



Centrum voor Wiskunde en Informatica

REPORTRAPPORT

PNA

Probability, Networks and Algorithms



Probability, Networks and Algorithms

Infinite volume limit for the stationary distribution of
Abelian sandpile models

Siva R. Athreya, Antal A. Járai

REPORT PNA-E0304 DECEMBER 8, 2003

CWI is the National Research Institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organization for Scientific Research (NWO).

CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

Copyright © 2003, Stichting Centrum voor Wiskunde en Informatica

P.O. Box 94079, 1090 GB Amsterdam (NL)

Kruislaan 413, 1098 SJ Amsterdam (NL)

Telephone +31 20 592 9333

Telefax +31 20 592 4199

ISSN 1386-3711

Infinite volume limit for the stationary distribution of Abelian sandpile models

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 $n \rightarrow \infty$, 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 Majumdar and Dhar, 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.

2000 Mathematics Subject Classification: 60K35; 82C22

Keywords and Phrases: Abelian sandpile model; BTW model; self-organized criticality; thermodynamic limit

Note: This work was carried out under project PNA3 - "Dynamic percolation phenomena near criticality". A slightly different version of this work will appear in Communications in Mathematical Physics.

Infinite volume limit for the stationary distribution of Abelian sandpile models

Siva R. Athreya and Antal A. Járai

November 10, 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 $n \rightarrow \infty$, 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 Majumdar and Dhar, 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

The Abelian sandpile model (ASM) was introduced by Bak, Tang and Wiesenfeld [2] as a model exhibiting self-organized criticality (SOC). Roughly speaking, SOC arises when a dynamics drives a system towards a stationary state characterized by power law correlations in space and time. The concept of SOC was proposed in [1, 2], as a mechanism that could explain the occurrence of fractal structures in diverse natural phenomena. Various physical situations where the concept may apply are discussed in the book [11]. The ASM is one of the simplest models in which the complex phenomenon of SOC can be studied. Due to its rich mathematical structure and tractability, the model has received substantial interest in the physics literature and in recent years in the mathematical literature as well; see the review papers [9, 6] and [14].

The Abelian sandpile 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 . In finite volume Λ , the model is defined as follows.

Every site $i \in \Lambda$ is occupied by a number of particles $z_i \in \{1, 2, \dots\}$. 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_Λ , with $q_\Lambda(i) > 0$, $i \in \Lambda$. 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*, and it can be concisely written as $z_j \rightarrow z_j - \Delta_{ij}$, where Δ is the discrete Laplacian in Λ ,

$$\Delta_{ij} = \begin{cases} 2d & \text{if } i = j, \\ -1 & \text{if } |i - j| = 1, \\ 0 & \text{otherwise.} \end{cases}$$

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. The order in which we do the topplings does not matter. One can show that any possible sequence of topplings leads to the same stable configuration [5, 20]. This new stable configuration is the state of the system after a single time-step. The result of particle addition at i and subsequent relaxation is given by an operator $a_i : \Omega_\Lambda \rightarrow \Omega_\Lambda$, where $\Omega_\Lambda = \{1, \dots, 2d\}^\Lambda$. Due to the random choice of i , we have a Markov-chain with state space Ω_Λ . The operators a_i commute (hence the name Abelian), which makes it possible to analyze the chain in some detail. In particular, there is a unique stationary distribution ν_Λ , which is uniform on the set of recurrent states of the Markov chain, and is independent of q_Λ [5, 20]. For this reason, it is quite natural to fix q_Λ to be the uniform measure. We note that the above definitions and results carry over to a general graph [6, 20].

The first mathematical results about the ASM, including the statements above, were proved by Dhar, see [4, 5, 8]. Additional background is provided by [6, 9]. For a detailed introduction to the basic properties of the model we refer the reader to [20]. 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, including sandpile and forest fire models, can be found in [25]. Further background about SOC is provided by [11, 6, 9] and the references therein.

The main object of study in the model is the sequence of topplings performed in one time-step, 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), and
- (c) the distance of the furthest affected site from the initial toppling (radius).

It is often assumed that these quantities have distributions with a power law tail in the limit $\Lambda \nearrow \mathbb{Z}^d$. Numerical results in $d = 2$ indicate a rich fractal and multi-fractal structure of the distributions of (a) and (b) [24]. Also, it has been argued that above the upper critical dimension $d_u = 4$, the probability of an avalanche of size s decays like $s^{-3/2}$ (again in the large volume limit) [23]. To the best of our knowledge, there is no rigorous proof of power law behavior, either in $d = 2$ or higher. Exact computations are possible for $d = 1$ [7] and on the Bethe lattice [8]. In the former case, the probability of an avalanche of size s occurring goes to 0 for fixed s . In the latter, the probability of an avalanche of size s is asymptotic to a multiple of $s^{-3/2}$ (See (6.14) in [8]).

As a step in analyzing the above distributions, in this paper we study some aspects of the limit $\Lambda \nearrow \mathbb{Z}^d$, and define avalanche characteristics in the infinite volume. In the two-dimensional case, Priezzhev [22] calculated the exact values of $\lim_{\Lambda \nearrow \mathbb{Z}^2} \nu_\Lambda(z_0 = k)$, $k = 1, \dots, 4$. By an idea of Majumdar and Dhar [18, 17], it is possible to compute, in principle, the limiting probability of any finite height configuration that satisfies a certain minimality property. In this paper, we prove that ν_Λ converges weakly to a limit ν in dimensions $d \geq 2$ (see Theorem 1), which implies the existence of the thermodynamic limit of the full height configuration in the stationary state. 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, and to determine the effect of the boundary in finite volumes. Recently, infinite volume versions of the sandpile process have been constructed on the one-dimensional lattice [16], on an infinite tree [14], and for a dissipative model [15]. 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 Majumdar and Dhar [19], 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 [22, 23]. 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 $\Lambda \nearrow \mathbb{Z}^d$ [21, 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 $d > 4$, the correspondence becomes non-local, and making the argument precise requires effort. The non-locality is due to the fact that the uniform spanning forest has infinitely many components when $d > 4$. As a consequence, the correspondence between sandpile configurations and trees breaks down in the infinite volume when $d > 4$, and a bit of extra randomness is necessary to describe the limit. This leads to the extra permutation in (19) of Lemma 3.

The rest of the paper is organized as follows. 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 our main result.

Acknowledgment. This work was started at the University of British Columbia where both authors were postdoctoral fellows and the research was supported in part by NSERC of Canada and the Pacific Institute for the Mathematical Sciences. We thank: Akira Sakai and Bálint Tóth for useful and stimulating discussions; Frank Redig for sharing with us ongoing work by him and his co-authors, and for suggesting the current general argument in the proof of the case $2 \leq d \leq 4$; the referees, whose detailed suggestions have made the article comprehensive and up to date; Deepak Dhar and the referees for pointing out an erroneous remark in the first version, their input is incorporated in Remark 1 (vi).

1.1 Notation and Preliminaries

We let P_Λ denote the product of the measures ν_Λ and q_Λ . We think of P_Λ as the joint law of the stationary height configuration and the position of particle dropping. We write X_Λ for the random site of Λ specified by q_Λ . We sometimes restrict our attention to volumes 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 the natural σ -algebra on Ω by \mathcal{G} .

By a cylinder event we mean an event in \mathcal{G} depending on the heights 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 function $f(\Lambda)$ taking values in a metric space with metric ρ , and 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 $\rho(f(\Lambda), a) < \varepsilon$.

When i and j are neighbors in \mathbb{Z}^d we denote this by $i \sim j$.

1.2 Results

Our main result is concerned with the limit of ν_Λ as $\Lambda \nearrow \mathbb{Z}^d$. In its statement, we assume that X_n satisfies

$$\lim_{\varepsilon \downarrow 0} \limsup_{n \rightarrow \infty} P_n(\text{dist}(X_n, \partial \Lambda_n) \leq \varepsilon n) = 0. \quad (1)$$

This condition is clearly satisfied when X_n is uniform 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

$$\nu(E) = \lim_{n \rightarrow \infty} \nu_n(E) = \lim_{n \rightarrow \infty} \nu_n(\tau_v E) = \lim_{n \rightarrow \infty} P_n(\tau_{X_n} E). \quad (2)$$

Remark 1. (i) The first two limits in (2) exhibit the weak convergence and translation invariance. In (10) and (34) we give expressions for $\nu(E)$ in terms of the USF on \mathbb{Z}^d . The third equality in (2), which is a consequence of translation invariance, says that the configuration at the position X_n has the same limiting law as at 0. See remark (v) below.

(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) For $2 \leq d \leq 4$, we establish the first two limits even as $\Lambda \nearrow \mathbb{Z}^d$. In $d > 4$, the first two limits hold for growing regions of the form $\Lambda_n = (nG) \cap \mathbb{Z}^d$, where G is an open set in \mathbb{R}^d with smooth boundary. We believe the former 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 (2), the restriction to volumes with regular boundary is necessary, if we want condition (1) to apply when X_n is uniform. In the case $2 \leq d \leq 4$, our proof allows us to relax condition (1) to $\lim_{n \rightarrow \infty} P_n(\text{dist}(X_n, \partial \Lambda_n) \leq N) = 0$ for any $N > 0$.

(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* [20, 5]. Since forbidden sub-configurations are finite, they do not occur in the limit, and hence ν is supported on allowed height configurations.

(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 $p(s) = \lim_{n \rightarrow \infty} \nu_n(S = s)$, $0 \leq s < \infty$ is well defined, since the event $\{S = s\}$ is a cylinder event. By the third equality in (2), $p(s)$ also equals the limiting probability of $\{\tau_{X_n} S = s\}$ when a particle is dropped at a random site X_n . In particular, when $S =$ avalanche size, $p(s)$ is the asymptotic avalanche-size distribution. It remains an open problem to determine whether $\sum_{s=0}^{\infty} p(s) < 1$ or $= 1$, the latter case being equivalent to the absence of infinite avalanches. The absence of infinite avalanches for $d > 4$ will be investigated in [10].

(vi) It is possible to show that $p(s) > 0$ for $s \geq 0$ when $d \geq 2$. To see this, we give an explicit finite configuration C_s such that $\nu(C_s) > 0$, and C_s produces an avalanche of size s on addition at 0. For $s = 0$, C_0 consists of a single 1 at the origin. Let e_1, \dots, e_d denote the coordinate vectors. For $s \geq 1$, we consider a string of s sites $i_k = (k-1)e_1$ ($1 \leq k \leq s$), and we set $z_{i_k} = 2d$. We also set $z_{i_1 - e_1} = 1$;

$z_{i_1 \pm e_j} = 1$ ($2 \leq j \leq d$); $z_{i_k \pm e_j} = 2$ ($2 \leq j \leq d$, $2 \leq k \leq s$); and $z_{i_s + e_1} = 2$. These values make up the configuration C_s . We denote the support of C_s by A . It is easy to check that C_s is allowed (does not contain any forbidden sub-configurations), and produces an avalanche of size s if we add at 0. Also, C_s is minimal, in the sense that decreasing any value creates a forbidden sub-configuration. By the technique of Majumdar and Dhar [18, 17], recurrent configurations in Λ containing C_s are in one-to-one correspondence with recurrent configurations containing C_s in a modified graph Λ' (with toppling matrix Δ'). In Λ' , the set A is connected by the single edge $\{i_s + e_1, i_s + 2e_1\}$ to the rest of the lattice. This gives $\nu_\Lambda(C_s) = \nu_{\Lambda'}(C_s) \det(\Delta') / \det(\Delta)$. The ratio of the two determinants can be evaluated in terms of the Green function by the method of [18]. In any case, the ratio remains strictly positive in the limit $\Lambda \nearrow \mathbb{Z}^d$ (one way to see this is by counting spanning trees, and using Wilson's algorithm [3]). In the cases $s = 0, 1$, one has $\nu_{\Lambda'}(C_s) = 1$, and for $s \geq 2$, $\nu_{\Lambda'}(C_s)$ is a positive number that only depends on A . These observations imply $\nu(C_s) > 0$.

(vii) In [18], the authors compute the correlation between the events that sites 0 and x (respectively) have height 1, in the large volume limit. Their computation directly implies that

$$\nu(z_0 = 1, z_x = 1) - \nu(z_0 = 1)\nu(z_x = 1) \sim |x|^{-2d}, \quad \text{as } |x| \rightarrow \infty.$$

That is, under ν , at least the random field $I[z_x = 1]$ has power law correlations.

(viii) It is natural to ask if one can define dynamics in the infinite volume. This question has been addressed in the one-dimensional case [16], for the Bethe lattice [14] and for a dissipative model [15]. In the last two cases, the absence of infinite avalanches was an important ingredient (see remark (v)). Construction of infinite volume dynamics for $d > 4$ will be addressed in [10]. There the authors will also investigate ergodic properties of ν , based on tail triviality of the USF [3].

2 Relation to the uniform spanning tree

Below we review the correspondence between the ASM and the uniform spanning tree [19], and then quote the necessary results about the USF.

2.1 The burning test

The following algorithm, called the ‘‘burning test’’ [5, 19, 20], 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 $\tilde{\Lambda}$ 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 $t \geq 1$ 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

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

Given a recurrent configuration $\mathbf{z} = (z_i)_{i \in \Lambda}$ we define a spanning tree $T_\Lambda = \phi(\mathbf{z})$ of $\tilde{\Lambda}$ 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} ($t \geq 1$). Therefore, 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 $t \geq 1$ we have

$$j \in A_t \quad \text{if and only if} \quad n_{t-1}(j) \geq z_j > n_t(j), \quad (3)$$

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

$$r(j) = n_{t-1}(j) - n_t(j), \quad (4)$$

is therefore equal to the number of possible values of z_j that are allowed by (3). 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 $\tilde{\Lambda}$ 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 [19, 20].

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 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 key 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 $\tilde{\Lambda}$. 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. 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. The theorem was proved by Pemantle [21], except for an extension proved in [3]. 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*

$$\mu_\Lambda(B \subset T_\Lambda) \leq \mu_{\Lambda'}(B \subset T_{\Lambda'}). \quad (5)$$

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

$$\mu(T \cap K = B) \stackrel{\text{def}}{=} \lim_{\Lambda \nearrow \mathbb{Z}^d} \mu_\Lambda(T_\Lambda \cap K = B) \quad (6)$$

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, except for the last statement of (iv), are either proved in [21] or are implicitly present there. The last statement of (iv) is proved in [3], and proofs of the other statements can be found there as well. In particular, (i) follows directly from [3, Corollary 4.3]. For the special case $K = B$, the existence of the limit in (6) follows from the monotonicity in (5). The general case $B \subset K$ follows by inclusion-exclusion. Statements (iii) and (iv) follow from [3, Corollary 9.6], and [3, Theorem 10.1] \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 $\text{eca}(B; F)$ 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 , $\text{eca}(B; F) = x$. Let $\text{desc}(B; F)$ 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

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

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 $x^* = \text{eca}(A; T)$ and $F^* = \text{desc}(x^*; T)$.
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}_Λ defined by

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

Secondly, Theorem 2 (iii), ensures that (F^*, x^*) is μ -a.s. well-defined for $2 \leq d \leq 4$, 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

$$\{(F^*, x^*) = (F, x)\} = \{T \cap H_{F,x} = F\}, \quad \mu\text{-a.s.}, \quad (7)$$

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 $\psi : \mathcal{F} \rightarrow \Omega_{A_0}$. The correspondence in Section 2.1 can be recast in terms of functions $f_\Lambda, f : \Omega' \rightarrow \bar{\Omega}_{A_0}$ 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 $u \in \Omega_{A_0}$

$$\limsup_{\Lambda \nearrow \mathbb{Z}^d} \int |I[f_\Lambda = u] - I[f = u]| d\mu_\Lambda \leq \lim_{k \rightarrow \infty} \lim_{\Lambda \nearrow \mathbb{Z}^d} \mu_\Lambda(H^* \not\subset \Lambda_k) = 0. \quad (8)$$

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 [13, Section 12]

$$\lim_{\Lambda \nearrow \mathbb{Z}^d} \int I[f = u] d\mu_\Lambda = \int I[f = u] d\mu. \quad (9)$$

Now (8) and (9) 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

$$\nu(E) = \sum_{(F,x) \in \mathcal{F}_E} \mu(T \cap H_{F,x} = F), \quad (10)$$

where

$$\mathcal{F}_E = \{(F, x) \in \mathcal{F} : \psi(F, x) \in E\}. \quad (11)$$

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

$$\lim_{\Lambda \nearrow \mathbb{Z}^d} \nu_\Lambda(\tau_v E) = \lim_{\Lambda \nearrow \mathbb{Z}^d} \nu_{\tau_{-v}\Lambda}(E) = \lim_{\Lambda \nearrow \mathbb{Z}^d} \nu_\Lambda(E) = \nu(E). \quad (12)$$

For the third equality in (2), observe that for fixed N and $n > N$

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

The first term is bounded by $P_n(\text{dist}(X_n, \partial\Lambda_n) \leq N)$, and goes to 0 as $n \rightarrow \infty$. By (12) the second term is arbitrarily close to $\nu(E)$ when N is large, and $n \rightarrow \infty$. \square

Proof of (7). 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

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

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 , $1 \leq i \leq r$, where $u, v \in A$ belong to the same A_i if and only if $\text{eca}(\{u, v\}; T_\Lambda) \neq \delta_\Lambda$. 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 $x_{\Lambda, i}^* = \text{eca}(A_i; T_\Lambda)$ and $F_{\Lambda, i}^* = \text{desc}(x_{\Lambda, i}^*; T_\Lambda)$. 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 , $1 \leq i \leq r$ denote the non-empty intersections of A with a component of T , we define $x_i^* = \text{eca}(A_i; T)$ and $F_i^* = \text{desc}(x_i^*; T)$. 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

$$X_{\Lambda, i} = \text{dist}_{T_\Lambda}(x_{\Lambda, i}^*, \delta_\Lambda), \quad 1 \leq i \leq r,$$

where dist_{T_Λ} 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 $X_{\Lambda, \sigma_\Lambda^*(1)} \leq \dots \leq X_{\Lambda, \sigma_\Lambda^*(r)}$, where in case of ties we make a choice for σ_Λ^* in a fixed but arbitrary manner. We also define

$$Y_\Lambda = \min_{1 \leq i < j \leq r} |X_{\Lambda, i} - X_{\Lambda, j}|.$$

For convenience, we set $Y_\Lambda = \infty$ when $r = 1$.

4. We need some more notation in order to formulate the analogue of (11). We define the events

$$D_\Lambda(\bar{x}) = D_\Lambda((x_i)_{i=1}^r) = \{x_1, \dots, x_r \text{ belong to distinct components of } T_\Lambda\},$$

for $x_1, \dots, x_r \in \mathbb{Z}^d$, and

$$B_\Lambda(\bar{F}, \bar{x}) = D_\Lambda(\bar{x}) \cap \{T_\Lambda \cap H_{F_i, x_i} = F_i, 1 \leq i \leq r\}, \quad (13)$$

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 $2 \leq d \leq 4$, we can show

$$\{(\bar{F}_\Lambda^*, \bar{x}_\Lambda^*) = (\bar{F}, \bar{x})\} = B_\Lambda(\bar{F}, \bar{x}) \quad \{(\bar{F}^*, \bar{x}^*) = (\bar{F}, \bar{x})\} = B(\bar{F}, \bar{x}), \quad (14)$$

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 (14) 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 $(\bar{F}_\Lambda^*, \bar{x}_\Lambda^*)$ and the value of all differences $X_{\Lambda, i} - X_{\Lambda, j}$, $1 \leq i < j \leq r$. In view of Remark 2 (ii), even less information about the $X_{\Lambda, i}$ is sufficient. The configuration in A_0 is determined by the relative order of the distances $\text{dist}_{T_\Lambda}(w, \delta_\Lambda)$ for $w \in A$. For $w \in A_i$ we have

$$\text{dist}_{T_\Lambda}(w, \delta_\Lambda) = \text{dist}_{F_{\Lambda, i}^*}(w, x_{\Lambda, i}^*) + X_{\Lambda, i}.$$

Therefore, the relative order within the i -th component only depends on $(F_{\Lambda,i}^*, x_{\Lambda,i}^*)$. To determine the relative order between $w_1 \in A_i$ and $w_2 \in A_j$, $i \neq j$, we need to consider

$$\text{dist}_{T_\Lambda}(w_1, \delta_\Lambda) - \text{dist}_{T_\Lambda}(w_2, \delta_\Lambda) = [\text{dist}_{F_{\Lambda,i}^*}(w_1, x_{\Lambda,i}^*) - \text{dist}_{F_{\Lambda,j}^*}(w_2, x_{\Lambda,j}^*)] + [X_{\Lambda,i} - X_{\Lambda,j}]. \quad (15)$$

We can expect that the fluctuations of the differences $X_{\Lambda,i} - X_{\Lambda,j}$ grow as $\Lambda \nearrow \mathbb{Z}^d$, and that the second term on the right hand side of (15) will dominate, and determine the order. For this, it is in fact enough if $Y_\Lambda > \max_{1 \leq i \leq r} \text{diam}(F_{\Lambda,i}^*)$, where diam denotes the graph diameter. When this happens, we say that the $X_{\Lambda,i}$ are ‘well separated’. If the $X_{\Lambda,i}$ are well separated, then by Remark 2 (ii), already the permutation σ_Λ^* defined above 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 $K(\bar{F}) = \max_{1 \leq i \leq r} \text{diam}(F_i)$. Let

$$\bar{\mathcal{F}}_E = \left\{ (\bar{F}, \bar{x}, \sigma) : \begin{array}{l} B_\Lambda(\bar{F}, \bar{x}) \text{ and } \sigma_\Lambda^* = \sigma \text{ imply } \phi^{-1}(T_\Lambda) \in E, \\ \text{whenever } Y_\Lambda > K(\bar{F}) \text{ and } \Lambda \supset \cup_{i=1}^r H_{F_i, x_i} \end{array} \right\}. \quad (16)$$

The family $\bar{\mathcal{F}}_E$ collects those spanning tree configurations and permutations, that contribute to the event E , given that the $X_{\Lambda,i}$ are well-separated. It will be part of Lemma 3 below to show that configurations with $Y_\Lambda \leq K(\bar{F})$ do not contribute in the limit; see (18).

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*

$$\lim_{k \rightarrow \infty} \liminf_{\Lambda \nearrow \mathbb{Z}^d} \mu_\Lambda(H_\Lambda^* \subset \Lambda_k) = 1. \quad (17)$$

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

$$\lim_{n \rightarrow \infty} \mu_n(B_n(\bar{F}, \bar{x}), Y_n \leq K(\bar{F})) = 0, \quad (18)$$

and

$$\lim_{n \rightarrow \infty} \mu_n(B_n(\bar{F}, \bar{x}), \sigma_n^* = \sigma, Y_n > K(\bar{F})) = \frac{1}{r!} \mu(B(\bar{F}, \bar{x})). \quad (19)$$

Due to (18), the event $Y_n > K(\bar{F})$ in the last statement could be omitted, without affecting the limit. However, it is instructive to keep it in for its use in the proof of Theorem 1; see (35).

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 (17) is to prove that for any $x, y \in A$

$$\limsup_{\Lambda \nearrow \mathbb{Z}^d} \mu_\Lambda(x \leftrightarrow y, \text{ but not inside } \Lambda_m) \rightarrow 0, \quad \text{as } m \rightarrow \infty. \quad (20)$$

To see this, note that when $x \leftrightarrow y$, there is a unique self-avoiding path $\omega : x \rightarrow y$ in Λ (or in \mathbb{Z}^d) such that the edges of ω belong to T_Λ (or T). Therefore, using (5), the expression in (20) can be bounded by

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

Here the right hand side goes to 0 as $m \rightarrow \infty$. 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 $\text{desc}(\Lambda_m; T_\Lambda) \subset \Lambda_{k-1}$. Since each component of T has a single end, for fixed m we have $\lim_{k \rightarrow \infty} \mu(\text{desc}(\Lambda_m; T) \subset \Lambda_{k-1}) = 1$. By Theorem 2 (ii) this implies

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

This proves (17).

We next turn to the proof of (18). For a site x , let $Z_\Lambda(x) = \text{dist}_{T_\Lambda}(x, \delta_\Lambda)$. Then it is sufficient to prove that for any $x, y \in A$

$$\lim_{n \rightarrow \infty} \mu_n(x \not\leftrightarrow y, |Z_n(x) - Z_n(y)| \leq 2K(\bar{F})) = 0. \quad (21)$$

Indeed, $Y_\Lambda \leq K(\bar{F})$ 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 $|Z_\Lambda(x) - Z_\Lambda(y)| \leq 2K(\bar{F})$. 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, 26].

Wilson's method gives a construction of T_Λ via loop-erased random walks [12]. 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 $T^{(i)} = T_\Lambda^{(i)}$ 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 $\gamma_\Lambda^{(i)} = \mathcal{LE}\{S^{(i)}[0, T^{(i)}]\}$, $i = 1, 2$. Then conditional on $G_\Lambda = \{S^{(2)}[0, T^{(2)}] \cap \gamma_\Lambda^{(1)} = \emptyset\}$, the joint law of $(\gamma_\Lambda^{(1)}, \gamma_\Lambda^{(2)})$ 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 [12, Theorem 7.7.2], that for $d \geq 5$ there exists a constant $a > 0$ such that

$$\lim_{n \rightarrow \infty} \frac{\rho(n)}{n} = a, \quad \text{a.s.} \quad (22)$$

We claim

$$Z_\Lambda(x) = \rho^{(1)}(T_\Lambda^{(1)}) + E_1, \quad (23)$$

where $E_1/T^{(1)} \rightarrow 0$ a.s. as $\Lambda \nearrow \mathbb{Z}^d$. We use the notion of a (two-sided) loop-free point, a concept introduced in [12, Lemma 7.7.1]. 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

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

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

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

where in the last step we used $d \geq 5$ [12]. This implies that for any $\delta > 0$, as $\Lambda \nearrow \mathbb{Z}^d$ we have

$$(1 + \delta)b j_\Lambda^{(1)} \geq \pi^{(1)}(j_\Lambda^{(1)}) = \pi^{(1)}(T_\Lambda^{(1)}) - 1 \geq (1 - \delta)b T_\Lambda^{(1)}$$

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

$$Z_\Lambda(y) = \rho^{(2)}(T_\Lambda^{(2)}) + E_2, \quad (24)$$

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

It follows from (22), (23) and (24), that

$$\frac{Z_\Lambda(x) T_\Lambda^{(2)}}{Z_\Lambda(y) T_\Lambda^{(1)}} \rightarrow 1, \quad \text{as } \Lambda \nearrow \mathbb{Z}^d \text{ a.s. on } G, \quad (25)$$

where $G = \{S^{(2)}[0, \infty) \cap \mathcal{LE}\{S^{(1)}[0, \infty)\} = \emptyset\}$. Since for $d \geq 5$ the walks $S^{(1)}[0, \infty)$ and $S^{(2)}[0, \infty)$ have finitely many intersections a.s. [12, 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 $n \rightarrow \infty$.

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 [13, Section 42.2] that

$$\frac{T_n^{(1)}}{T_n^{(2)}} \Rightarrow \frac{\tau^{(1)}}{\tau^{(2)}}, \quad \text{as } n \rightarrow \infty, \quad (26)$$

where \Rightarrow denotes weak convergence. It is simple to deduce from (26), (25) and $I[G_n] \rightarrow I[G]$ that

$$\lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} P\left(G_n, \frac{Z_n(x)}{Z_n(y)} \in [1 - \delta, 1 + \delta]\right) = 0. \quad (27)$$

This in turn implies (21), since $Z_n(x), Z_n(y) \rightarrow \infty$ as $n \rightarrow \infty$. Note that the fact that $Z_n(x)/Z_n(y)$ does not ‘concentrate mass at 1’ provides the proof that the probability of $Y_n \leq K(\bar{F})$ vanishes.

Finally, we show that a strengthening of the preceding argument also proves (19). For this we describe the event in (19) in terms of Wilson’s algorithm. Enumerate the sites in $\cup_{i=1}^r V(F_i)$ starting with x_1, \dots, x_r and followed by an arbitrary list y_1, y_2, \dots 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 $S^{(i)}, i = 1, \dots, r$ 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

$$G_\Lambda = \left\{ S^{(i)}[0, T^{(i)}) \cap \left(\cup_{j=1}^{i-1} \gamma_\Lambda^{(j)} \right) = \emptyset, i = 1, \dots, r \right\}. \quad (28)$$

In addition, $B_\Lambda(\bar{F}, \bar{x})$ gives conditions on the paths starting at y_1, y_2, \dots , 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 sub-event 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, \dots 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 G_Λ and C_Λ .

Applying Wilson’s algorithm in \mathbb{Z}^d with root at infinity, it is clear that $P(C) = \mu(B(\bar{F}, \bar{x}))$. As before, (27) takes care of the condition $Y_n > K(\bar{F})$ in (19). Therefore, specializing to $\Lambda = \Lambda_n$ and using (27), (19) will be proved, once we show

$$\lim_{n \rightarrow \infty} P(C_n, Z_n(x_{\sigma(i)}) < Z_n(x_{\sigma(i+1)}), i = 1, \dots, r-1) = P(C) \frac{1}{r!}. \quad (29)$$

Arguing as in the proof of (18), we have

$$\lim_{n \rightarrow \infty} \frac{Z_n(x_{\sigma(i)})}{Z_n(x_{\sigma(i+1)})} \frac{T_n^{\sigma(i+1)}}{T_n^{\sigma(i)}} = 1, \quad \text{a.s. on } G, i = 1, \dots, r-1. \quad (30)$$

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

Next we show $I[C_\Lambda] \rightarrow I[C]$ a.s. We may assume the occurrence of G , since as before, we already know $I[G_\Lambda] \rightarrow I[G]$. When C occurs, the random walks started at y_1, y_2, \dots remain inside a finite (random) box up to their respective hitting times. This implies that C_Λ occurs for large enough Λ . 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 hitting 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 paragraph, we can replace C_n by C in (29) without affecting the limit. Also, by (30) and (27) we can replace each Z_n by the corresponding hitting time without affecting the limit. Therefore we are left to show

$$\lim_{n \rightarrow \infty} P\left(C, T_n^{\sigma(i)} < T_n^{\sigma(i+1)}, i = 1, \dots, r-1\right) = P(C) \frac{1}{r!}, \quad (31)$$

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 (31) can be written as

$$P\left(C_m, T_n^{\sigma(i)} < T_n^{\sigma(i+1)}, i = 1, \dots, r-1\right) + \eta(m, n), \quad (32)$$

where $\lim_{m \rightarrow \infty} \limsup_{n \rightarrow \infty} \eta(m, n) = 0$ by $I[C_m] \rightarrow I[C]$. Also, we can replace $T_n^{\sigma(j)}$ by $T_n^{\sigma(j)} - T_m^{\sigma(j)}$ in (32). 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 (26) we have

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

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

$$\begin{aligned} \lim_{n \rightarrow \infty} P\left(C_m, T_n^{\sigma(i)} < T_n^{\sigma(i+1)}, i = 1, \dots, r-1\right) &= P(C_m)P\left(\tau^{\sigma(i)} < \tau^{\sigma(i+1)}, i = 1, \dots, r-1\right) \\ &= P(C_m) \frac{1}{r!}. \end{aligned} \quad (33)$$

Since $P(C_m) \rightarrow P(C)$, (32) and (33) proves (31) by letting $m \rightarrow \infty$. 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

$$\nu(E) \stackrel{\text{def}}{=} \sum_{(\bar{F}, \bar{x}, \sigma) \in \bar{\mathcal{F}}_E} \frac{1}{r!} \mu(B(\bar{F}, \bar{x})), \quad (34)$$

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

For $k \geq 1$ we let

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

Let k be large, and isolate contributions to the event E where $H^* \not\subset \Lambda_k$, or where separation of the $X_{\Lambda, i}$ does not occur. By the discussions preceding (16), we have

$$\begin{aligned} &\left| \nu_n(E) - \sum_{(\bar{F}, \bar{x}, \sigma) \in \bar{\mathcal{F}}_E(k)} \mu_n(B_n(\bar{F}, \bar{x}), \sigma_n^* = \sigma, Y_n > K(\bar{F})) \right| \\ &\leq \mu_n(H^* \not\subset \Lambda_k) + \sum_{(\bar{F}, \bar{x}, \sigma) \in \bar{\mathcal{F}}_E(k)} \mu_n(B_n(\bar{F}, \bar{x}), Y_n \leq K(\bar{F})). \end{aligned} \quad (35)$$

Given $\varepsilon > 0$, by (17) we can choose k large, so that the limsup of the first term on the right hand side, as $n \rightarrow \infty$, is at most ε . Fixing such a k , and noting that $\bar{\mathcal{F}}_E(k)$ is finite, the second term on the

right hand side of (35) is less than ε , if n is large enough, by (18). Also, for each $(\bar{F}, \bar{x}) \in \bar{\mathcal{F}}_E(k)$, the summand on the left hand side of (35) approaches $\mu(B(\bar{F}, \bar{x}))/r!$ by (19). Now letting $\varepsilon \rightarrow 0$ proves that $\lim_{n \rightarrow \infty} \nu_n(E) = \nu(E)$.

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 $t > 0$

$$P(T_n(z)/(2dn^2) < t) \rightarrow P(\tau(z) < t) \quad \text{uniformly in } z. \quad (36)$$

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 (36) can be recast as

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

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 $\alpha \rightarrow 0$. \square

References

- [1] Bak, P., Tang, C. and Wiesenfeld, K.: Self-organized criticality: An explanation of the 1/f noise. Phys. Rev. A **59**, 381-384 (1987)
- [2] Bak, P., Tang, C. and Wiesenfeld, K.: Self-organized criticality. Phys. Rev. A **38**, 364-374 (1988)
- [3] Benjamini, I., Lyons, R., Peres, Y. and Schramm, O.: Uniform spanning forests. Ann. Probab. **29**, 1-65 (2001)
- [4] Dhar, D. and Ramaswamy, R.: Exactly solved model of self-organized critical phenomena. Phys. Rev. Lett. **63**, 1659-1662 (1989)

- [5] Dhar, D.: Self-organized critical state of sandpile automaton models. *Phys. Rev. Lett.* **64**, 1613–1616 (1990)
- [6] Dhar, D.: The Abelian sandpile and related models. *Phys. A* **263**, 4–25 (1999)
- [7] Dhar, D.: Studying Self-organized criticality with exactly solved models. Preprint (1999) <http://arXiv.org/abs/cond-mat/9909009>
- [8] Dhar, D. and Majumdar, S.N.: Abelian sandpile models on the Bethe lattice. *J. Phys. A* **23**, 4333–4350 (1990)
- [9] Ivashkevich, E.V. and Priezzhev, V.B.: Introduction to the sandpile model. *Phys. A* **254**, 97–116 (1998)
- [10] Járαι, A.A. and Redig, F.: Infinite volume limits of high-dimensional sandpile models. In preparation
- [11] Jensen, H.J.: *Self-organized criticality. Emergent complex behavior in physical and biological systems.* (Cambridge Lecture Notes in Physics, 10) Cambridge University Press (2000)
- [12] Lawler, G.F.: *Intersections of random walks.* Birkhäuser, softcover edition (1996)
- [13] Loève, M.: *Probability theory I–II.* (Graduate Texts in Mathematics, 45–46) Springer-Verlag, 4th edition (1977)
- [14] Maes, C., Redig, F. and Saada, E.: The Abelian sandpile model on an infinite tree. *Ann. Probab.* **30**, 2081–2107 (2002)
- [15] Maes, C., Redig, F. and Saada, E.: The infinite volume limit of dissipative abelian sandpiles. Preprint (2003) To appear in *Commun. Math. Phys.* http://www.ma.utexas.edu/mp_arc
- [16] 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**, 1–22 (2000)
- [17] Mahieu, S. and Ruelle, P.: Scaling fields in the two-dimensional Abelian sandpile model. *Phys. Rev. E* **64**, 066130 (2001)
- [18] Majumdar, S.N. and Dhar, D.: Height correlations in the Abelian sandpile model. *J. Phys. A* **24**, L357–L362 (1991)
- [19] Majumdar, S.N. and Dhar, D.: Equivalence between the Abelian sandpile model and the $q \rightarrow 0$ limit of the Potts model. *Physica A* **185**, 129–145 (1992)
- [20] Meester, R., Redig, F. and Znamenski, D.: The Abelian sandpile; a mathematical introduction. *Markov Process. Related Fields* **7**, 509–523 (2002)

- [21] Pemantle, R.: Choosing a spanning tree for the integer lattice uniformly. *Ann. Probab.* **19**, 1559–1574 (1991)
- [22] Priezzhev, V.B.: Structure of two-dimensional sandpile. I. Height Probabilities. *J. Stat. Phys.* **74**, 955–979 (1994)
- [23] Priezzhev, V.B.: The upper critical dimension of the Abelian sandpile model. *J. Stat. Phys.* **98**, 667–684 (2000)
- [24] Tebaldi, C., De Menech, M. and Stella, A.L.: Multifractal scaling in the Bak-Tang-Wiesenfeld sandpile and edge events. *Phys. Rev. Letters* **83**, 3952–3955 (1999)
- [25] Vespignani, A. and Zapperi S.: How Self-organised criticality works: A unified mean-field picture. *Phys. Rev. A* **57**, 6345–6361 (1988).
- [26] 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*, ACM, New York, 296–303 (1996)

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