as $(F^{(i)}, v^{(i)})_{i=1}^k$ is a possible sequence for $(F_A^{(i)}, v_A^{(i)})_{i=1}^k$, we have

$$C \cap C' = \{K_A = k, (F_A^{(i)}, v_A^{(i)}) = (F^{(i)}, v^{(i)}), 1 \leq i \leq k\}. \quad (7.21)$$

We also introduce

$$C_n := \left(\bigcap_{i=1}^k \{T_n^i = \tau_n^i\} \right) \cap \{\mathcal{F}_{n,k} \cap B_1 = \{v^{(1)}, \dots, v^{(k)}\}\},$$

$$C'_n := \left\{ \bigcup_{j=k+1}^L \text{LE}(S^j[0, T_n^j]) = \bigcup_{i=1}^k F^{(i)} \right\},$$

and observe that

$$C_n \cap C'_n = \{K_{n,A} = k, (F_{n,A}^{(i)}, v_{n,A}^{(i)}) = (F^{(i)}, v^{(i)}), 1 \leq i \leq k\}. \quad (7.22)$$

Here is the outline of the proof. The restriction involving B_1 has little effect on the walks S^i , $i = 1, \dots, k$, once they are far away from B_1 , and likewise, the condition $\{T^i = \infty\}$, $i = 1, \dots, k$. Therefore, for some large n' , once these walks leave $V_{n'}$, they can be treated as independent. The point where Assumption 7.1(iv) (bounded harmonic functions are constant) becomes crucial, is to show that the walks can also be treated as having the same distribution. Namely, we show that Assumption 7.1(iv) implies that for some $n'' > n'$, the exit measures of the walks on $\partial V_{n''}$ are nearly identical in total variation distance. Therefore, their continuations are nearly i.i.d. This will imply the near uniformity of $\sigma_{n,A}$ for $n \gg n''$.

Let $\varepsilon > 0$ be given. As in the proof of Lemma 7.1, we deduce that

$$\lim_{n \rightarrow \infty} \mathbf{P} \left[\left(\bigcap_{i=1}^k \{T^i = \infty\} \right) \Delta \left(\bigcap_{i=1}^k \{T_n^i = \tau_n^i\} \right) \right] = 0, \quad (7.23)$$

where Δ denotes symmetric difference. Letting $\hat{\tau}_{B_1}^i$ denote the last visit by S^i to the set B_1 , transience implies that for each $i = 1, \dots, k$ we have

$$\lim_{n \rightarrow \infty} \mathbf{P} [S^i[\tau_n^i, \infty) \cap S^i[0, \hat{\tau}_{B_1}^i] \neq \emptyset] = 0. \quad (7.24)$$

It follows from (7.23) and (7.24) that there exists n_1 such that for all $n \geq n_1$ we have

$$\mathbf{P}[C \Delta C_n] < \varepsilon. \quad (7.25)$$

Since on the event $C \cap C'$ we have $T^j < \infty$ for $j = k+1, \dots, L$, we can find a large enough finite set $B_2 \subset V$ ($B_2 \supset B_1$) such that with $G_1 := \cap_{j=k+1}^L \{S^j[0, T^j] \subset B_2\}$ we have

$$\mathbf{P}[C \cap C' \cap G_1^c] < \varepsilon. \quad (7.26)$$

Let

$$G_{n,2} := \bigcap_{i=1}^k \{S^i[\tau_n^i, \infty) \cap S^i[0, \hat{\tau}_{B_2}^i] = \emptyset\}.$$

Applying (7.24) for B_2 in place of B_1 , we get that there exists n_2 such that for all $n \geq n_2$ we have

$$\mathbf{P}[G_{n,2}] > 1 - \varepsilon. \quad (7.27)$$

It follows from (7.25), (7.26) and (7.27) that there exists n_3 such that for all $n \geq n_3$ we have

$$\mathbf{P}[(C \cap C') \Delta (C_n \cap C'_n)] < 3\varepsilon. \quad (7.28)$$

Let S denote a simple random walk independent of the S^j 's. Lévy's 0–1 law implies that for each $i = 2, \dots, k$, almost surely we have

$$\lim_{n \rightarrow \infty} \mathbf{P} \left[S[0, \infty) \cap \left(\bigcup_{\substack{1 \leq j \leq k \\ j \neq i}} S^j[0, \infty) \right) = \emptyset \mid S(0) = S^i(\tau_n^i), \right. \\ \left. S^j, j = 1, \dots, k, j \neq i \right] = 1.$$

Hence we can find n_4 such that with

$$G_{n,3} := \bigcap_{i=1}^k \left\{ S^i[\tau_n^i, \infty) \cap \left(\bigcup_{\substack{1 \leq j \leq k \\ j \neq i}} S^j[0, \infty) \right) = \emptyset \right\}$$

for all $n \geq n_4$ we have

$$\mathbf{P}[G_{n,3}] > 1 - \varepsilon. \quad (7.29)$$

Let $n' := \max\{n_1, n_2, n_3, n_4\}$. Given $D \subset V$, $B \subset \partial D$ and $z \in \bar{D} := D \cup \partial D$, let

$$h_D(z, B) := \mathbf{P}[S(\tau_D) \in B \mid S(0) = z]$$

denote the exit measure of simple random walk on the boundary of D . Note that $h_D(\cdot, B)$ is harmonic in D for any $B \subset \partial D$, and $0 \leq h_D(z, B) \leq 1$. Here

we write $\partial D = \{y \in V \setminus D : y \sim x \text{ for some } x \in D\}$. We show that we can find an index $n'' > n'$ such that

$$\sup_{z_1, z_2 \in \partial V_{n''}} \|h_{V_{n''}}(z_1, \cdot) - h_{V_{n''}}(z_2, \cdot)\| \leq \varepsilon, \quad (7.30)$$

where $\|\cdot\|$ denotes total variation distance. Indeed, if this was not the case we could find $z_1, z_2 \in \partial V_{n'}$, and a sequence $r_1 < r_2 < \dots$ and subsets $A_i \subset \partial V_{r_i}$ such that

$$|h_{V(r_i)}(z_1, A_i) - h_{V(r_i)}(z_2, A_i)| \geq \varepsilon, \quad i = 1, 2, \dots \quad (7.31)$$

By passing to a subsequence, we may assume that the limit

$$h(z) := \lim_{i \rightarrow \infty} h_{V(r_i)}(z, A_i)$$

exists for all $z \in V$. From (7.31) we have $|h(z_1) - h(z_2)| \geq \varepsilon$. However, h is a bounded harmonic function, so it must be constant by Assumption 7.1 (iv). This contradiction proves (7.30).

Consider now $n > n''$, and let

$$f(y) := \mathbf{P}[S(\tau_{n''}) = y \mid S(0) = o], \quad y \in \partial V_{n''}.$$

It follows from (7.30) that

$$\|f(\cdot) - h_{V_{n''}}(x_i, \cdot)\| \leq \varepsilon, \quad i = 1, \dots, k. \quad (7.32)$$

Hence, for any $x_1, \dots, x_k \in \partial V_{n'}$ there exists a coupling $g_{x_1, \dots, x_k}(z_1, \dots, z_k, y_1, \dots, y_k)$ with marginals

$$\begin{aligned} \sum_{y_1, \dots, y_k} g_{x_1, \dots, x_k}(z_1, \dots, z_k, y_1, \dots, y_k) &= \prod_{i=1}^k h_{V_{n''}}(x_i, z_i), \\ \sum_{z_1, \dots, z_k} g_{x_1, \dots, x_k}(z_1, \dots, z_k, y_1, \dots, y_k) &= \prod_{i=1}^k f(y_i), \end{aligned}$$

where

$$\sum_{z_1=y_1, \dots, z_k=y_k} g_{x_1, \dots, x_k}(z_1, \dots, z_k, y_1, \dots, y_k) \geq 1 - O(\varepsilon).$$

Let $\{\tilde{S}^i(n)\}_{n \geq 0}$, $i = 1, \dots, k$ be independent simple random walks with initial distribution f . We couple the initial distribution of the \tilde{S}^i 's to the distribution of $S^i(\tau_{n''}^i)$'s using g , where $x_i = S^i(\tau_{n''}^i)$. In this coupling, we have $S^i(\tau_{n''}^i + m) = \tilde{S}^i(m)$, $m \geq 0$ with probability at least $1 - O(\varepsilon)$.

We define the random permutation $\tilde{\sigma} \in \Sigma_k$ by the condition

$$\text{length}(\text{LE}(\tilde{S}^{\tilde{\sigma}(1)}[0, \tilde{\tau}_n^{\tilde{\sigma}(1)}])) < \dots < \text{length}(\text{LE}(\tilde{S}^{\tilde{\sigma}(k)}[0, \tilde{\tau}_n^{\tilde{\sigma}(k)}])). \quad (7.33)$$

Here, if there are ties, we break them in a uniformly random way. That is, if $\{i_1, \dots, i_r\} \subset \{1, \dots, k\}$ is a maximal set of indices such that the loop-erasures of the paths $\tilde{S}^{i_j}[0, \tilde{\tau}_n^{i_j}]$, $j = 1, \dots, r$ have equal lengths, we pick an ordering on them uniformly at random, and use that ordering in (7.33). This way of breaking ties ensures that $\tilde{\sigma}$ is exactly uniformly distributed on Σ_k .

It follows from Lemma 7.5, and the almost sure finiteness of $\tau_{n''}^i$, that there exist $n_5 > n''$ and an $M_1 < \infty$ such that with

$$G_{n,4} := \bigcap_{i=1}^k \left\{ \text{there exists a loop-free point for } S^i \text{ in } [\tau_{n''}^i, \tau_{n''}^i + M_1] \right\}$$

for all $n \geq n_5$ we have

$$\mathbf{P}[G_{n,4}] > 1 - \varepsilon. \quad (7.34)$$

Occurrence of the event $G_{n,4}$ ensures that when $n \gg n''$, most of the length of $\text{LE}(S^i[0, \tau_n^i])$ comes from the length of $\text{LE}(\tilde{S}^i[0, \tilde{\tau}_n^i])$, for $i = 1, \dots, k$. In particular, there exists a deterministic $M_2 = M_2(M_1, n'')$, such that whenever $C_n \cap G_{n,4} \cap \{|d_{n,A}^{(i,j)}| > M_2\}$ occurs, we have $\sigma_{n,A} = \tilde{\sigma}$.

We are ready to start analyzing the event on the left hand side of (7.12). A straightforward computation shows that for any events $A \in \sigma(S^i[0, \tau_n^i])$, $i = 1, \dots, k$ and $B \in \sigma(\tilde{S}^i[0, \infty])$, $i = 1, \dots, k$ we have

$$\begin{aligned} & |\mathbf{P}[A \cap B] - \mathbf{P}[A] \mathbf{P}[B]| \quad (7.35) \\ & \leq \sup_{\substack{x_1, \dots, x_k \\ y_1, \dots, y_k}} \sum_{\substack{z_1, \dots, z_k \\ y_1, \dots, y_k}} \left| g_{x_1, \dots, x_k}(z_1, \dots, z_k, y_1, \dots, y_k) - \prod_{i=1}^k \delta(z_i, y_i) f(y_i) \right| \leq O(\varepsilon). \end{aligned}$$

We apply this with $A = C_{n'}$ and $B = \{\tilde{\sigma} = s\}$, where $s \in \Sigma_k$ is fixed. To be precise, due to the breaking of ties for $\tilde{\sigma}$, this B is defined on a slightly larger σ -field than in (7.35). But this has no consequence. Using (7.35) and (7.25), for $n > n''$ we obtain

$$\mathbf{P}[C_{n'} \cap \{\tilde{\sigma} = s\}] = \frac{1}{k!} \mathbf{P}[C_{n'}] + O(\varepsilon) = \frac{1}{k!} \mathbf{P}[C] + O(\varepsilon). \quad (7.36)$$

Our goal now is to show that on a slightly different event $\tilde{\sigma}$ can be replaced by $\sigma_{n,A}$.

Due to Proposition 7.1, we can find $n_6 > n''$ such that for all $n \geq n_6$ we have

$$\mathbf{P}\left[\min_{1 \leq i < j \leq k} d_{n,A}^{(i,j)} > M_2\right] \geq 1 - \varepsilon. \quad (7.37)$$

Consider for $n \geq \max\{n_5, n_6\}$ the event

$$\tilde{C}_n := C_{n'} \cap G_{n',2} \cap G_{n',3} \cap G_{n,4} \cap \left\{ \min_{1 \leq i < j \leq k} d_{n,A}^{(i,j)} > M_2 \right\}. \quad (7.38)$$

Observe that $\tilde{C}_n \subset C_n$ and that on \tilde{C}_n , we have $\sigma_{n,A} = \tilde{\sigma}$. Due to the estimates (7.25), (7.27), (7.29) and (7.37), for $n \geq \max\{n_5, n_6\}$ we have

$$\mathbf{P}[\tilde{C}_n \cap \{\sigma_{n,A} = s\}] = \mathbf{P}[C_{n'} \cap \{\tilde{\sigma} = s\}] + O(\varepsilon). \quad (7.39)$$

The presence of the event $G_{n,2}$ in (7.38) ensures that on the event $\tilde{C}_n \cap \{\sigma_{n,A} = s\}$, we have $\mathcal{F}_k \cap B_2 = \mathcal{F}_{n,k} \cap B_2$. Therefore,

$$\mathbf{P}[C'_n \cap G_1 \mid \tilde{C}_n \cap \{\sigma_{n,A} = s\}] = \mathbf{P}[C' \cap G_1 \mid C] = \mathbf{P}[C' \mid C] + O(\varepsilon). \quad (7.40)$$

It follows from (7.36), (7.39) and (7.40) that

$$\mathbf{P}[\tilde{C}_n \cap \{\sigma_{n,A} = s\} \cap C'_n \cap G_1] = \frac{1}{k!} \mathbf{P}[C \cap C'] + O(\varepsilon). \quad (7.41)$$

Since $\mathbf{P}[\tilde{C}_n \triangle C_n] = O(\varepsilon)$ and $\mathbf{P}[G_1^c \cap C \cap C'] < \varepsilon$, (7.41) implies that

$$\mathbf{P}[C_n \cap C'_n \cap \{\sigma_{n,A} = s\}] = \frac{1}{k!} \mathbf{P}[C \cap C'] + O(\varepsilon).$$

Comparing with (7.21) and (7.22), this completes the proof of the Proposition.

Proof of Theorem 7.2. Due to Lemma 7.1, for any $\varepsilon > 0$ there exists a finite $B \subset V$ such that

$$\liminf_{n \rightarrow \infty} \mathbf{P} \left[\bigcup_{i=1}^{K_{n,A}} (F_x^{(i)}, v_x^{(i)}) \subset B \right] \geq 1 - \varepsilon.$$

Hence we can restrict our attention to the finite collection of rooted trees (F, v) that lie inside B . Let $(F^{(1)}, v^{(1)}), \dots, (F^{(k)}, v^{(k)})$ be a possible value of $(F_A^{(1)}, v_A^{(1)}), \dots, (F_A^{(k)}, v_A^{(k)})$, with $K_A = k$, such that all these trees lie inside B .

Lemma 7.1 shows that in a suitable coupling, the events $(F_{n,A}^{(i)}, v_{n,A}^{(i)})_{i=1}^k = (F^{(i)}, v^{(i)})_{i=1}^k$ and $(F_A^{(i)}, v_A^{(i)})_{i=1}^k = (F^{(i)}, v^{(i)})_{i=1}^k$ are asymptotically equal, when this occurs, we have $(F_{n,x}^{(i)}, v_{n,x}^{(i)})_{i=1}^{K_{n,x}} = (F_x^{(i)}, v_x^{(i)})_{i=1}^{K_x}$ for all $x \in A$. Proposition 7.1 implies that for all $x \in A$ and for large enough n the condition (7.8) of Lemma 7.2 holds with high probability. This implies that with high probability, the permutations $\{\sigma_{n,x}\}_{x \in A}$ are determined by $\sigma_{n,A}$. Moreover, the dependence of the collection $\{\sigma_{n,x}\}_{x \in A}$ on $\sigma_{n,A}$ is given by the same (deterministic) function as the dependence of $\{\sigma_x\}_{x \in A}$ on σ_A . Proposition 7.2 implies that conditioned on $(F_A^{(i)}, v_A^{(i)}) = (F^{(i)}, v^{(i)})$, $i = 1, \dots, k$, the distribution of $\sigma_{n,A}$ is close to

uniform. This implies that for each $x \in A$, the joint distribution of $\{\sigma_{n,x}\}_{x \in A}$ is close to the joint distribution of $\{\sigma_x\}_{x \in A}$.

The above considerations, Lemma 7.2, and the definition of η in (7.9) imply that the joint distribution of $\{\eta_{n,x}\}_{x \in A}$ converges to the joint distribution of $\{\eta_x\}_{x \in A}$ as $n \rightarrow \infty$. Hence the theorem follows. \square

8. Infinite volume limits on regular trees

In this section we consider infinite d -regular trees. The paper [22] proved the existence of the limit $\nu_n \Rightarrow \nu$ along sufficiently regular exhaustions (see condition (24) there). It was also claimed that the limit exists along any exhaustion, however this does not follow from the arguments in [22] (note that statement (25) of [22] does not imply the Cauchy net property claimed there). In this section we prove the general convergence result. Note that the proof of Theorem 7.2 does not apply to the infinite d -regular tree, for more than one reason: Assumption 7.1 (iv) is not satisfied, and there is no fluctuation in the lengths of paths, so Proposition 7.1 fails. Nevertheless, the Majumdar–Dhar bijection can still be used to show that $\nu_n \Rightarrow \nu$ along any exhaustion.

Let $G = (T^d, E)$ be the infinite d -regular tree, with $d \geq 3$. We will denote by o an arbitrary fixed vertex of G .

Theorem 8.1. *For any $d \geq 3$ and any exhaustion $V_1 \subset V_2 \subset \dots \subset T^d$, we have $\nu_n \Rightarrow \nu$ for a unique measure ν , independent of the exhaustion.*

We begin with some preparations for the proof. Fix a finite $A \subset T^d$. We need to consider the convergence of $\nu_{G_n}[\eta_{n,x} = h_x, x \in A]$, as $n \rightarrow \infty$, where $h \in \{0, 1, \dots, d-1\}^A$ is fixed.

Let $\{S(n)\}_{n \geq 0}$ denote a simple random walk in T^d . Let $\tau_n := \inf\{k \geq 0 : S(k) \notin V_n\}$, and for $B \subset T^d$, let $\xi_B := \inf\{k \geq 0 : S(k) \in B\}$. The following notation will be useful: given $x \in \partial A$ and $V_n \supset A$, let

$$q_{n,x} := \mathbf{P}[\tau_n < \xi_A \mid S(0) = x].$$

Given a self-avoiding path σ from x to V_n^c that does not visit A , we also define:

$$q_{n,x}(\sigma) := \mathbf{P}[\tau_n < \xi_A \text{ and } \text{LE}(S[0, \tau_n]) = \sigma \mid S(0) = x]. \quad (8.1)$$

Recall that \mathcal{T}_{G_n} is the set of spanning trees of G_n . We will orient edges towards the sink, and view trees as arrow configurations. For $t_n \in \mathcal{T}_{G_n}$ let

$$\mathcal{C}(t_n) := \{y \in \partial A : \exists e \in t_n \text{ such that } e_- = y \text{ and } e_+ \in A\}.$$

Note that we always have $\mathcal{C}(t_n) \subsetneq \partial A$. We classify trees according to the value of \mathcal{C} . Fix $C \subsetneq \partial A$, and consider trees t_n with $\mathcal{C}(t_n) = C$. In any such tree, the path from a vertex $y \in (\partial A) \setminus C$ to s , that is the path $\pi_{n,y}(t_n)$, does not visit A .

Due to Lemma 3.3, the occurrence or not of the event $\{\eta_{n,x} = h_x, x \in A\}$ depends on: the lengths of $\{\pi_{n,y}(t_n)\}_{y \in (\partial A) \setminus C}$ and the position of arrows with tails in A . We will refer to the latter simply as “the arrows in A ”. We denote by $m_{n,y}$ the length of $\pi_{n,y}$. The key to convergence is a remarkable symmetry property of the bijection stated in the next two lemmas.

Lemma 8.1. *For any $C \subset \partial A$ and $\{h_x\}_{x \in A}$, the following alternative holds. Either*

- (A) *for any collection $m_{n,y} \geq 1, y \in (\partial A) \setminus C$, the event $\{\eta_{n,x} = h_x, x \in A\}$ does not occur for any choice of arrows in A ; or*
- (B) *for any collection $m_{n,y} \geq 1, y \in (\partial A) \setminus C$, the event $\{\eta_{n,x} = h_x, x \in A\}$ occurs for exactly one choice of arrows in A .*

Lemma 8.2. *Suppose that $C \subset \partial A$ and that Case (B) holds in Lemma 8.1. Let $\sigma_y, y \in (\partial A) \setminus C$ be fixed self-avoiding paths from each y to s that avoid A . Then*

$$\begin{aligned} \mu_n [\mathcal{C}(t_n) = C; \pi_{n,y}(t_n) = \sigma_y, y \in (\partial A) \setminus C; \eta_{n,x}(t_n) = h_x, x \in A] \\ = f_{A,C}(q_{n,y'}, y' \in \partial A) \prod_{y \in (\partial A) \setminus C} q_{n,y}(\sigma_y) \end{aligned} \quad (8.2)$$

for some function $f_{A,C} : [0, 1]^{\partial A} \rightarrow [0, 1]$, whose form only depends on the pair (A, C) , and not on the σ_y 's. The statement extends to Case (A), by taking $f_{A,C}$ to be the 0 function.

Proof of Lemma 8.1. Consider the following auxiliary graph. We start with the subgraph of G induced by $A \cup \partial A$. For each $y \in (\partial A) \setminus C$ we glue a path of length $m_{n,y}$ at y . All glued on paths end at the common endpoint s , that serves as the sink. No new edges are added for vertices in C . We denote this graph by $G_{A,C}$ (the dependence on the $m_{n,y}$'s is suppressed in the notation). Consider the following sandpile configuration $\eta(h)$ on $G_{A,C}$. On the set A , $\eta(h)$ equals h , on C it equals 0, and on the rest of the vertices it equals 1. It is easy to see using the Burning Algorithm, that whether $\eta(h) \in \mathcal{R}_{G_{A,C}}$ or not is independent of the values of $m_{n,y}$. We claim that if $\eta(h) \notin \mathcal{R}_{G_{A,C}}$ then the statements in Case (A) hold, and if $\eta(h) \in \mathcal{R}_{G_{A,C}}$ then the statements in Case (B) hold.

Consider any $\eta_n \in \mathcal{R}_{G_n}$, for which $\eta_{n,x} = h_x, x \in A$, and for which the Burning Algorithm produces a tree t_n with $\mathcal{C}(t_n) = C$, and paths $\pi_{n,y}$ with lengths $m_{n,y}$. We consider the burning of η_n in G_n in parallel to the burning of $\eta(h)$ in $G_{A,C}$. We show that inside $\partial A \cup A$, each site will burn at the same time in the two processes.

Since the time of burning equals graph distance from the sink in the tree produced by the algorithm, in both configurations the first time when a vertex

of ∂A burns is $m_1 := \min\{m_{n,y} : y \in (\partial A) \setminus C\}$. Let $y_{1,1}, \dots, y_{1,r_1}$ be the list of y 's for which the minimum is achieved. After time m_1 , the status of vertices in the subtree of V_n emanating from each $y_{1,i}$ away from A has no influence on the burning of vertices in $A \cup \partial A$ (they have been disconnected by the burning of the vertex $y_{1,i}$). Hence we may discard these subtrees from V_n for the rest of the process. Let $m_2 := \min\{m_{n,y} : m_{n,y} > m_1, (y \in \partial A) \setminus C\}$, and let $y_{2,1}, \dots, y_{2,r_2}$ be the list of y 's for which the minimum is achieved.

We claim that at all times $m_1 \leq m \leq m_2$, the two burning processes agree in $A \cup \partial A$. We show this by induction on m . The claim holds for $m = m_1$, as in both processes precisely $y_{1,1}, \dots, y_{1,r_1}$ are burnt at time m_1 . Assume the claim holds for some m with $m_1 \leq m < m_2$. Let $z \in A$ be a vertex that is unburnt at time m (in both configurations, necessarily). The equality $\eta(h)_z = h_z = \eta_n, z$ and the induction hypothesis ensures that z burns at time $m+1$ in $\eta(h)$ if and only if it burns in η_n . Let now $z \in C$, and let $z' \in A$ be the unique neighbour of z in A . Since $[z, z'] \in t_n$, z will burn in η_n at time $m+1$ if and only if z' burnt at time m . By the induction hypothesis, the latter occurs if and only if z' burnt in $\eta(h)$ at time m . Then by the definition $\eta(h)_z = 0$ we get that this happens if and only if z burns in $\eta(h)$ at time $m+1$. Finally, consider a vertex $z \in (\partial A) \setminus C$ that is unburnt at time m (in both configurations, necessarily). Let $z' \in A$ be its unique neighbour in A . Since $[z', z] \in t_n$, z' burns after z in η_n , and hence by the induction hypothesis z' is unburnt at time m in both η_n and $\eta(h)$. In η_n , z will burn at time $m+1$ if and only if $m+1 = m_2 = m_{n,z}$, and $z = y_{2,i}$ for some $1 \leq i \leq r_2$. Due to the definition $\eta(h)_z = 1$ and the fact that z' is unburnt in $\eta(h)$ at time m , this is equivalent to z burning in $\eta(h)$ at time $m+1$. This completes the induction.

We can now iterate the above argument until there are no more burnable vertices in $A \cup \partial A$, showing that the two burning processes are identical in $A \cup \partial A$.

The equality of the burning processes gives that if $\eta(h) \notin \mathcal{R}_{G_{A,C}}$, then there can be no tree with the given h , C and $m_{n,y}$'s. If $\eta(h) \in \mathcal{R}_{G_{A,C}}$, then there is exactly one possible arrow configuration in A , namely the one given by the burning of $\eta(h)$ in $G_{A,C}$ (here we use the same $\alpha_{P,K}$'s in the graphs $G_{A,C}$ and G_n). This completes the proof. \square

Proof of Lemma 8.2. Consider the following auxiliary weighted graph $G' = G'_{A,C}$. We add to the graph induced by $A \cup \partial A$ the vertex s , and the following edges: for any $y \in C$ there is an edge e_y between y and s with weight $w(e_y) = q_{n,y}(1 - q_{n,y})^{-1}$; and for any $y \in (\partial A) \setminus C$ there are edges $f_{y,1}$ and $f_{y,2}$ between y and s , with respective weights $w(f_{y,1}) = q_{n,y}(\sigma_y)(1 - q_{n,y})^{-1}$ and $w(f_{y,2}) = (q_{n,y} - q_{n,y}(\sigma_y))(1 - q_{n,y})^{-1}$. All other edges have weight 1.

Observe that the weights have been chosen in such a way that the probability for the network random walk started at $y \in \partial A$ to reach s before reaching A is $q_{n,y}$, the same as it was in G_n . Let S be a network random on G_n stopped at

time τ_n (the hitting time of s), and let S' be a network random walk on $G'_{A,C}$, stopped at the hitting time τ' of s . Let ξ_k be the time of the k -th visit by S to the set $A \cup \partial A \cup \{s\}$. The choice of weights implies that if $S(0) = S'(0) \in A \cup \partial A$, then $\{S(\xi_k)\}_{k \geq 0}$ has the same distribution as $\{S'(k)\}_{k \geq 0}$. We can couple the two walks in such a way that we have $S(\xi_k) = S'(k)$ for all $k \geq 0$. Moreover, by the choice of the weights, the coupling can be arranged in such a way that for each $y \in (\partial A) \setminus C$, the edge $f_{y,1}$ is the last edge traversed by S' if and only if $\text{LE}(S[\hat{\tau}_y, \tau_n]) = \sigma_y$, where $\hat{\tau}_y$ is the time of the last visit to y by S .

Let $a = |A|$, $b = |\partial A|$ and $c = |C|$. Let u_1, \dots, u_{a+b} be an enumeration of the vertices in $A \cup \partial A$, where $\{u_1, \dots, u_{b-c}\} = (\partial A) \setminus C$, $\{u_{b-c+1}, \dots, u_{b+a-c}\} = A$ and $\{u_{b+a-c+1}, \dots, u_{b+a}\} = C$. Let S^i and S'^i be network random walks on G_n and $G'_{A,C}$, respectively, with $S^i(0) = u_i = S'^i(0)$, coupled as above, and assume that these pairs are independent for $1 \leq i \leq a+b$. We use Wilson's method on G_n and $G'_{A,C}$ with the above enumeration of vertices and the coupled random walks to generate t_n distributed according to μ_n and t' distributed according to $\mu_{G'_{A,C}}$.

Our assumptions and Lemma 8.1 imply that there is a unique arrow configuration in A that realizes the event $\{\eta_{n,x} = h_x, x \in A\}$, given the restrictions $\mathcal{C}(t_n) = C$, $\pi_{n,y} = \sigma_y$. Let $\vec{F}_A := \{[v_1, v'_1], \dots, [v_a, v'_a]\}$ be this arrow configuration, where $v_j = u_{b-c+j}$. Observe that the edges leading from A to $(\partial A) \setminus C$ are always in \vec{F}_A , and hence we may assume without loss of generality that the indexing is such that $v'_1 = u_1, \dots, v'_{b-c} = u_{b-c}$. We define $\vec{F}_0 = \{f_{y,1} : y \in (\partial A) \setminus C\}$, where these edges are oriented away from ∂A , let $\vec{F}_{A,1} = \{[v_1, u_1], \dots, [v_{b-c}, u_{b-c}]\}$, let $\vec{F}_{A,2} = \vec{F}_A \setminus \vec{F}_{A,1}$, and let $\vec{F}_C := \{e_y : y \in C\}$, where these edges are oriented towards C . The coupling ensures that the event in (8.2) occurs if and only if t' consists of the edges:

$$\vec{F} := \vec{F}_0 \cup \vec{F}_A \cup \vec{F}_C. \quad (8.3)$$

In order to complete the proof, we need to show that the probability that Wilson's method on $G'_{A,C}$ produces \vec{F} is of the claimed form.

For $i = 1, \dots, b-c$, the conditional probability of the event $\{\mathcal{F}'_i = \mathcal{F}'_{i-1} \cup \{f_{u_i,1}\}\}$ given the event $\{\mathcal{F}'_{i-1} = \{f_{u_j,1} : 1 \leq j < i\}\}$ is of the form

$$f_{i,A,C}(q_{n,y} : y \in \partial A) q_{n,u_i}(\sigma_{u_i}),$$

where the form of $f_{i,A,C} : [0,1]^{\partial A} \rightarrow [0,1]$ only depends on the pair (A,C) . In particular,

$$\begin{aligned} p_0 &:= \mathbf{P}[\mathcal{F}_{b-c} = \vec{F}_0] \\ &= f'_{A,C}(q_{n,y} : y \in \partial A) \prod_{y \in (\partial A) \setminus C} q_{n,y}(\sigma_y), \end{aligned}$$

where the form of the function $f'_{A,C}$ only depends on the pair (A, C) . Let

$$\begin{aligned} p_{A,1} &:= \mathbf{P}[\mathcal{F}_{2(b-c)} = \vec{F}_0 \cup \vec{F}_{A,1} \mid \mathcal{F}_{b-c} = \vec{F}_0], \\ p_{A,2} &:= \mathbf{P}[\mathcal{F}_{b+a-c} = \vec{F}_0 \cup \vec{F}_{A,1} \cup \vec{F}_{A,2} \mid \mathcal{F}_{2(b-c)} = \vec{F}_0 \cup \vec{F}_{A,1}], \\ p_C &:= \mathbf{P}[t' = \vec{F} \mid \mathcal{F}_{b+a-c} = \vec{F}_0 \cup \vec{F}_A]. \end{aligned}$$

Here $p_{A,1} = f''_{A,C}$, where again the form of the function $f''_{A,C}$ only depends on the pair (A, C) , and $p_C = \prod_{y \in C} (1 - q_{n,y})$. We show that $p_{A,2}$ is a constant depending on (A, C) , and this will complete the proof. The operation of contracting an edge in a graph means identifying its endpoints to a single vertex. Let $G''_{A,C}$ denote the graph obtained from $G'_{A,C}$ by contracting all edges in $\vec{F}_0 \cup \vec{F}_{A,1} \cup \vec{F}_C$. Conditional on the event $\{\vec{F}_0 \cup \vec{F}_{A,1} \cup \vec{F}_C \subset t'\}$, the distribution of t' is equal to the distribution of $t'' \cup \vec{F}_0 \cup \vec{F}_{A,1} \cup \vec{F}_C$, where t'' is a sample from $\mu_{G''_{A,C}}$. Since all non-loop edges in $G''_{A,C}$ have weight 1, $\mu_{G''_{A,C}}$ is uniform on $\mathcal{T}_{G''_{A,C}}$. It follows that $p_{A,2} = |\mathcal{T}_{G''_{A,C}}|^{-1}$, that is a constant depending only on the pair (A, C) .

Since $\mathbf{P}[t' = \vec{F}] = p_0 p_{A,1} p_{A,2} p_C$, the proof of the lemma is complete. \square

Proof of Theorem 8.1. Let $C \subsetneq \partial A$ and suppose that the event $\{\eta_{n,x} = h_x, x \in A\}$ is realized by some tree t_n with $\mathcal{C}(t_n) = C$. For each $y \in (\partial A) \setminus C$, let σ_y be any self-avoiding path from y to s that avoids A . Lemma 8.1 shows that there is a unique arrow configuration \vec{F}_A in A (possibly depending on C and the σ_y 's), such that any tree t'_n with $\mathcal{C}(t'_n) = C$ such that t'_n contains the σ_y 's and \vec{F}_A also realizes the event $\{\eta_{n,x} = h_x, x \in A\}$. Lemma 8.2 shows that the probability mass of all such trees is given by the expression on the right hand side of (8.2). It follows that we have

$$\begin{aligned} \nu_{G_n}[\eta_{n,x} = h_x, x \in A] & \\ &= \sum_{C \subsetneq \partial A} \sum_{\{\sigma_y : y \in (\partial A) \setminus C\}} f_{A,C}(q_{n,y} : y \in \partial A) \prod_{y \in (\partial A) \setminus C} q_{n,y}(\sigma_y). \end{aligned} \tag{8.4}$$

It is crucial here that the second sum is over *all* collections self-avoiding paths from the y 's to s , avoiding A , and that $f_{A,C}$ is independent of the paths.

Observe that $\sum_{\sigma_y} q_{n,y}(\sigma_y) = q_{n,y}$, and therefore, performing the sum over the σ_y 's in (8.4) gives

$$\nu_{G_n}[\eta_{n,x} = h_x, x \in A] = \sum_{C \subsetneq \partial A} f_{A,C}(q_{n,y} : y \in \partial A) \prod_{y \in (\partial A) \setminus C} q_{n,y}.$$

Regardless of what the exhaustion is, we have

$$\lim_{n \rightarrow \infty} q_{n,y} = \mathbf{P}[S[0, \infty) \cap A = \emptyset \mid S(0) = y] = \frac{d-2}{d-1}.$$

Hence we have

$$\nu[\eta_{n,x} = h_x, x \in A] = \sum_{\substack{C \subseteq \partial A \\ C \neq \emptyset}} f_{A,C} \left(\frac{d-2}{d-1}, \dots, \frac{d-2}{d-1} \right) \left(\frac{d-2}{d-1} \right)^{|\partial A \setminus C|}.$$

□

Remark 8.1. As an example, taking $A = \{o\}$, one can recover the computation of height probabilities by Dhar and Majumdar [7] from the above.

9. Concluding remarks

9.1. Finiteness of avalanches

Following the program introduced by Maes, Redig and Saada [22], once the existence of the limit ν has been established, it is natural to ask if one can define sandpile dynamics on the infinite graph G . The first question is whether adding a particle at a vertex o in a sample configuration from the measure ν produces an avalanche that is finite ν -a.s. (that is, only finitely many topplings are necessary to stabilize it). In [16] a sufficient condition was given, in the case of \mathbb{Z}^d , $d \geq 3$, in terms of a certain modification $\tilde{\mu}$ of the measure μ . Let \tilde{G}_n be the graph obtained from G_n by wiring o to the sink, and let $\tilde{\mu} = \lim_{n \rightarrow \infty} \mu_{\tilde{G}_n}$ be the limiting wired spanning forest measure. Let t_o denote the component of o under the measure $\tilde{\mu}$. It was shown in [16] that if $\tilde{\mu}[|t_o| < \infty] = 1$, then avalanches on G are finite ν -a.s.

The arguments in [16] apply without change to show that on any transient graph, avalanches are ν -a.s. finite, if $\tilde{\mu}[|t_o| < \infty] = 1$. Lyons, Morris and Schramm [19] analyzed t_o under general conditions, in particular have shown that t_o is finite in any transitive graph with at least cubic volume growth. They have also shown that finiteness of t_o is equivalent to the one-end property. Since we assumed the one-end property in Theorems 6.1 and 7.2, it follows that the limiting measures constructed in these theorems have a.s. finite avalanches, if the graph is transient.

Let us also discuss the case of the d -regular tree. Let $o \in T^d$ be a fixed vertex. Take a sample configuration from the measure ν , and add a particle at o . Let N denote the number distinct vertices that topple in the stabilization of this configuration. The computations in [7] show that $\nu[N = n] \sim cn^{-3/2}$ as $n \rightarrow \infty$, and also that $\sum_{n=0}^{\infty} \nu[N = n] = 1$. Hence avalanches are ν -a.s. finite. As above, finiteness can also be derived from the well-known fact that each tree in the WSF has one end [10] (see also [3, Section 11]).

Open question 2. Are avalanches ν -a.s. finite on \mathbb{Z}^2 ? More generally, are avalanches ν -a.s. finite on a recurrent graph G such that the WSF has one end a.s.? See [15] for a related open question regarding a weaker property.

9.2. Stationary Markov process

Having established finiteness of avalanches, one can apply the general machinery developed in [22] to show the existence of a natural stationary Markov process with invariant measure ν . Suppose that $G = (V, E)$ is transient. Let $\varphi : V \rightarrow (0, \infty)$ be a function such that $\sum_{x \in V} \varphi(x)G(x, o) < \infty$, where $G = \Delta^{-1}$. Given an exhaustion $V_1 \subset V_2 \subset \dots \subset V$, consider the continuous time sandpile Markov chain on G_n , where particles are added at $x \in V_n$ at Poisson rate $\varphi(x)$. This Markov chain has invariant measure ν_{G_n} , and it follows from the results of [22] that its semigroup strongly converges in $L^2(\nu)$ to the semigroup of a Markov process with invariant measure ν .

Appendix

In this appendix we give the proof that the path-valued process $\{X_k\}_{k \geq 0}$ introduced in the proof of Lemma 7.4 is stationary and mixing. We will write $p(x, y)$ for the transition probability of S .

Let $[x_k(-m), \dots, x_k(-1), x_k(0), x_k(1), \dots, x_k(m)]$ be fixed finite paths in G , for $k = 0, \dots, K-1$, such that $x_k(0) = o$. Without loss of generality, we assume that $m > M$. Call two finite paths $y_1 = [y_1(-\ell_1), \dots, y_1(0) = o, \dots, y_1(\ell_2)]$ and $y_2 = [y_2(-\ell_1), \dots, y_2(0) = o, \dots, y_2(\ell_2)]$ *equivalent*, $y_1 \equiv y_2$, if there exists $\bar{\psi} \in \text{AUT}_o$ such that $\bar{\psi}y_1(j) = y_2(j)$ for each j . This is clearly an equivalence relation.

We will use the following simple lemma, whose proof is obvious.

Lemma A.1. *Suppose that T is a transformation from a finite set of paths \mathcal{P}_1 into a finite set of paths \mathcal{P}_0 , where \mathcal{P}_0 and \mathcal{P}_1 have the same number of elements. Suppose that T has the property that $y_1 \equiv y_2$ if and only if $Ty_1 \equiv Ty_2$. Let $\lambda_i : \mathcal{P}_i \rightarrow \mathbb{R}$, $i = 0, 1$ be functions that are constant on equivalence classes, such that $\lambda_1(y) = \lambda_0(Ty)$. Then $\sum_{y \in \mathcal{P}_1} \lambda_1(y) = \sum_{y' \in \mathcal{P}_0} \lambda_0(y)$.*

Consider the probability

$$\begin{aligned} & \mathbf{P} [X_k(j) = x_k(j), -m \leq j \leq m, 0 \leq k \leq K-1] & (A.1) \\ &= \mathbf{P} [\Psi_k \phi_{S(kM)}^{-1} S(kM + j) = x_k(j), -m \leq j \leq m, 0 \leq k \leq K-1] \\ &= \sum_{y'} \mathbf{P} [S(j) = y'(j), -m \leq j \leq (K-1)M + m] \\ & \quad \times \prod_{k=0}^{K-1} \mathbf{P} [\Psi_k \phi_{y'(kM)}^{-1} y'(kM + j) = x_k(j), -m \leq j \leq m], \end{aligned}$$

where the summation is over all paths y' with parameter set $\{-m, \dots, (K-1)M + m\}$ that are at o at time 0. The first factor in the right hand side

of (A.1) equals $\prod_{j=-m}^{(K-1)M+m-1} p(y'(j), y'(j+1))$. In order to abbreviate the second factor, introduce the notation $U_k y'(j) = \phi_{y'(kM)}^{-1} y'(kM+j)$, $-m \leq j \leq m$. Then the right hand side of (A.1) is

$$\lambda_0(y') := \prod_{j=-m}^{(K-1)M+m-1} p(y'(j), y'(j+1)) \prod_{k=0}^{K-1} \mathbf{P}[\Psi_k U_k y'(j) = x_k(j), -m \leq j \leq m].$$

Now consider

$$\begin{aligned} & \mathbf{P}[X_k(j) = x_{k-1}(j), -m \leq j \leq m, 1 \leq k \leq K] \\ &= \mathbf{P}[\psi_k \phi_{S(kM)}^{-1} S(kM+j) = x_{k-1}(j), -m \leq j \leq m, 1 \leq k \leq K] \\ &= \sum_y \mathbf{P}[S(j) = y(j), -m+M \leq j \leq KM+m] \\ &\quad \times \prod_{k=1}^K \mathbf{P}[\Psi_k \phi_{y(kM)}^{-1} y(kM+j) = x_{k-1}(j), -m \leq j \leq m] =: \sum_y \lambda_1(y), \end{aligned} \tag{A.2}$$

where the summation is over paths y with parameter set $\{-m+M, \dots, KM+m\}$ that are at o at time 0. Introduce the map T , $Ty(j) = \phi_{y(M)}^{-1} y(M+j)$. The first factor in the right hand side of (A.2) equals

$$\prod_{j=-m+M}^{KM+m-1} p(y(j), y(j+1)) = \prod_{j=-m}^{(K-1)M+m-1} p(Ty(j), Ty(j+1)). \tag{A.3}$$

The second factor equals

$$\prod_{k=0}^{K-1} \mathbf{P}[\Psi_{k+1} \phi_{y((k+1)M)}^{-1} y((k+1)M+j) = x_k(j), -m \leq j \leq m]. \tag{A.4}$$

We claim that the path $\{\phi_{y((k+1)M)}^{-1} y((k+1)M+j)\}_{j=-m}^m$ is equivalent to the path $U_k T y$. Indeed, it is easy to check that the automorphism

$$\tilde{\psi} = \phi_{Ty(kM)}^{-1} \phi_{y(M)}^{-1} \phi_{y((k+1)M)}$$

does the job. Hence, using right invariance of λ , the k th factor in (A.4) equals

$$\begin{aligned} \mathbf{P}[\Psi_{k+1} \tilde{\psi}^{-1} U_k T y = x_k] &= \lambda(\{\Psi : \Psi \tilde{\psi}^{-1} U_k T y = x_k\}) = \lambda(\{\Psi : \Psi U_k T y = x_k\} \tilde{\psi}) \\ &= \lambda(\{\Psi : \Psi U_k T y = x_k\}) = \mathbf{P}[\Psi_k U_k T y = x_k]. \end{aligned}$$

This and (A.3) shows that $\lambda_1(y) = \lambda_0(Ty)$. A similar computation shows that λ_0 is constant on equivalence classes.

It is left to show that $y_1 \equiv y_2$ if and only if $Ty_1 \equiv Ty_2$. Indeed, $y_2 = \bar{\psi}y_1$ if and only if $Ty_2 = \phi_{y_2(M)}^{-1}\bar{\psi}\phi_{y_1(M)}Ty_1$. An application of Lemma A.1 now shows that the expressions in (A.1) and (A.2) equal, and this is sufficient to conclude stationarity.

The proof of mixing can be carried out using a similar computation. Suppose that $(K-1)M+m < tM-m$, and consider the probability:

$$\begin{aligned} & \mathbf{P} [X_k(j) = x_k(j), -m \leq j \leq m, k \in \{0, \dots, K-1\} \cup \{t, \dots, t+K-1\}] \\ &= \sum_{y_0} \sum_y \sum_{y_t} p(y_0)p(y)p(y_t)\nu(y_0) \\ & \quad \times \prod_{k=0}^{K-1} \mathbf{P} [\Psi_{t+k}\phi_{y_t((t+k)M)}^{-1}y_t((t+k)M+j) = x_{t+k}(j), -m \leq j \leq m]. \end{aligned} \quad (\text{A.5})$$

Here the y_0 -sum is over paths with parameter set $\{-m, \dots, (K-1)M+m\}$ that are at o at time 0, the y -sum is over paths with parameter set

$$\{(K-1)M+m, \dots, tM-m\}$$

starting at $y_0((K-1)M+m)$, and the y_t -sum is over paths with parameter set $\{tM-m, \dots, (t+(K-1))M+m\}$ that start at $y(tM-m)$. The expressions $p(y_0)$, $p(y)$ and $p(y_t)$ stand for the products of random walk transition probabilities for these paths, and $\nu(y_0)$ is the expression containing $\Psi_0, \dots, \Psi_{K-1}$. Keeping y_0 and y fixed, we introduce the map T ,

$$Ty_t(i) = \phi_{y_t(tM)}^{-1}y_t(tM+i), \quad -m \leq i \leq (K-1)M+m.$$

Then T maps paths starting at $y(tM-m)$ to paths with parameter set

$$\{-m, \dots, (K-1)M+m\}$$

that are at o at time 0. A straightforward computation then shows that

$$\phi_{y_t((t+k)M)}^{-1}y_t((t+k)M+\cdot) \equiv U_kTy_t(\cdot).$$

An application of Lemma A.1 yields that (keeping y_0 and y fixed)

$$\begin{aligned} & \sum_{y_t} p(y_t) \prod_{k=0}^{K-1} \mathbf{P} [\Psi_{t+k}\phi_{y_t((t+k)M)}^{-1}y_t((t+k)M+j) = x_{t+k}(j), -m \leq j \leq m] \\ &= \sum_{y'} \tilde{\lambda}_0(y'), \end{aligned}$$

where $\tilde{\lambda}_0$ has the same form as λ_0 , with x_k replaced by x_{t+k} . We can now carry out the summation over y to yield a factor 1, and conclude that the expression

in (A.5) equals

$$\begin{aligned} & \mathbf{P} [X_k(j) = x_k(j), -m \leq j \leq m, 0 \leq k \leq K - 1] \\ & \times \mathbf{P} [X_{t+k}(j) = x_{t+k}(j), -m \leq j \leq m, 0 \leq k \leq K - 1]. \end{aligned} \quad (\text{A.6})$$

The equality in (A.6) is sufficient to conclude mixing.

Acknowledgments

I thank François Dunlop, Thierry Gobron and Ellen Saada for giving me the opportunity to present this talk. I am grateful for a question of David B. Wilson that prompted me to consider the problems in Section 8.

References

- [1] S.R. ATHREYA AND A.A. JÁRAI (2004) Infinite volume limit for the stationary distribution of Abelian sandpile models. *Commun. Math. Phys.* **249**, 197–213.
- [2] P. BAK, C. TANG AND K. WIESENFELD (1987) Self-organized criticality: an explanation of $1/f$ noise. *Phys. Rev. Lett.* **59**, 381–384.
- [3] I. BENJAMINI, R. LYONS, Y. PERES AND O. SCHRAMM (2001) Uniform spanning forests. *Ann. Prob.* **29**, 1–65.
- [4] B. BOLLOBÁS (1998) *Modern Graph Theory*. Graduate Texts in Mathematics **184**, Springer, New York.
- [5] D. DHAR (1990) Self-organized critical state of sandpile automaton models. *Phys. Rev. Lett.* **64**, 1613–1616.
- [6] D. DHAR (2006) Theoretical studies of self-organized criticality. *Physica A* **369**, 29–70.
- [7] D. DHAR AND S.N. MAJUMDAR (1990) Abelian sandpile model on the Bethe lattice. *J. Phys. A: Math. and Gen.* **23** 4333–4350.
- [8] E.B. DYNKIN AND M.B. MALJUTOV (1961) Random walk on groups with a finite number of generators. *Dokl. Akad. Nauk SSSR* **137**, 1042–1045 (in Russian). English translation: *Soviet Mathematics, Doklady* **2**, 399–402.
- [9] O. HÄGGSTRÖM (1995) Random-cluster measures and uniform spanning trees. *Stoch. Process. Appl.* **59**, 267–275.
- [10] O. HÄGGSTRÖM (1998) Uniform and minimal essential spanning forests on trees. *Random Struct. and Algorithms* **12**, 27–50.
- [11] M. HALL JR (1976) *The Theory of Groups*. 2nd edition. AMS Chelsea Publishing Co., New York.
- [12] A.E. HOLROYD, L. LEVINE, K. MÉSZÁROS, Y. PERES, J. PROPP AND D.B. WILSON (2008) Chip-firing and rotor-routing on directed graphs. In: *In and out of Equilibrium*. 2. V. Sidoravicius and M.E. Vares (eds.), Progr. Probab. **60**, Birkhäuser, Basel, 331–364.

- [13] A.A. JÁRAI (2005) Thermodynamic limit of the abelian sandpile model on \mathbf{Z}^d . *Markov Processes Relat. Fields* **11**, 313–336.
- [14] A.A. JÁRAI AND H. KESTEN (2004) A bound for the distribution of the hitting time of arbitrary sets by random walk. *Electron. Commun. Probab.* **9**, 152–161.
- [15] A.A. JÁRAI AND R. LYONS (2007) Ladder sandpiles. *Markov Processes Relat. Fields* **13**, 493–518.
- [16] A.A. JÁRAI AND F. REDIG (2008) Infinite volume limit of the abelian sandpile model in dimensions $d \geq 3$. *Probab. Theory and Relat. Fields* **141**, 181–212.
- [17] A.A. JÁRAI, F. REDIG AND E. SAADA (2010) Zero dissipation limit in the Abelian sandpile model. Preprint. [arXiv:0906.3128v2](https://arxiv.org/abs/0906.3128v2).
- [18] G.F. LAWLER (1996) *Intersections of Random Walks*. Birkhäuser, Boston.
- [19] R. LYONS, B.J. MORRIS AND O. SCHRAMM (2008) Ends in uniform spanning forests. *Electron. J. Probab.* **13**, 1702–1725.
- [20] R. LYONS, Y. PERES AND O. SCHRAMM (2003) Markov chain intersections and the loop-erased walk. *Ann. Inst. H. Poincaré, Probabilités et Statistiques* **39**, 779–791.
- [21] C. MAES, F. REDIG, E. SAADA AND A. VAN MOFFAERT (2000) On the thermodynamic limit for a one-dimensional sandpile process. *Markov Processes Relat. Fields* **6**, 121.
- [22] C. MAES, F. REDIG AND E. SAADA (2002) The abelian sandpile model on an infinite tree. *Ann. Probab.* **30**, 20812107.
- [23] C. MAES, F. REDIG AND E. SAADA (2004) The infinite volume limit of dissipative abelian sandpiles. *Commun. Math. Phys.* **244**, 395417.
- [24] C. MAES, F. REDIG AND E. SAADA (2005) Abelian sandpile models in infinite volume. *Sankhy* **67**, 634661 (2005).
- [25] S.N. MAJUMDAR AND D. DHAR (1992) Equivalence between the Abelian sandpile model and the $q \rightarrow 0$ limit of the Potts model. *Physica A* **185**, 129–145.
- [26] R. PEMANTLE (1991) Choosing a spanning tree for the integer lattice uniformly. *Ann. Probab.* **19**, 15591574.
- [27] F. REDIG (2006) Mathematical aspects of the abelian sandpile model. In: *Mathematical Statistical Physics: Lecture Notes of the Les Houches Summer School 2005*, A. Bovier, A. van Enter and F. Dunlop (eds.), Elsevier.
- [28] E.R. SPEER (1993) Asymmetric Abelian sandpile models. *J. Stat. Phys.* **71**, 61–74.
- [29] F. SPITZER (1976) *Principles of Random Walks*. 2nd edition. Graduate Texts in Mathematics **34**. Springer-Verlag, New York-Heidelberg.
- [30] D.B. WILSON (1996) Generating random spanning trees more quickly than the cover time. In: *Proceedings of the Twenty-Eighth ACM Symposium on the Theory of Computing*, ACM, New York, 296–303.
- [31] W. WOESS (2000) *Random Walks on Infinite Graphs and Groups*. Cambridge Tracts in Mathematics, **138**, Cambridge University Press, Cambridge.