isid/ms/2003/16
June 18, 2003
http://www.isid.ac.in/~statmath/eprints

Infinite volume limit for the stationary distribution of Abelian sandpile models

Siva R. Athreya
Antal A. Járai

Indian Statistical Institute, Delhi Centre
7, SJSS Marg, New Delhi–110 016, India

Infinite volume limit for the stationary distribution of Abelian sandpile models

Siva R. Athreya and Antal A. Járai

June 19, 2003

Abstract

We study the stationary distribution of the standard Abelian sandpile model in the box $\Lambda_n = [-n,n]^d \cap \mathbb{Z}^d$ for $d \geq 2$. We show that as $n \to \infty$, the finite volume stationary distributions weakly converge to a translation invariant measure on allowed sandpile configurations in $\mathbb{Z}^d$. This allows us to define infinite volume versions of the avalanche-size distribution and related quantities. The proof is based on a mapping of the sandpile model to the uniform spanning tree due to Dhar and Majumdar, and the existence of the wired uniform spanning forest measure on $\mathbb{Z}^d$. In the case $d > 4$ we also make use of Wilson's method.

1 Introduction

In this paper we study an infinite volume limit of the Abelian sandpile model (ASM) on $\mathbb{Z}^d$. The ASM was introduced in [2] as a simple model exhibiting self-organized criticality (SOC). The model has received substantial interest in the physics literature and in recent years in the mathematical literature as well. The concept of SOC was originally proposed in [1], [2].

The ASM is an interacting particle system (particles $\equiv$ "grains of sand") living on a finite subset Λ of the d-dimensional integer lattice $\mathbb{Z}^d$. Every site $i \in \Lambda$ is occupied by a number of particles $z_i \in \mathbb{N}^d$. If $1 \leq z_i \leq 2d$ then the site i is called *stable*, if $z_i > 2d$, it is called *unstable*. The value z_i is also called the *height* of the site. The value $z_c = 2d$ is called the *critical* height. The height configuration undergoes the following discrete-time dynamics. Given a configuration in which all sites are stable, we add a particle at a random site $i \in \Lambda$ which is chosen according to a distribution q, with $q(i) > 0$, $i \in \Lambda$. If as a result, i becomes unstable, $2d$ particles jump from site i, one to each adjacent site, decreasing the height of i by $2d$ and increasing the height at each nearest neighbor by 1. If the unstable site i was on the boundary of Λ, we still decrease the height of i by $2d$, and one or more particles leave the system through the boundary. This operation is called *toppling*. It may happen that new unstable sites are created by the toppling of i. We topple them as well, until eventually all sites become stable again. (One can show that any possible sequence of topplings leads to the same stable configuration.) This new stable configuration is the state of the system after a single time-step. The above description defines a Markov-chain with state space $\Omega_\Lambda = \{1, \dots, 2d\}^\Lambda$. It is known that there is a unique stationary distribution ν_Λ, independent of q [14]. In what follows we let q be the uniform measure and denote the law of the stationary chain by P_Λ.

For a definition of the ASM on a general graph, see [7, 14]. A thorough review of exactly solvable models exhibiting SOC is carried out in the lecture notes by Dhar [7]. A unified mean field study of SOC models, such as forest fires, can be found in [18]. Further background about SOC is provided by [12, 8] and the references therein. The first mathematical results about the ASM has been obtained in [4, 5, 6]. We refer the reader to [7] for background on basic facts regarding the ASM, and to [14] for a more detailed introduction with complete proofs.

The sequence of topplings performed in one time-step is called an *avalanche*. A basic problem is to determine the properties of avalanches under the stationary distribution ν_Λ. Some quantities of interest are (a) the number of topplings in an avalanche (size), (b) the number of sites affected by an avalanche (range), (c) the radius of the affected region (radius). These quantities are believed (and have numerically been observed) to have distributions falling off as a power. It is believed that the power law holds on some interval depending on the size of Λ, where the interval grows as $\Lambda \nearrow \mathbb{Z}^d$. To the best of our knowledge, there is no precise formulation or proof of such a statement, except when $d = 1$ [7].

As a step in analyzing the above question, we consider the limit $\Lambda \nearrow \mathbb{Z}^d$, and define avalanche characteristics in the infinite volume. In the two-dimensional case, Priezzhev [16] calculated the exact values of $\lim_{\Lambda \nearrow \mathbb{Z}^2} \nu_\Lambda(z_0 = k)$, $k = 1, \dots, 4$. In this paper we prove the more general, although less explicit, statement that ν_Λ converges weakly to a limit ν in dimensions $d \geq 2$ (see Theorem 1). Since the distributions of the quantities (a)–(c) above can be defined in terms of ν_Λ alone (without referring to the dynamics), we obtain that limiting distributions for (a)–(c) exist. It remains an important open problem to describe the limit in more detail (when it is non-trivial), and determine the effect of the boundary (which disappears in the limit). Recently, infinite volume versions of the sandpile process have been constructed on the one-dimensional lattice [12] and on an infinite tree [11]. Unlike in these articles, we do not construct a dynamics in the limit. However, our Theorem 1 is a necessary ingredient in such constructions.

Our proof is based on the deep observation of [13], that the set of recurrent states of the ASM can be mapped onto the set of spanning trees on Λ. This observation has also been used in [16, 17]. It is known that ν_Λ is the uniform measure on the set of recurrent states, and therefore ν_Λ corresponds to the uniform spanning tree measure on Λ. It is also known that the uniform spanning tree has a limit as $\Lambda \nearrow \mathbb{Z}^d$ [15, 3], called the uniform spanning forest (USF). Therefore it is not surprising that ν_Λ converges as well, and in fact, when $2 \leq d \leq 4$, a continuity property of the correspondence is indeed sufficient to prove this. However, in the case $d > 4$, due to the non-locality of the correspondence, making the argument precise requires effort. In particular, the correspondence breaks down in the infinite volume when $d > 4$, and a bit of extra randomness is necessary to describe the limit.

The rest of the paper is organized in the following manner. In the next section we state some basic notation and preliminaries. In the following Section 1.2 we state our main theorem and comment on its implications. Section 2 contains a review of the burning test and the connections of the ASM with the uniform spanning tree. Finally in Section 3, we provide a proof of Theorem 1, our main result.

Acknowledgment. We thank Akira Sakai and Bálint Tóth for useful and stimulating discussions.

Frank Redig shared with us ongoing work by him and his co-authors, and suggested a different proof in the case $2 \leq d \leq 4$. This work was initiated at the University of British Columbia. The research was supported in part by NSERC of Canada and the Pacific Institute for the Mathematical Sciences.

1.1 Notation and Preliminaries

When i and j are neighbors in $\mathbb{Z}^d$ we write $i \sim j$. We sometimes restrict our attention to graphs of the form $\Lambda_n = [-n,n]^d \cap \mathbb{Z}^d$, and write $\nu_n = \nu_{\Lambda_n}$, $P_n = P_{\Lambda_n}$ etc. We regard ν_Λ as a measure on the space $\Omega = \{1,\dots,2d\}^{\mathbb{Z}^d}$ in the natural way. We denote by $\mathcal{G}$ the natural σ-algebra on Ω.

By a cylinder event we mean an event in $\mathcal{G}$ depending on the states of finitely many sites only. For $v \in \mathbb{Z}^d$ let τ_v denote translation by v. If E is a cylinder event depending on a set of sites A, then $\tau_v E$ depends on the set of sites $\tau_v A = \{u + v : u \in A\}$. For a random variable Y we define $\tau_v Y$ similarly.

Given a quantity $f(\Lambda)$ defined for all (or all sufficiently large) finite subsets Λ of $\mathbb{Z}^d$ we say that

$$\lim_{\Lambda \nearrow \mathbb{Z}^d} f(\Lambda) = a,$$

if given any $\varepsilon > 0$ there is a finite $\Lambda_0 \subset \mathbb{Z}^d$ such that for all finite $\Lambda \supset \Lambda_0$ we have $|f(\Lambda) - a| < \varepsilon$.

1.2 Results

Our main result is concerned with the limit of ν_Λ as $\Lambda \nearrow \mathbb{Z}^d$. In its statement X_n denotes the position where the first particle is dropped in the (stationary) sandpile process on Λ_n.

Theorem 1. *Let $d \geq 2$. The measures ν_n weakly converge to a translation invariant measure ν on Ω. For any cylinder event E and any $v \in \mathbb{Z}^d$ we have*

$$\nu(E) = \lim_{n\to\infty} \nu_n(E) = \lim_{n\to\infty} \nu_n(\tau_v E) = \lim_{n\to\infty} P_n(\tau_{X_n} E). \tag{1}$$

Remark 1. (i) The first two limits in (1) exhibit the weak convergence and translation invariance. In (9) and (32) we give expressions for $\nu(E)$ in terms of the USF on $\mathbb{Z}^d$. The third equality in (1) says that in the stationary chain, the configuration at the position where the first (or any) particle is dropped has the same limiting law as at the origin. This is a consequence of translation invariance.

(ii) As mentioned earlier, there is a difference in the proof according to whether $2 \leq d \leq 4$ or $d > 4$. In the former case the USF is a.s. a single tree, and in this case the one-to-one correspondence between spanning trees and allowed configurations extends to configurations on $\mathbb{Z}^d$. When $d > 4$, the correspondence breaks down on $\mathbb{Z}^d$, due to the fact that the USF has multiple trees. However, the limit can still be described in terms of trees using extra randomness.

(iii) When $2 \leq d \leq 4$, we establish the first two limits even as $\Lambda \nearrow \mathbb{Z}^d$. We believe this stronger result to hold also when $d > 4$, but it was convenient to restrict to regular volumes at certain points in the proof. In the case of the third limit in (1) the restriction to volumes with regular boundary is necessary.

(iv) It is known that for any Λ the set of recurrent states can be characterized as those that do not contain any *forbidden sub-configurations* [14]. Since forbidden sub-configurations are finite, they do not occur in the limit ν-a.s.

(v) Given a configuration in Ω, it makes sense to talk about the size, range, radius, etc. of an avalanche when a particle is dropped at a fixed site, let's say the origin. Let S denote one of these quantities. Then Theorem 1 implies that $p(s) = \lim_{n\to\infty} \nu_n(S = s)$, $0 \le s < \infty$ is well defined, since the event $\{S = s\}$ is a cylinder event. By the third equality in (1), $p(s)$ also equals the limiting probability of $\{S = s\}$ when a particle is dropped at a random site X_n. In particular, when $S =$ avalanche size, $p(s)$ is the limiting avalanche-size distribution. It remains an open problem to determine whether $\sum_{s=0}^{\infty} p(s) < 1$ or $= 1$, or even whether $p(s) > 0$ for $s \ge 1$. We expect that the latter holds when $d \ge 3$, and fails when $d = 2$, and that transience/recurrence of random walk plays a role.

(vi) A detailed study of the case $d = 1$ has been done in [12], where an infinite volume dynamics has been constructed. We do not know whether a similar construction works when $d \ge 2$. A first step towards understanding the dynamics would be to determine whether there are infinite avalanches. This is not the case on an infinite regular tree, and this fact was important for the construction in [11].

2 Relation to the uniform spanning tree

Below we review the correspondence between the ASM and the uniform spanning tree [13], and then quote the necessary results about the USF.

2.1 The burning test

The following algorithm, called the "burning test" [13, 14], checks whether a configuration in Ω_Λ is recurrent. At the same time, it establishes a one-to-one map between recurrent configurations and spanning trees on a suitable modification of Λ. Define the graph $\widetilde{\Lambda}$ by adding a new site δ_Λ to Λ which is joined to each i in the boundary $\partial\Lambda$ by $2d - \deg(i)$ edges. Given a stable configuration, we set $A_0 = \{\delta_\Lambda\}$, and call A_0 the set of sites burning at time 0. For $t \ge 1$ we recursively define A_t (the set of sites burning at time t) as follows. Site i is burning at time t if its height is larger than the number of its unburnt neighbors.

In other words, for $j \in \Lambda$ let

$$n_t(j) = \#\{i \in \Lambda : i \sim j,\ i \notin \cup_{r=0}^{t-1} A_r\} \qquad \text{and} \qquad A_t = \{j \in \Lambda : z_j > n_t(j),\ j \notin \cup_{r=0}^{t-1} A_r\}.$$

Given a recurrent configuration $\mathbf{z} = (z_i)_{i\in\Lambda}$ we define a spanning tree $T_\Lambda = \phi(\mathbf{z})$ of $\widetilde{\Lambda}$ rooted at δ_Λ. We build the tree in such a way that A_t is the set of sites at graph distance t from the root. It is easy to see from the definitions that any site $j \in A_t$ has at least one neighbor in A_{t-1} $(t \ge 1)$. So to complete the definition of T_Λ, we only need to specify how to choose the parent of $j \in A_t$, when there are more than one neighbors in A_{t-1}. For this first observe that for $t \ge 1$ we have

$$j \in A_t \qquad \text{if and only if} \qquad n_{t-1}(j) \ge z_j > n_t(j), \tag{2}$$

where we set $n_0(j) \equiv z_c = 2d$. The number of possible parents of j, that is

$$r(j) = n_{t-1}(j) - n_t(j), \tag{3}$$

is therefore equal to the number of possible values of z_j that are allowed by (2). Thus we can choose the parent of j depending on the value of z_j in a one-to-one fashion according to some fixed rule.

The above algorithm produces a tree T_Λ, which spans $\widetilde{\Lambda}$ if and only if the sets $(A_t)_{t\geq 1}$ exhaust Λ. It is known that this happens if and only if $\mathbf{z}$ was recurrent [14].

The procedure can be reversed to show that ϕ is one-to-one and onto. We also describe ϕ^{-1} in detail. Given a spanning tree T_Λ, let B_t denote the set of sites at graph distance t from the root, $t \geq 0$. Let $m_t(j) = \#\{i : i \sim j,\ i \notin \cup_{r=0}^{t-1} B_r\}$. For any $j \in B_t$ the number of neighbors of j in B_{t-1} is $m_{t-1}(j) - m_t(j)$, and one of these neighbors is the parent of j. We set the value of z_j in such a way that for $j \in B_t$ the inequalities $m_{t-1}(j) \geq z_j > m_t(j)$ are satisfied, and we pick that value which corresponds to the parent of j according to our fixed rule. It is clear that the resulting configuration $\mathbf{z}$ is such that in the burning test $A_t = B_t$, $n_t(j) = m_t(j)$ and $\phi(\mathbf{z}) = T_\Lambda$.

Remark 2. (i) In order to reconstruct z_j, it is enough to know the distance of j from the root of T_Λ relative to the distances of its neighbors from the root. This usually allows one to reconstruct z_j knowing only a small portion of T_Λ. Let v denote the earliest common ancestor of all the neighbors of j (earliest means furthest from δ_Λ), and let F denote the subtree consisting of all descendants of v. We regard the site v as the root of F. The pair (F, v) already determines the value of z_j. This is because the distances of j and its neighbors from v in F give us the necessary information about $m_{t-1}(j)$ and $m_t(j)$, even without knowing for which t we have $j \in B_t$.

(ii) By the argument of (i), it is enough to know, in fact, the relative order of the distances from each neighbor of j to the root. This observation will play a role in the case $d > 4$.

Since all recurrent states have equal weight under ν_Λ, the image of ν_Λ under ϕ is uniform on all spanning trees of $\widetilde{\Lambda}$. It is called the uniform spanning tree on Λ with wired boundary conditions. We denote its law by μ_Λ. It is known (see Theorem 2 below) that as $\Lambda \nearrow \mathbb{Z}^d$, μ_Λ weakly converges to a limit called the wired uniform spanning forest [3]. We refer to the limit simply as the USF. (On $\mathbb{Z}^d$ the wired and free spanning forests coincide [3].)

2.2 Properties of the USF

The theorem below summarizes the results we need about the USF. For more background on spanning trees see [3]. In the statement of the theorem below, μ_Λ is the law of a random subset T_Λ of edges of $\mathbb{Z}^d$.

Theorem 2. *Let $d \geq 1$.*

(i) If B is any finite set of edges in $\mathbb{Z}^d$, and $B \subset \Lambda \subset \Lambda'$ with Λ' finite, then

$$\mu_\Lambda(B \subset T_\Lambda) \leq \mu_{\Lambda'}(B \subset T_{\Lambda'}). \tag{4}$$

(ii) For any finite sets $B \subset K$ of edges in $\mathbb{Z}^d$ the limit

$$\mu(T \cap K = B) \stackrel{\text{def}}{=} \lim_{\Lambda \nearrow \mathbb{Z}^d} \mu_\Lambda(T_\Lambda \cap K = B) \tag{5}$$

exists, and defines a translation invariant probability measure, called the USF.

(iii) The USF has no cycles μ-a.s. If $d \leq 4$, the USF is a single tree a.s. For $2 \leq d \leq 4$ the USF has one end a.s., meaning that any two infinite paths in T have infinitely many vertices in common.

(iv) If $d > 4$ then a.s. the USF has infinitely many components, each component is infinite and has a single end.

Proof. All statements are proved in [3]. In particular, (i) follows directly from [3, Corollary 4.3]. For the special case $K = B$, the existence of the limit in (5) follows from the monotonicity in (4). The general case $B \subset K$ follows by inclusion-exclusion. The statements (iii) and (iv) (with free boundary conditions) already appear in [15], except for the last statement in (iv), which is proved in [3]. $\square$

3 Proof of Theorem 1

3.1 The case $2 \leq d \leq 4$

As indicated earlier, the proof of Theorem 1 in this case is accomplished by exploiting the continuity of the correspondence between spanning trees and the sandpile model. In Section 3.2 we use a more concrete approach that would also apply here. We begin by listing some conventions and definitions.

1. It will be convenient to regard μ_Λ and μ (from Theorem 2) as measures on the space $\Omega' = \{0,1\}^{\mathbb{E}^d}$, where $\mathbb{E}^d$ denotes the set of all bonds of $\mathbb{Z}^d$, and 1 represents an edge being present. We consider Ω' with the metrizable product topology. For $\omega \in \Omega'$ let $\omega|_\Lambda$ denote the restriction of ω to edges joining vertices in Λ. Let $\mathcal{X} \subset \Omega'$ denote the set of spanning trees of $\mathbb{Z}^d$ with one end.

2. Let F be a finite rooted tree in $\mathbb{Z}^d$ with root x. F will be assumed to denote the edge set and $V(F)$ the vertex set. For a set of sites $B \subset V(F)$, we define $\mathrm{eca}(B;F)$ as the 'earliest common ancestor' of B in F. More formally, this can be described as the unique site furthest from x and common to all paths that start in B, end at x and stay in F. It may so happen that for certain B, $\mathrm{eca}(B;F) = x$. Let $\mathrm{desc}(B;F)$ denote the tree (or forest) consisting of all descendants of B in F.

3. We consider the sandpile configuration in a fixed finite set $A_0 \subset \mathbb{Z}^d$ for $\Lambda \supset A_0$. Let A denote the set of sites that are either in A_0 or have a neighbor in A_0. Let

$$\mathcal{F} = \mathcal{F}(A) = \left\{ (F,x) : \begin{array}{c} F \text{ is a finite rooted tree in } \mathbb{Z}^d \text{ with} \\ \text{root } x,\ A \subset V(F),\ \mathrm{eca}(A;F) = x \end{array} \right\}.$$

Given $(F,x) \in \mathcal{F}$, let $H_{F,x}$ denote the set of edges incident on a site in $V(F)$, excluding those edges incident on x that do not belong to F. In particular, $F \subset H_{F,x}$.

4. We write T for the USF, that is, $T \in \Omega'$ with distribution μ. If we define the "root" of T to be at infinity, we call $x^* = \mathrm{eca}(A; T)$ and $F^* = \mathrm{desc}(x^*; T)$.

5. We use the notation $H^*(\omega)$ for the set valued random variable whose value is $H_{F,x}$ on the event $\omega \cap H_{F,x} = F$ and $\mathbb{Z}^d$ otherwise. We also extend the definition of F^* and x^* whenever $H^*(\omega)$ is finite by letting $F^*(\omega) = F$, $x^*(\omega) = x$ on the event $\omega \cap H_{F,x} = F$.

Before we proceed to the proof, we observe the consistency of the above list. First, note that due to Theorem 2 (iii), $\mu(\mathcal{X}) = 1$. Similarly, $\mu_\Lambda(\mathcal{X}_\Lambda) = 1$ for the set $\mathcal{X}_\Lambda$ defined by

$$\mathcal{X}_\Lambda = \left\{ \omega \in \Omega' : \begin{matrix} \omega|_\Lambda \text{ has no cycles, and each component} \\ \text{of } \omega|_\Lambda \text{ is joined to } \Lambda^c \text{ by a unique edge} \end{matrix} \right\}.$$

Secondly, Theorem 2 (iii), ensures that (F^*, x^*) is μ-a.s. well-defined for $2 \le d \le 4$, and we have $(F^*, x^*) \in \mathcal{F}$. Thirdly, for different $(F, x) \in \mathcal{F}$, the events $\{\omega \cap H_{F,x} = F\}$ are disjoint, which implies that H^* is well defined. Finally, observe that

$$\{(F^*, x^*) = (F, x)\} = \{T \cap H_{F,x} = F\}, \qquad \mu\text{-a.s.}, \tag{6}$$

which means that the extended definition of F^* and x^* makes sense. We will assume the last observation for now and provide a proof at the end of this subsection.

Proof of Theorem 1. We observe that by Remark 2 (i), for $\omega \in \mathcal{X}_\Lambda$, the sandpile configuration in A_0 is already determined by (F^*, x^*), independently of Λ, when $H^* \subset \Lambda$. More precisely, defining the auxiliary space $\Omega_{A_0} = \{1, \dots, 2d\}^{A_0}$, the configuration is given in terms of a function $\psi : \mathcal{F} \to \Omega_{A_0}$. The correspondence in Section 2.1 can be recast in terms of functions $f_\Lambda, f : \Omega' \to \bar{\Omega}_{A_0}$ defined below. Let $\bar{\Omega}_{A_0} = \Omega_{A_0} \cup \{*\}$ (endowed with the discrete topology), and define

$$f_\Lambda(\omega) = \begin{cases} \phi_\Lambda^{-1}(\omega|_\Lambda)|_{A_0} & \omega \in \mathcal{X}_\Lambda, \\ * & \omega \in \Omega' \setminus \mathcal{X}_\Lambda. \end{cases}$$

$$f(\omega) = \begin{cases} \psi(F, x) & \text{when } H^*(\omega) = H_{F,x}, \\ * & \text{otherwise.} \end{cases}$$

By the observations above, for $\omega \in \mathcal{X}_\Lambda$ and $H^*(\omega) \subset \Lambda_k \subset \Lambda$ we have $f_\Lambda(\omega) = f(\omega) = \psi(F^*, x^*)$. This implies that for $u \in \Omega_{A_0}$

$$\limsup_{\Lambda \nearrow \mathbb{Z}^d} \int |I[f_\Lambda = u] - I[f = u]| \, d\mu_\Lambda \le \lim_{k \to \infty} \lim_{\Lambda \nearrow \mathbb{Z}^d} \mu_\Lambda(H^* \not\subset \Lambda_k) = 0. \tag{7}$$

Here in the last step we used that $\{H^* \subset \Lambda_k\}$ is a cylinder event, and that H^* is finite μ-a.s. It is easy to see using the definition of H^* that f is continuous at every $\omega \in \mathcal{X}$, and therefore by the general theory of weak convergence [10, Section 12]

$$\lim_{\Lambda \nearrow \mathbb{Z}^d} \int I[f = u] d\mu_\Lambda = \int I[f = u] d\mu. \tag{8}$$

Now (7) and (8) imply that for any $u \in \Omega_{A_0}$

$$\lim_{\Lambda \nearrow \mathbb{Z}^d} \nu_\Lambda(\mathbf{z}|_{A_0} = u) = \lim_{\Lambda \nearrow \mathbb{Z}^d} \int I[f_\Lambda = u] d\mu_\Lambda = \int I[f = u] d\mu \stackrel{\text{def}}{=} \nu(\mathbf{z}|_{A_0} = u).$$

This exhibits the weak convergence of ν_Λ to a limit ν. For a cylinder E depending on the set of sites A_0 we have

$$\nu(E) = \sum_{(F,x)\in\mathcal{F}_E} \mu(T \cap H_{F,x} = F), \tag{9}$$

where

$$\mathcal{F}_E = \{(F,x) \in \mathcal{F} : \psi(F,x) \in E\}. \tag{10}$$

Translation invariance of the limit follows, since for any fixed $v \in \mathbb{Z}^d$ we have

$$\lim_{\Lambda\nearrow\mathbb{Z}^d} \nu_\Lambda(\tau_v E) = \lim_{\Lambda\nearrow\mathbb{Z}^d} \nu_{\tau_{-v}\Lambda}(E) = \lim_{\Lambda\nearrow\mathbb{Z}^d} \nu_\Lambda(E) = \nu(E). \tag{11}$$

For the third equality in (1), observe that for fixed N and $n > N$

$$P_n(\tau_{X_n}E) = P_n(\tau_{X_n}E, \mathrm{dist}(X_n, \partial\Lambda_n) \le N) + \sum_{\substack{v\in\Lambda_n:\\ \mathrm{dist}(v,\partial\Lambda_n)>N}} \nu_n(\tau_v E)q(v).$$

The first term is bounded by $P_n(\mathrm{dist}(X_n, \partial\Lambda_n) \le N)$, and goes to 0 as $n \to \infty$. By (11) the second term is arbitrarily close to $\nu(E)$ when N is large, and $n \to \infty$.. $\square$

Proof of (6). First we show that $(F^*, x^*) = (F, x)$ implies the event on the right hand side. Since $F = F^*$, we have $F \subset T$, and hence $F \subset T \cap H_{F,x}$. Consider an edge $f = \langle u_1, u_2\rangle \in H_{F,x} \setminus F$, with $u_1 \in V(F)$. We show that $f \notin T$. In the case when we also have $u_2 \in V(F)$, we are done, since T has no cycles. If $u_2 \notin V(F)$, then first note that $u_1 \neq x$, by the definition of $H_{F,x}$. Therefore, if we had $f \in T$, then u_2 would be a descendant of x^* in T, and we would have $f \in F^* = F$, a contradiction.

Now assume that $T \cap H_{F,x} = F$ occurs. First, this implies $F \subset T$. It also implies, by the definition of $H_{F,x}$, that if an edge incident on any $u \in V(F)$ with $u \neq x$ does not belong to F, then it does not belong to T either. Hence the only site in $V(F)$ that is connected (in T) to infinity without using edges of F is x. This implies that $V(F)$ is precisely the set of descendants of x in T, and that F consists precisely of those edges of T that are descendants of x. It is simple to deduce from this that $x^* = x$ and $F^* = F$. $\square$

3.2 The case $d > 4$

We will be borrowing most of the definitions and conventions from the previous case. The few modifications we will make are due to the fact that there are multiple components in the USF.

1. We need to modify the definition of the set $\mathcal{F}$. We let

$$\bar{\mathcal{F}} = \bar{\mathcal{F}}(A) = \left\{ (F_i, x_i)_{i=1}^r : \begin{array}{l} F_i \text{ are vertex-disjoint finite rooted trees in} \\ \mathbb{Z}^d \text{ with root } x_i,\ \mathrm{eca}(A \cap V(F_i); F_i) = x_i, \\ i = 1, \dots, r, \text{ and } A \subset \cup_{i=1}^r V(F_i),\ r \ge 1 \end{array} \right\}.$$

   We write $(\bar{F}, \bar{x})$ to denote an element of $\bar{\mathcal{F}}$.

2. For $\Lambda \supset A$, recall T_Λ from Section 2.1. Since T_Λ falls apart into multiple components as $\Lambda \nearrow \mathbb{Z}^d$, any two fixed sites u and v are either connected within a 'short distance', or the connection occurs through the root δ_Λ. We decompose T_Λ into vertex disjoint trees by removing δ_Λ. With slight abuse of language, we refer to these trees as the components of T_Λ. The decomposition of T_Λ induces a decomposition of A into (random) sets A_i, $1 \le i \le r$, where $u, v \in A$ belong to the same A_i if and only if $\mathrm{eca}(\{u,v\}; T_\Lambda) \neq \delta_\Lambda$. Here r is random, and the indexing of the A_i's is determined by some fixed rule that assigns a particular indexing to any partition of A. We let $x^*_{\Lambda,i} = \mathrm{eca}(A_i; T_\Lambda)$ and $F^*_{\Lambda,i} = \mathrm{desc}(x^*_{\Lambda,i}; T_\Lambda)$. It is clear that $(\bar{F}^*_\Lambda, \bar{x}^*_\Lambda) = \{(F^*_{\Lambda,i}, x^*_{\Lambda,i})_{i=1}^r\} \in \bar{\mathcal{F}}$.

   It is straightforward to extend these definitions to the case $\Lambda = \mathbb{Z}^d$, noting that each component of the USF has one end a.s. Letting A_i, $1 \le i \le r$ denote the non-empty intersections of A with a component of T, we define $x^*_i = \mathrm{eca}(A_i; T)$ and $F^*_i = \mathrm{desc}(x^*_i; T)$. By Theorem 2 (iv), $(\bar{F}^*, \bar{x}^*) = (F^*_i, x^*_i)_{i=1}^r$ is μ-a.s. well-defined, and is an element of $\bar{\mathcal{F}}$.

3. Define
$$X_{\Lambda,i} = \mathrm{dist}_{T_\Lambda}(x^*_{\Lambda,i}, \delta_\Lambda), \quad 1 \le i \le r,$$
   where $\mathrm{dist}_{T_\Lambda}$ denotes the graph distance in T_Λ. Let Σ_r denote the set of permutations of $\{1, \dots, r\}$. We define the random permutation $\sigma^*_\Lambda \in \Sigma_r$ by the conditions $X_{\Lambda,\sigma^*_\Lambda(1)} \le \dots \le X_{\Lambda,\sigma^*_\Lambda(r)}$, where in case of ties we make a choice for σ^*_Λ in a fixed but arbitrary manner. We also define
$$Y_\Lambda = \min_{1 \le i < j \le r} |X_{\Lambda,i} - X_{\Lambda,j}|.$$
   For convenience, we set $Y_\Lambda = \infty$ when $r = 1$.

4. We need some more notation in order to formulate the analogue of (10). We define the events
$$D_\Lambda(\bar{x}) = D_\Lambda\big((x_i)_{i=1}^r\big) = \{x_1, \dots, x_r \text{ belong to distinct components of } T_\Lambda\},$$
   for $x_1, \dots, x_r \in \mathbb{Z}^d$, and
$$B_\Lambda(\bar{F}, \bar{x}) = D_\Lambda(\bar{x}) \cap \{T_\Lambda \cap H_{F_i,x_i} = F_i, 1 \le i \le r\}, \tag{12}$$
   for $(\bar{F}, \bar{x}) = (F_i, x_i)_{i=1}^r \in \bar{\mathcal{F}}$. When $\Lambda = \mathbb{Z}^d$, we denote the corresponding events by $D(\bar{x})$ and $B(\bar{F}, \bar{x})$. Analogously to the case $2 \le d \le 4$, we can show
$$\{(\bar{F}^*_\Lambda, \bar{x}^*_\Lambda) = (\bar{F}, \bar{x})\} = B_\Lambda(\bar{F}, \bar{x}) \qquad \{(\bar{F}^*, \bar{x}^*) = (\bar{F}, \bar{x})\} = B(\bar{F}, \bar{x}), \tag{13}$$
   for any $(\bar{F}, \bar{x}) \in \bar{\mathcal{F}}$ and $\Lambda \supset \cup_{i=1}^r H_{F_i,x_i}$.

Remark 3. Note that the events on the right hand side of (13) are disjoint for different $(\bar{F}, \bar{x}) \in \bar{\mathcal{F}}$. By Remark 2 (i), the occurrence or not of E is already determined by $(\bar{F}^*_\Lambda, \bar{x}^*_\Lambda)$ *and* $(X_{\Lambda,i})_{i=1}^r$. In fact, it is enough to know the value of all differences $X_{\Lambda,i} - X_{\Lambda,j}$, $1 \le i < j \le r$. Moreover, we can expect that the fluctuations of these differences grow as $\Lambda \nearrow \mathbb{Z}^d$, and for large Λ there exists a random permutation σ of $\{1, \dots, r\}$, such that $X_{\Lambda,\sigma(i+1)} - X_{\Lambda,\sigma(i)} \gg 1$, $1 \le i \le r-1$ with large probability. If the variables $X_{\Lambda,i}$ are in this sense 'well separated' (we will give a more precise definition shortly), then by Remark 2 (ii), already the permutation σ and $(\bar{F}^*_\Lambda, \bar{x}^*_\Lambda)$ determine the occurrence or not of E.

5. Fix $(\bar{F}, \bar{x}) \in \bar{\mathcal{F}}$ and $\sigma \in \Sigma_r$, where r is the number of components of $(\bar{F}, \bar{x})$. Assume that the events $B_\Lambda(\bar{F}, \bar{x})$ and $\{\sigma^*_\Lambda = \sigma\}$ occur. By the above consideration, this already determines whether

E occurs or not, independently of Λ, whenever Y_Λ is larger than some constant $K = K(\bar{F})$. We take $K(\bar{F}) = \max_{1\le i\le r} \operatorname{diam}(F_i)$, where $\operatorname{diam}(F)$ denotes the graph diameter of F. Let

$$\bar{\mathcal{F}}_E = \left\{ (\bar{F}, \bar{x}, \sigma) : \begin{array}{l} B_\Lambda(\bar{F}, \bar{x}) \text{ and } \sigma^*_\Lambda = \sigma \text{ imply } \phi^{-1}(T_\Lambda) \in E, \\ \text{whenever } Y_\Lambda > K(\bar{F}) \text{ and } \Lambda \supset \cup_{i=1}^r H_{F_i,x_i} \end{array} \right\}. \tag{14}$$

6. Let H^*_Λ denote the random set whose value equals $\cup_{i=1}^r H_{F_i,x_i}$ on the event $B_\Lambda(\bar{F}, \bar{x})$, $(\bar{F}, \bar{x}) \in \bar{\mathcal{F}}$.

We will need the following Lemma.

Lemma 3. *Let $d > 4$. We have*

$$\lim_{k\to\infty} \liminf_{\Lambda \nearrow \mathbb{Z}^d} \mu_\Lambda(H^*_\Lambda \subset \Lambda_k) = 1. \tag{15}$$

For fixed $(\bar{F}, \bar{x}, \sigma) \in \bar{\mathcal{F}}_E$, we have

$$\lim_{n\to\infty} \mu_n\big(B_n(\bar{F}, \bar{x}),\, Y_n \le K(\bar{F})\big) = 0, \tag{16}$$

and

$$\lim_{n\to\infty} \mu_n\big(B_n(\bar{F}, \bar{x}),\, \sigma^*_n = \sigma,\, Y_n > K(\bar{F})\big) = \frac{1}{r!}\mu(B(\bar{F}, \bar{x})). \tag{17}$$

Proof of Lemma 3. Denote by $x \leftrightarrow y$ the event that sites x and y belong to the same component of T_Λ (or T). The first step in showing (15) is to prove that for any $x, y \in A$

$$\limsup_{\Lambda \nearrow \mathbb{Z}^d} \mu_\Lambda(x \leftrightarrow y, \text{ but not inside } \Lambda_m) \to 0, \quad \text{as } m \to \infty. \tag{18}$$

To see this, note that when $x \leftrightarrow y$, there is a unique self-avoiding path $\omega : x \to y$ in Λ (or in $\mathbb{Z}^d$) such that the edges of ω belong to T_Λ (or T). Therefore, using (4), the expression in (18) can be bounded by

$$\limsup_{\Lambda \nearrow \mathbb{Z}^d} \sum_{\substack{\omega : x \to y \\ \omega \subset \Lambda \\ \omega \not\subset \Lambda_m}} \mu_\Lambda(\omega \subset T_\Lambda) \le \sum_{\substack{\omega : x \to y \\ \omega \not\subset \Lambda_m}} \mu(\omega \subset T) = \mu(x \leftrightarrow y, \text{ but not inside } \Lambda_m).$$

Here the right hand side goes to 0 as $m \to \infty$. Now assume that m is such that for any $x, y \in A$ either $x \leftrightarrow y$ inside Λ_m, or $x \not\leftrightarrow y$. Then $H^*_\Lambda \subset \Lambda_k$ can be ensured if $\operatorname{desc}(\Lambda_m; T_\Lambda) \subset \Lambda_{k-1}$. Since each component of T has a single end, for fixed m we have $\lim_{k\to\infty} \mu(\operatorname{desc}(\Lambda_m; T) \subset \Lambda_{k-1}) = 1$. By Theorem 2 (ii) this implies

$$\lim_{k\to\infty} \liminf_{\Lambda \nearrow \mathbb{Z}^d} \mu_\Lambda(\operatorname{desc}(\Lambda_m; T_\Lambda) \subset \Lambda_{k-1}) = 1.$$

This proves (15).

We next turn to the proof of (16). For a site x, let $Z_\Lambda(x) = \operatorname{dist}_{T_\Lambda}(x, \delta_\Lambda)$. Then it is sufficient to prove that for any $x, y \in A$

$$\lim_{n\to\infty} \mu_n(x \not\leftrightarrow y,\, |Z_n(x) - Z_n(y)| \le 2K(\bar{F})) = 0. \tag{19}$$

Indeed, $Y_\Lambda \le K(\bar{F})$ and the occurrence of $B_\Lambda(\bar{F}, \bar{x})$ would imply that there exist $x, y \in A$ such that $x \not\leftrightarrow y$, and $|Z_\Lambda(x) - Z_\Lambda(y)| \le 2K(\bar{F})$. Therefore we are going to study the paths from x and y to the boundary of Λ conditional on $x \not\leftrightarrow y$. The key tool for this is Wilson's method. It is described for example in [3, 19].

Wilson's method gives a construction of T_Λ via loop-erased random walks [9]. In particular, using the method with root at δ_Λ, it follows that the paths from x and y to δ_Λ can be generated in the following way. Let $\{S^{(i)}(n)\}_{n\geq 0}$, $i=1,2$ be two independent simple random walks starting at $S^{(1)}(0)=x$ and $S^{(2)}(0)=y$. Let $T^{(i)}=T^{(i)}_\Lambda$ be the hitting time of Λ^c by the two walks. Let $\mathcal{LE}$ denote the operation of erasing loops from a path in sequence, as they are created, and let $\gamma^{(i)}_\Lambda=\mathcal{LE}\left\{S^{(i)}[0,T^{(i)})\right\}$, $i=1,2$. Then conditional on $G_\Lambda=\{S^{(2)}[0,T^{(2)})\cap\gamma^{(1)}_\Lambda=\emptyset\}$, the joint law of $(\gamma^{(1)}_\Lambda,\gamma^{(2)}_\Lambda)$ is the same as the joint law of the paths in T_Λ from x and y to δ_Λ conditional on $x\not\leftrightarrow y$. In the sequel we assume that the latter paths have been generated by the random walks in this way. In particular, we assume that the constructions in different volumes Λ are coupled by using the same infinite random walks $S^{(1)}$ and $S^{(2)}$.

Denote by $\rho(n)$ the number of points remaining of the first n points after loops are erased from a random walk $S[0,\infty)$. It is shown in [9, Theorem 7.7.2], that for $d\geq 5$ there exists a constant $a>0$ such that

$$\lim_{n\to\infty}\frac{\rho(n)}{n}=a,\quad \text{a.s.}\tag{20}$$

We claim

$$Z_\Lambda(x)=\rho^{(1)}(T^{(1)}_\Lambda)+E_1,\tag{21}$$

where $E_1/T^{(1)}\to 0$ a.s. as $\Lambda\nearrow\mathbb{Z}^d$. We use the notion of a (two-sided) loop-free point, a concept introduced in [9]. A random walk $S[0,\infty)$ has a natural extension to a two-sided random walk $S(-\infty,\infty)$. We call a point j *loop-free* for S, if $S(-\infty,j]\cap S(j,\infty)=\emptyset$. If $j_0<j_1<j_2$ are loop-free, then loop-erasure on $[j_0,j_1]$ does not interfere with loop-erasure on $[j_1,j_2]$. Let

$$j^{(1)}_\Lambda=\max\{j<T^{(1)}_\Lambda : j \text{ is loop-free}\}.$$

Note that $E_1\leq T^{(1)}_\Lambda-j^{(1)}_\Lambda$. Let $\pi^{(1)}(n)$ denote the number of loop-free points in $[0,n)$ for the random walk $S^{(1)}$. By the ergodic theorem,

$$\lim_{n\to\infty}\frac{\pi^{(1)}(n)}{n}=b=P(0\text{ is loop-free})>0,\quad \text{a.s.}$$

This implies that for any $\delta>0$, as $\Lambda\nearrow\mathbb{Z}^d$ we have

$$(1+\delta)bj^{(1)}_\Lambda\geq\pi^{(1)}(j^{(1)}_\Lambda)=\pi^{(1)}(T^{(1)}_\Lambda)-1\geq(1-\delta)bT^{(1)}_\Lambda$$

eventually a.s. This implies that $(1-\delta)/(1+\delta)\leq j^{(1)}_\Lambda/T^{(1)}_\Lambda\leq 1$ eventually a.s., and therefore $E_1/T^{(1)}\to 0$ follows. Similarly to the above one can show that on the event G_Λ,

$$Z_\Lambda(y)=\rho^{(2)}(T^{(2)}_\Lambda)+E_2,\tag{22}$$

where $E_2/T^{(2)}\to 0$ as $\Lambda\nearrow\mathbb{Z}^d$ a.s.

It follows from (20), (21) and (22), that

$$\frac{Z_\Lambda(x)}{Z_\Lambda(y)}\frac{T^{(2)}_\Lambda}{T^{(1)}_\Lambda}\to 1,\quad \text{as } \Lambda\nearrow\mathbb{Z}^d \text{ a.s. on } G,\tag{23}$$

where $G=\left\{S^{(2)}[0,\infty)\cap\mathcal{LE}\left\{S^{(1)}[0,\infty)\right\}=\emptyset\right\}$. Since for $d\geq 5$ the walks $S^{(1)}[0,\infty)$ and $S^{(2)}[0,\infty)$ have finitely many intersections a.s. [9, Proposition 3.2.3], we have $\lim_{\Lambda\nearrow\mathbb{Z}^d}I[G_\Lambda]=I[G]$ a.s.

For simplicity, let us restrict to $\Lambda=\Lambda_n$, and consider $n\to\infty$.

Consider two independent Brownian motions in $\mathbb{R}^d$ started at 0, and let $\tau^{(i)}$, $i = 1, 2$ denote their first exit times from $(-1,1)^d$.

It follows from Donsker's theorem [10, Section 42.2] that

$$\frac{T_n^{(1)}}{T_n^{(2)}} \Rightarrow \frac{\tau^{(1)}}{\tau^{(2)}}, \quad \text{as } n \to \infty, \tag{24}$$

where $\Rightarrow$ denotes weak convergence. It is simple to deduce from (24), (23) and $I[G_n] \to I[G]$ that

$$\lim_{\delta\to 0} \limsup_{n\to\infty} P\left(G_n,\, \frac{Z_n(x)}{Z_n(y)} \in [1-\delta, 1+\delta]\right) = 0. \tag{25}$$

This in turn implies (19).

Finally, we show that a strengthening of the preceding argument also proves (17). For this we describe the event in (17) in terms of Wilson's algorithm. Enumerate the sites in $\cup_{i=1}^r V(F_i)$ starting with $x_1, \ldots, x_r$ and followed by an arbitrary list $y_1, y_2, \ldots$ of the rest of the sites. We apply Wilson's method with root δ_Λ and with paths starting successively at the sites enumerated above. Let $S^{(i)}$, $i = 1, \ldots, r$ be independent simple random walks started at x_i, with $T^{(i)}$ the hitting time of Λ^c. Let $\gamma_\Lambda^{(i)}$ be the loop-erasure of $S^{(i)}[0, T^{(i)})$ as before. For the event $D_\Lambda(\bar{x})$ we require the occurrence of the event

$$G_\Lambda = \left\{ S^{(i)}[0, T^{(i)}) \cap \left(\cup_{j=1}^{i-1} \gamma_\Lambda^{(j)}\right) = \emptyset,\, i = 1, \ldots, r \right\}. \tag{26}$$

In addition, $B_\Lambda(\bar{F}, \bar{x})$ gives conditions on the paths starting at $y_1, y_2, \ldots$, namely these paths have to realize the events $T_\Lambda \cap H_{F_i, x_i} = F_i$, given the paths $\gamma_\Lambda^{(i)}$. We denote the latter event by C_Λ. Thus C_Λ is a subevent of G_Λ, which occurs if and only if given the paths implicit in the event G_Λ, the loop-erased random walks started at $y_1, y_2, \ldots$ realize $T_\Lambda \cap H_{F_i, x_i} = F_i$. Analogously we can define events G and $C \subset G$, which are the $\Lambda = \mathbb{Z}^d$ versions of these events.

Applying Wilson's algorithm in $\mathbb{Z}^d$ with root at infinity, it is clear that $P(C) = \mu(B(\bar{F}, \bar{x}))$. Therefore, specializing to $\Lambda = \Lambda_n$ and using (25), (17) will be proved once we show

$$\lim_{n\to\infty} P\left(C_n,\, Z_n(x_{\sigma(i)}) < Z_n(x_{\sigma(i+1)}),\, i = 1, \ldots, r-1\right) = P(C)\frac{1}{r!}. \tag{27}$$

Arguing as in the proof of (16), we have

$$\lim_{n\to\infty} \frac{Z_n(x_{\sigma(i)})}{Z_n(x_{\sigma(i+1)})} \frac{T_n^{\sigma(i+1)}}{T_n^{\sigma(i)}} = 1, \quad \text{a.s. on } G,\, i = 1, \ldots, r-1. \tag{28}$$

Since $C \subset G$, the above convergence also holds a.s. on C.

Next we show $I[C_\Lambda] \to I[C]$ a.s. We may assume the occurrence of G, since as before we already know $I[G_\Lambda] \to I[G]$. When C occurs as well, the random walks started at $y_1, y_2, \ldots$ remain inside a finite (random) box upto their respective hitting times. This implies that C_Λ occurs for large enough n. If G occurs but C does not, then two things can happen. One is that for some j the random walk started at y_j has infinite hiting time. In this case C_Λ cannot occur. The other is that all hitting times are finite, but at least one of the events $T \cap H_{F_i, x_i} = F_i$ is not realized. When this happens, it also happens for all large Λ, and thus C_Λ does not occur.

By the previous paragarph, we can replace C_n by C in (27) without affecting the limit. Also, by (28) and (25) we can replace each Z_n by the corresponding hitting time without affecting the limit. Therefore we are left to show

$$\lim_{n\to\infty} P\left(C, T_n^{\sigma(i)} < T_n^{\sigma(i+1)},\, i=1,\ldots,r-1\right) = P(C)\frac{1}{r!}, \tag{29}$$

We complete the proof by approximating C by C_m for $0 < m < n$, keeping m fixed but large. The probability on the left hand side of (29) can be written as

$$P\left(C_m, T_n^{\sigma(i)} < T_n^{\sigma(i+1)},\, i=1,\ldots,r-1\right) + \eta(m,n), \tag{30}$$

where $\lim_{m\to\infty}\limsup_{n\to\infty}\eta(m,n) = 0$ by $I[C_m]\to I[C]$. Also, we can replace $T_n^{\sigma(j)}$ by $T_n^{\sigma(j)} - T_m^{\sigma(j)}$ in (30). For fixed m, C_m and $(T_n^{(i)} - T_m^{(i)})_{i=1}^r$ are conditionally independent given $(S^{(i)}(T_m^{(i)}))_{i=1}^r$. Similarly to (24) we have

$$\frac{T_n^{(i)} - T_m^{(i)}}{T_n^{(j)} - T_m^{(j)}} \Rightarrow \frac{\tau^{(i)}}{\tau^{(j)}}, \quad 1 \le i < j \le r,$$

uniformly in $(S^{(i)}(T_m^{(i)}))_{i=1}^r$ for fixed m as $n\to\infty$. This gives

$$\begin{aligned}\lim_{n\to\infty} P\left(C_m, T_n^{\sigma(i)} < T_n^{\sigma(i+1)},\, i=1,\ldots,r-1\right) &= P(C_m)P\left(\tau^{\sigma(i)} < \tau^{\sigma(i+1)}, i=1,\ldots,r-1\right)\\ &= P(C_m)\frac{1}{r!}.\end{aligned} \tag{31}$$

Since $P(C_m)\to P(C)$, (30) and (31) proves (29) by letting $m\to\infty$. This completes the proof of the lemma. $\square$

We are now ready to present the proof of the Theorem.

Proof of Theorem 1. We write down an expression for the limit $\nu(E)$. In the lemma we have shown that conditioned on $B_\Lambda(\bar F,\bar x)$, σ_Λ^* is asymptotically uniform on Σ_r. Therefore we define

$$\nu(E) \stackrel{\text{def}}{=} \sum_{(\bar F,\bar x,\sigma)\in\bar{\mathcal F}_E} \frac{1}{r!}\mu(B(\bar F,\bar x)), \tag{32}$$

where the value of r in the summand is the number of components of $\bar F$.

For $k \ge 1$ we let

$$\begin{aligned}\bar{\mathcal F}(k) &= \left\{(F_i,x_i)_{i=1}^r \in \bar{\mathcal F} :\ \cup_{i=1}^r H_{F_i,x_i} \subset \Lambda_k\right\}\\ \bar{\mathcal F}_E(k) &= \left\{(\bar F,\bar x,\sigma)\in\bar{\mathcal F}_E : (\bar F,\bar x)\in\bar{\mathcal F}(k)\right\}.\end{aligned}$$

By the discussions preceding (14), we have

$$\begin{aligned}&\left|\nu_n(E) - \sum_{(\bar F,\bar x,\sigma)\in\bar{\mathcal F}_E(k)} \mu_n\big(B_n(\bar F,\bar x),\, \sigma_n^* = \sigma,\, Y_n > K(\bar F)\big)\right|\\ &\le \mu_n\big(H^* \not\subset \Lambda_k\big) + \sum_{(\bar F,\bar x,\sigma)\in\bar{\mathcal F}_E(k)} \mu_n\big(B_n(\bar F,\bar x),\, Y_n \le K(\bar F)\big).\end{aligned} \tag{33}$$

Let $\varepsilon > 0$, and choose k large, so that the lim sup of the first term on the right hand side is at most ε, as $n\to\infty$. Fixing this k, and noting that $\bar{\mathcal F}_E(k)$ is finite, the second term on the right hand side of (33) is less than ε, if n is large enough, by (16). Also, for each $(\bar{F}, \bar{x}) \in \bar{\mathcal{F}}_E(k)$, the summand on the left hand side of (33) approaches $\mu(B(\bar{F}, \bar{x}))/r!$ by (17). Now letting $\varepsilon \to 0$ proves that $\lim_{n\to\infty} \nu_n(E) = \nu(E)$.

For the second limit in Theorem 1, we can apply the same argument, using a minor modification of Lemma 3. Note that the convergence is in fact uniform in v, as long as the distance of v from the boundary is at least αn, for any fixed $\alpha > 0$. To see this, first note that under this condition, no problem arises where we have shown convergence as $\Lambda \nearrow \mathbb{Z}^d$. Therefore we only need to verify that there is uniformity in the application of Donsker's theorem as well.

To make the last observation more precise, let $z = (z_1, \dots, z_d) \in [-(1-\alpha), 1-\alpha]^d$, and let $T_n(z)$ be the exit time from $[-n, n]^d$ for simple random walk started at $v = nz$. Then $T_n(z)/(2dn^2) \Rightarrow \tau(z)$ where $\tau(z)$ is the exit time from $(-1, 1)^d$ for Brownian motion started at z. What we need to verify is that for any $t > 0$

$$P(T_n(z)/(2dn^2) < t) \to P(\tau(z) < t) \quad \text{uniformly in } z. \tag{34}$$

Let $(S_n)_{n\geq 0} = (S_{n,1}, \dots, S_{n,d})_{n\geq 0}$ be simple random walk started at 0.The event on the left hand side of (34) can be recast as

$$\bigcup_{i=1}^{d} \left[\left\{ \max_{0\leq m\leq 2dn^2 t} \frac{S_{m,i}}{n} \geq 1 - z_i \right\} \cup \left\{ \min_{0\leq m\leq 2dn^2 t} \frac{S_{m,i}}{n} \leq -1 - z_i \right\} \right].$$

Thus the claim follows from the weak convergence of the joint law of the maxima and minima in this event.

With this observation we can prove the third equality of the theorem arguing similarly to the case $2 \leq d \leq 4$, and letting $\alpha \to 0$. □

References

[1] Bak, P., Tang, C. and Wiesenfeld, K., *Self-organized criticality: An explanation of the 1/f noise* Phys. Rev. A **59** (1987) 381-384.

[2] Bak, P., Tang, C. and Wiesenfeld, K., *Self-organized criticality.* Phys. Rev. A **38** (1988) 364–374.

[3] Benjamini, I., Lyons, R., Peres, Y. and Schramm, O.: *Uniform spanning forests.* Ann. Prob.**29** (2001) 1–65.

[4] Dhar, D. and Ramaswamy, R.: *Exactly solved model of self-organized critical phenomena.* Phys. Rev. Lett. **63** (1989) 1659–1662.

[5] Dhar, D.: *Self-organized critical state of sandpile automaton models.* Phys. Rev. Lett. **64** (1990) 1613–1616.

[6] Dhar, D. and Majumdar, S.N.: *Abelian sandpile models on the Bethe lattice.* J. Phys. A **23** (1990) 4333–4350.

[7] Dhar, D.: *Studying Self-organized criticality with exactly solved models. preprint* `http://arXiv.org/abs/cond-mat/9909009` (1999).

[8] Jensen, H.J.: Self-organized criticality. Emergent complex behavior in physical and biological systems. (Cambridge Lecture Notes in Physics, 10) Cambridge University Press (2000).

[9] Lawler, G.F.: Intersections of random walks. Birkhäuser, softcover edition (1996).

[10] Loève, M.: Probability theory I–II. (Graduate Texts in Mathematics, 45–46) Springer-Verlag, 4th edition (1977).

[11] Maes, C., Redig, F. and Saada, E.: *The Abelian sandpile model on an infinite tree.* Ann. Probab. **30** (2002) 2081–2107.

[12] Maes, C., Redig, F., Saada, E. and Van Moffaert, A.: *On the thermodynamic limit for a one-dimensional sandpile process.* Markov Process. Related Fields **6** (2000) 1–22.

[13] Majumdar, S.N. and Dhar, D.: *Equivalence between the Abelian sandpile model and the $q \to 0$ limit of the Potts model.* Physica **A 185** (1992) 129–145.

[14] Meester, R., Redig, F. and Znamenski, D.: *The Abelian sandpile; a mathematical introduction.* Markov Proccess. Related Fields **7** (2002) 509–523.

[15] Pemantle, R.: *Choosing a spanning tree for the integer lattice uniformly.* Ann. Prob.**19** (1991) 1559–1574.

[16] Priezzhev, V.B.: *Structure of two-dimensional sandpile. I. Height Probabilities.* J. Stat. Phys.**74** (1994) 955–979.

[17] Priezzhev, V.B.: *The upper critical dimension of the Abelian sandpile model.* J. Stat. Phys.**98** (2000) 667–684.

[18] Vespignani, A. and Zapperi S.: *How Self-organised criticality works: A unified mean-field picture* Phys. Rev. A **57** (1988) 6345–6361.

[19] Wilson, D.B.: *Generating random spanning trees more quickly than the cover time.* In Proceedings of the Twenty-Eighth ACM Symposium on the Theory of Computing (1996) 296–303. ACM, New York.

Siva R. Athreya,
7 SJSS marg, Indian Statistical Institute, New Delhi, 110016, India. e-mail: `athreya@isid.ac.in.`

Antal A. Járai,
CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands. e-mail: `jarai@cwi.nl`