Thermodynamic Limit of the Abelian Sandpile Model on $\mathbf{Z}^d$

A.A. Járai

Carleton University, School of Mathematics and Statistics, 1125 Colonel By Drive, Ottawa, ON K1S 5B6, Canada. E-mail: jarai@math.carleton.ca

Abstract. We review basic properties of the Abelian sandpile model[1] and describe recent progress made regarding its infinite volume limit on $\mathbf{Z}^d$. In particular, we discuss the existence of the infinite volume limit of the stationary measure for $d \geq 2$, existence of infinite volume addition operators for $d \geq 3$, and construction of an infinite volume process for $d \geq 5$. We give an overview of the techniques relevant for these constructions.

1. Introduction

The Abelian sandpile, also known as the BTW model, was introduced by Bak, Tang and Wiesenfeld [2, 3], as an example of a concept they termed *self-organized criticality* (SOC). Although there is no mathematical definition of SOC, we can describe to what sort of models this term is usually applied. The general feature of SOC models is that they possess a stochastic dynamics that drives them towards a stationary state characterized by power law correlations in space and time. The behaviour should be compared to that of lattice systems of statistical physics at phase transition, such as critical percolation, where these models exhibit power law spatial correlations.

In SOC models, criticality appears with a new flavour: correlations build up as a result of the dynamics. Various physical situations where the concept of SOC may apply are described in the book [14].

---

[1]This note is a somewhat extended version of the talk given at the workshop.

The Abelian sandpile model (ASM) in particular has attracted a lot of attention, due to its simple definition and rich structure. The dynamics is given by a Markovian evolution of discrete *height variables* indexed by a finite set Λ. The mathematical study of the ASM was initiated by Dhar [5], who discovered its Abelian property, and coined the name Abelian. He gave an algebraic and combinatorial characterization of the configurations that occur in the stationary state with positive probability. Since then several analytic results have been obtained, for example, derivation of critical exponents on the Bethe lattice [8] and evidence of power law correlations [20]. Work of Majumdar and Dhar [21] revealed a correspondence between the sandpile model and the uniform spanning tree. Using this connection, exact height probabilities were obtained in two dimensions [24], and it was argued that the critical dimension of the model is four [25]. See [6,7,10] for reviews.

From the mathematical point of view, it is natural to consider the limit of infinite Λ, and try to define the model in infinite volume. This question was studied in various settings by Maes et al. [16,17,19]. The content of these works is reviewed in [18], which our paper is meant to complement.

The first part of our paper gives an introduction to the ASM, while the second part reports on recent progress regarding its infinite volume (thermodynamic) limit on $\mathbf{Z}^d$. We give a fairly detailed introduction to the basic properties of the model in Section 2. We focus on properties that are relevant for the infinite volume limit; no attempt was made to give a complete review. See [6,10,22] for additional introductory expositions. In Section 3, we give an overview of important techniques, as well as a brief description of infinite volume results of Maes et al. In Section 4, we state the main results for the infinite volume limit on $\mathbf{Z}^d$. Section 5 presents the main ideas of the proofs. A key role in the proofs is played by the correspondence with spanning trees; reviewed in Section 3.3.

2. The model and its basic properties

2.1. The model

Let $\Lambda \subset \mathbf{Z}^d$ be finite. Each site $i \in \Lambda$ is assigned a positive integer *height variable* z_i. Values between 1 and $2d$ are called *stable*, values larger than $2d$ are *unstable*. We let $\Omega_\Lambda = \{1, \ldots, 2d\}^\Lambda$ denote the space of stable configurations. We are also given a probability distribution $\{q_\Lambda(i)\}_{i \in \Lambda}$ on the set Λ. For much of our discussion, the reader may assume that q_Λ is the uniform distribution.

The configuration undergoes the following discrete time dynamics. If the current state is $z_\Lambda \in \Omega_\Lambda$, we choose a site i at random according to q_Λ, and increase the height at i by 1, that is

$$z_i \to z_i + 1.$$

If the new height is stable ($\leq 2d$), we have a new configuration in Ω_Λ, and that is the new state of the system. If i became unstable ($> 2d$), the site i *topples*, which means the operation of moving one particle from i to each of its neighbours. This can be conveniently written as

$$z_j \to z_j - \Delta_{ij} \quad \text{for all } j \in \Lambda, \tag{2.1}$$

where Δ is the discrete Laplacian in Λ:

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

As a result of toppling i, it may happen that one or more neighbours of i become unstable. We topple them as well, and continue as long as there are unstable sites. Note that (2.1) implies that when we topple a site on the boundary of Λ, one or more particles leave the system. This ensures that a stable configuration is reached in finitely many steps.

It is a basic fact that irrespective of the order in which topplings are carried out, the same sites will topple the same number of times, and therefore a unique stable configuration is reached in the relaxation process [5, 22]. This new stable configuration is the state of the system after a single time step (we think of the sequence of topplings to be carried out instantaneously). See the example in Figure 1. The sequence of topplings performed in a single time step is called an *avalanche*. The model has a natural generalization defined on an arbitrary finite directed graph via a general *toppling matrix* Δ [5].

2.2. Abelian property

We define the operators a_i: $\Omega_\Lambda \to \Omega_\Lambda$ as the result of particle addition at i and subsequent relaxation through an avalanche. The following Abelian property is crucial for the model [5]:

$$a_i a_j = a_j a_i \quad \text{for all } i, j \in \Lambda.$$

Abelianness is a consequence of the property that the rule whether to topple a site only depends on the height at that site, and not on its neighbours. A physically more realistic model of a sandpile is to topple a site if the discrete gradient along some edge is larger than a critical value, by moving particles in the direction that decreases the gradient. In such models, topplings of two unstable sites do not commute in general, which makes them less tractable than Abelian models. The phenomenon of SOC, however, is already present in the Abelian model.

In the dynamics defined above, the operator a_i is applied with probability $q_\Lambda(i)$ at each time step, defining a Markov chain with state space Ω_Λ. It can be

3	3	2	1		3	3	2	1		3	3	2	1
3	4	4	3	addition →	3	4	4	3	addition →	3	5	4	3
4	4	2	2		4	4	3	2		4	4	3	2
3	3	3	2		3	3	3	2		3	3	3	2

3	4	2	1		3	4	3	1
4	1	5	3	two topplings →	4	3	1	4
4	5	3	2		5	1	5	2
3	3	3	2		3	4	3	2

(toppling ↝)

3	4	3	1		4	4	3	1
5	3	2	4	toppling ↝	1	4	2	4
1	3	1	3		2	3	1	3
4	4	4	2		4	4	4	2
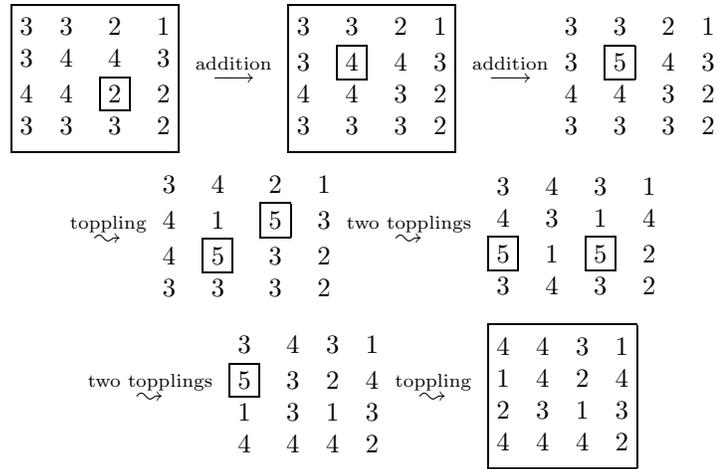
(two topplings ↝)

Figure 1. Illustration of the dynamics on a 4×4 lattice. Framed configurations are stable. Framed sites represent particle addition or toppling. The second particle addition results in an avalanche of 6 topplings.

shown using the Abelian property that as long as $q_\Lambda(i) > 0$ for all $i \in \Lambda$, there is a unique stationary distribution ν_Λ. The stationary distribution is independent of q_Λ, and uniform on the set of recurrent states of the Markov chain [5]. It is common to take q_Λ to be uniform.

In studying the infinite volume dynamics, we run into the problem that there is no uniform probability measure on $\mathbf{Z}^d$. A potential solution is to pass to a continuous time model, where additions occur according to independent Poisson processes with rates $\{\varphi(i)\}_{i \in \Lambda}$. One needs an assumption on the behaviour of $\varphi(i)$ at infinity for the infinite volume dynamics to be well-defined. In particular, for the standard model on $\mathbf{Z}^d$, φ has to go to zero sufficiently fast; see Theorem 4.2 in Section 4. A uniform addition rate $\varphi(i) \equiv$ constant is possible for so-called dissipative models; see Section 3.7.

2.3. Recurrent states

Let $\mathcal{R}_\Lambda$ denote the set of recurrent states of the Markov chain. Then the operators a_i restricted to $\mathcal{R}_\Lambda$ generate an Abelian group K_Λ, and $|K_\Lambda| = |\mathcal{R}_\Lambda| = \det(\Delta)$ [5]. In fact, K_Λ is isomorphic to the factor group $\mathbf{Z}^\Lambda / S_\Lambda$, where S_Λ is the additive subgroup of $\mathbf{Z}^\Lambda$ generated by the rows of Δ. This property reflects the rule (2.1). In fact, $\mathcal{R}_\Lambda$ contains one representative from each co-set of S_Λ, and therefore can be identified with K_Λ in a natural way.

In addition to its algebraic aspect, $\mathcal{R}_\Lambda$ has a combinatorial description. Namely, one can characterize the transient states $\Omega_\Lambda \setminus \mathcal{R}_\Lambda$ via *forbidden sub-configurations* (FSC). As an example, first consider a configuration in which there are two ones next to each other:

$$\boxed{1}\boxed{1} \tag{2.2}$$

A simple argument shows that such a state cannot be recurrent. Indeed, after sufficiently long time, at least one toppling has occurred at both sites. After this time, either site can only be 1 between a toppling and the next particle addition there. Therefore, the other site cannot be 1 at the same time. We call the configuration in (2.2) an FSC.

The above example can be generalized. Let $|A|$ denote the number of elements of a set A. A configuration on a finite set F is called an FSC, if

$$z_i \leq |\{j \in F : j \sim i\}| \quad \text{for all} \ \ i \in F, \tag{2.3}$$

where $j \sim i$ denotes that j and i are neighbours. For example,

			1
2	2	2	3
1			1

is an FSC. A configuration in Λ is called *allowed*, if it contains no FSC. It was shown in [5] that recurrent states are allowed. The converse is also true, and can be proved in a number of ways. First, for sandpiles with a symmetric toppling matrix, [21] gives a correspondence between allowed configurations and spanning trees, whose number is $\det(\Delta)$, the same as $|\mathcal{R}_\Lambda|$. See Section 3.3 below about this correspondence. A more direct and more general proof was later given by Speer [26], who generalized the notion of allowed configurations to sandpiles with an arbitrary (not necessarily symmetric) toppling matrix, and showed its equivalence to recurrence. Finally, [22] gives an alternative proof of the equivalence between allowed and recurrent configurations in the symmetric case, without invoking the correspondence with spanning trees.

2.4. Expected number of topplings

We close the list of basic properties with a very useful observation of Dhar [5]. Let

$$N_\Lambda(i, j) = \text{number of topplings at } j \text{ caused by particle addition at } i.$$

By the definition of the toppling procedure:

$$(a_i z)_j = z_j + \delta_{ij} - \sum_{k \in \Lambda} N_\Lambda(i, k) \Delta_{kj}. \tag{2.4}$$

Let E_μ denote expectation with respect to a probability measure μ. Averaging (2.4) with respect to the invariant measure ν_Λ, we get

$$\sum_{k \in \Lambda} \mathsf{E}_{\nu_\Lambda} N_\Lambda(i,k)\Delta_{kj} = \delta_{ij}.$$

Therefore

$$\mathsf{E}_{\nu_\Lambda} N_\Lambda(i,j) = (\Delta^{-1})_{ij} \overset{\text{def}}{=} G_\Lambda(i,j) \qquad \text{(Dhar's formula).} \qquad (2.5)$$

Here G_Λ is $(2d)^{-1}$ times the Green function of simple random walk in Λ, killed on exit from Λ.

2.5. Notation

We fix some more notation for the rest of the paper. We let $\Omega = \{1,\ldots,2d\}^{\mathbf{Z}^d}$ denote the set of stable configurations in infinite volume. We endow Ω with the product topology, which makes it into a compact metric space. Given a function $f(\Lambda)$ defined on all sufficiently large finite subsets of $\mathbf{Z}^d$, and taking values in a metric space with metric ρ, we say that $\lim_{\Lambda \to \mathbf{Z}^d} f(\Lambda) = a$, if for any $\varepsilon > 0$ there exists Λ_0 such that for all $\Lambda \supset \Lambda_0$ we have $\rho(f(\Lambda),a) < \varepsilon$.

For $K \subset \mathbf{Z}^d$, we let $\mathcal{F}_K$ denote the σ-field on Ω generated by the height variables in K. A *local event* (resp. *local function*) is an event (resp. function) that belongs to $\mathcal{F}_K$ for some finite K.

For $i \in \Lambda$, we let $\deg_\Lambda(i) = |\{j \in \Lambda : j \sim i\}|$.

2.6. Basic questions

A basic problem is to describe properties of avalanches under stationarity. One can consider the following avalanche characteristics:

1) number of topplings in the avalanche (size);

2) number of sites toppled during the avalanche (range);

3) distance of the furthest toppled site from the start of the avalanche (radius);

4) the maximum number of times a site topples (duration).

It is generally believed that the distributions of the above quantities obey power laws when $d \geq 2$, in the thermodynamic limit $\Lambda \to \mathbf{Z}^d$. More precisely, this means the following. Let S_Λ denote one of the above four characteristics for an avalanche in the stationary ASM. That is, sample a configuration from ν_Λ, drop a particle at a fixed site, let's say the origin, and let S_Λ denote the value of the chosen quantity for this avalanche. It is expected that

$$\lim_{\Lambda \to \mathbf{Z}^d} \text{Prob}\{S_\Lambda = s\} = p(s),$$

where $p(s) \sim cs^{-\tau}$, as $s \to \infty$, for some constant $c > 0$ and *critical exponent* $\tau > 0$ depending on d and the chosen quantity.

In trying to build a mathematical framework in which the above power laws can be studied, one is led to the questions below; a program started in [16,17,19].

1) Does ν_Λ have a (weak) limit ν as $\Lambda \to \mathbf{Z}^d$?

2) Can one define the operators a_i on ν-typical configurations?

3) Are avalanches ν-a.s. finite?

4) Does a_i leave ν invariant?

5) Is there a natural infinite volume Markov process with invariant measure ν? Does it have good ergodic properties?

In the next section we give an overview of some results and techniques developed for the ASM that are relevant for studying these questions.

3. An overview of results and techniques

3.1. The burning algorithm

There is a simple algorithm that checks whether a given configuration is allowed (and hence recurrent) [5]. We start the algorithm by declaring all sites in Λ to be unburnt, and then successively burn sites whose height is larger than the number of their unburnt neighbours. More formally, let $A_0 = \emptyset$, the set of sites burning at time 0. If A_s has been defined for $0 \le s \le t$, we let $F_t = \Lambda \setminus (\cup_{s=0}^t A_s)$, the set of unburnt sites after time t. We set

$$A_{t+1} = \{i \in F_t : z_i > \deg_{F_t}(i)\}.$$

A_t is the set of sites burning at time t; see Figure 2.

It is not hard to show that the burning process recursively removes sites that cannot be part of an FSC. If the process stops without burning all sites, the remaining configuration is an FSC. If all sites are burnt, then the initial configuration was allowed.

Note that the burning always starts at the boundary of Λ. In fact, we can think of the fire starting at an artificially added site δ, usually called the *sink*. We connect δ to each $i \in \partial\Lambda$ by $2d - \deg_\Lambda(i)$ edges. We denote this new graph by $\widetilde{\Lambda}$. We declare δ to be burning at time 0, set $A_0 = \{\delta\}$, and then proceed with the same burning rule as before.

There is an equivalent way of thinking about the burning process. The addition operators satisfy the relations

$$\prod_{j \in \Lambda} a_j^{\Delta_{ij}} = \mathrm{id}, \qquad i \in \Lambda,$$

$$
\begin{array}{cccc}
\boxed{4} & 3 & 2 & 2 \\
\boxed{4} & 2 & 1 & 3 \\
1 & 4 & 2 & 3 \\
\boxed{3} & 3 & 2 & \boxed{3}
\end{array}
\longrightarrow
\begin{array}{cccc}
\bullet & \boxed{3} & 2 & 2 \\
\bullet & 2 & 1 & 3 \\
1 & 4 & 2 & \boxed{3} \\
\bullet & \boxed{3} & 2 & \bullet
\end{array}
\longrightarrow
\begin{array}{cccc}
\bullet & \bullet & 2 & 2 \\
\bullet & 2 & 1 & \boxed{3} \\
1 & \boxed{4} & 2 & \bullet \\
\bullet & \bullet & \boxed{2} & \bullet
\end{array}
\longrightarrow
$$

$$
\begin{array}{cccc}
\bullet & \bullet & 2 & \boxed{2} \\
\bullet & \boxed{2} & 1 & \bullet \\
\boxed{1} & \bullet & \boxed{2} & \bullet \\
\bullet & \bullet & \bullet & \bullet
\end{array}
\longrightarrow
\begin{array}{cccc}
\bullet & \bullet & \boxed{2} & \bullet \\
\bullet & \bullet & 1 & \bullet \\
\bullet & \bullet & \bullet & \bullet \\
\bullet & \bullet & \bullet & \bullet
\end{array}
\longrightarrow
\begin{array}{cccc}
\bullet & \bullet & \bullet & \bullet \\
\bullet & \bullet & \boxed{1} & \bullet \\
\bullet & \bullet & \bullet & \bullet \\
\bullet & \bullet & \bullet & \bullet
\end{array}
$$
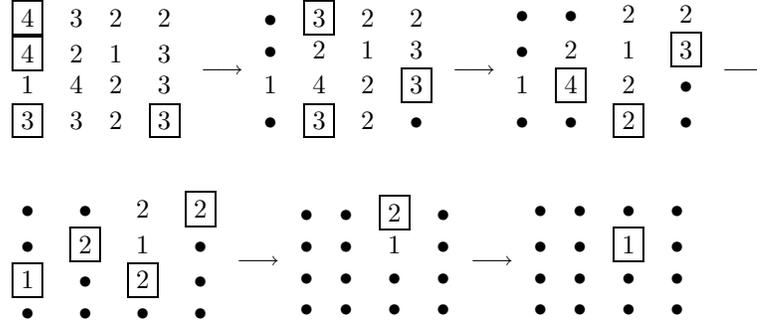
Figure 2. Illustration of the burning algorithm on a 4×4 lattice. At each step, framed sites are burnt.

where id is the identity element of K_Λ. Multiplying these relations over $i \in \Lambda$, factors containing any $i \in \Lambda \setminus \partial\Lambda$ cancel. We get the relation

$$
\prod_{i \in \partial\Lambda} a_i^{v_i} = \mathrm{id},
$$

where $v_i = \sum_{j \in \Lambda} \Delta_{ij} = 2d - \deg_\Lambda(i)$. The latter expression is the number of neighbours of i in Λ^c. This means that adding v_i particles to each boundary site in a recurrent configuration triggers an avalanche that recreates the original configuration. Alternatively, we can think of the avalanche being triggered by sending one particle along each edge emanating from δ. It is not hard to see that in this avalanche, each site topples exactly once. Following the sequence of topplings is equivalent to the burning process.

3.2. Exact probabilities via determinants

We already mentioned that the number of recurrent states is $|\mathcal{R}_\Lambda| = \det(\Delta_\Lambda)$. Using this determinantal formula in a subtle way, Majumdar and Dhar [20] gave a method for calculating the probabilities of certain configurations in the thermodynamic limit $\Lambda \to \mathbf{Z}^d$. To illustrate the method, consider the probability $\nu_\Lambda(z_0 = 1)$ of having height 1 at a given site. Let $x_1, \ldots, x_{2d}$ denote the $2d$ neighbours of the origin. We modify Λ by removing the $2d - 1$ edges connecting 0 to $x_2, \ldots, x_{2d}$. We denote this modified graph by Λ'. In Λ', 0 is connected by the single edge $0x_1$ to the rest of the lattice. We modify the toppling matrix accordingly, by setting

$$
\begin{aligned}
\Delta'_{00} &= \Delta_{00} - (2d - 1) = 1; \\
\Delta'_{x_j x_j} &= \Delta_{x_j x_j} - 1 = 2d - 1, \quad (2 \le j \le 2d);
\end{aligned}
$$

and

$$\Delta'_{0x_j} = \Delta'_{x_j 0} = 0, \quad (2 \leq j \leq 2d).$$

We can write $\Delta' = \Delta + B$, where the matrix B satisfies $B_{yz} = 0$, unless $y, z \in \{0, x_2, \ldots, x_{2d}\}$.

One can prove that there is a one-to-one correspondence between recurrent configurations in Λ satisfying $z_0 = 1$, and recurrent configurations in Λ' satisfying $z'_0 = 1$. (The correspondence is given by decreasing the height at x_j by 1 for $2 \leq j \leq 2d$.) This implies that

$$\nu_\Lambda(z_0 = 1) = \nu_{\Lambda'}(z'_0 = 1)\frac{\det(\Delta')}{\det(\Delta)} = \nu_{\Lambda'}(z_0 = 1)\det(I + GB),$$

where $G = \Delta^{-1}$. The key idea is that due to the properties of B, the determinant on the right hand side reduces to a finite determinant, whose dimension does not depend on Λ. The value of the determinant only depends on G_{ij}, $i, j \in \{0, x_1, \ldots, x_{2d}\}$. Also, one easily checks $\nu_{\Lambda'}(z'_0 = 1) = 1$, hence

$$\lim_{\Lambda \to \mathbf{Z}^d} \nu_\Lambda(z_0 = 1) = \lim_{\Lambda \to \mathbf{Z}^d} \det(I + GB).$$

In two dimensions, explicit values of the potential kernel [27, page 148] yield the exact result $\nu(z_0 = 1) = (2/\pi^2) - (2/\pi^3)$ [20].

The method can be generalized to *minimally allowed* configurations, that is, finite configurations that do not contain an FSC, but decreasing any of the heights creates an FSC. For example,

$$\begin{matrix} & 1 & 2 & \\ 1 & 4 & 4 & 2 \\ & 1 & 2 & \end{matrix}$$

is a minimally allowed configuration in two dimensions. In particular, the method can be used to calculate the correlation between the events that sites 0 and x have height 1, respectively, giving the power law decay

$$\lim_{\Lambda \to \mathbf{Z}^d} \left[\nu_\Lambda(z_0 = 1, z_x = 1) - \nu_\Lambda(z_0 = 1)\nu_\Lambda(z_x = 1)\right] \simeq \frac{1}{|x|^{2d}}.$$

The method does not apply to probabilities of heights greater than 1. In two dimensions, the probabilities of heights 2,3 and 4 were evaluated by Priezzhev [24] who gave exact analytical expressions for them. His method is based on a correspondence with spanning trees. This correspondence also plays a crucial role in our results, and is reviewed in the next section.
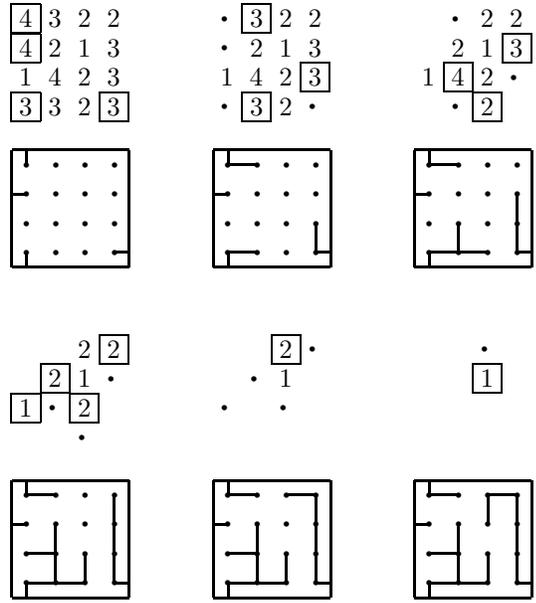
Figure 3. Spanning tree representation of the burning process on a 4×4 lattice. Framed sites are burnt at each step. Dots represent sites that were burnt in the previous step. The outer box around each tree represents the vertex δ.

3.3. Spanning trees

In this section we describe a deep observation of Majumdar and Dhar [21] that relates the ASM to the uniform spanning tree (UST).

We follow the spread of fire in the burning process, and build a tree by connecting each site to a neighbour that was burnt in the previous step. Here it will be convenient to use the graph $\widetilde{\Lambda}$ introduced in Section 3.1. Recall that A_t denotes the set of sites burning at time t. Observe that each $i \in A_t$ ($t \geq 1$) has at least one neighbour in A_{t-1}, because sites become burnable after a certain number of their neighbours have been burnt. Below we define a rule that assigns to any $i \in A_t$ a unique element of A_{t-1}, called the *parent* of i. Joining each site to its parent then defines a spanning tree on $\widetilde{\Lambda}$.

The construction of the spanning tree is illustrated in Figure 3. In the first step, four boundary sites were burnt (that is $|A_1| = 4$), and these are connected to δ (the only element of A_0) represented by a square drawn around Λ. Note that for the three corner sites there are two possible edges to choose from, and we have not yet specified how a choice is made. At each subsequent step, newly burnt sites are connected to sites burnt in the previous step, marked by dots.

Now we specify how to choose the parent of $i \in A_t$. Let

$$n_s(i) = |\{\text{unburnt neighbours of } i \text{ at time } s\}|.$$

The number of neighbours of $i \in A_t$ that are in A_{t-1} is $n_{t-1}(i) - n_t(i)$. By the burning rule, the possible height values at i consistent with $i \in A_t$ are $n_t < z_i \leq n_{t-1}$. Since the number of available neighbours is the same as the number of possible values of z_i, we can assign a unique parent $j \in A_{t-1}$ to i depending on the value of z_i. To be specific, let us order the $2d$ lattice directions in a fixed way. Associating larger z_i with larger lattice direction defines the parent uniquely.

In Figure 3, the directions were ordered as N > W > S > E. During the first step, the site in the upper left corner having height 4 is burnt. The available edges are in the directions N and W, and possible heights consistent with this site being burnt at this stage are 3 and 4. Therefore the larger direction, N was chosen. At the two other corners having height 3, the smaller direction was chosen. The only other application of the rule occurs in step four, when a site with height 2 has two neighbours burnt in the previous step.

The above procedure results in a tree $T_\Lambda = \phi(z_\Lambda)$, which spans $\widetilde{\Lambda}$ if and only if the sets $(A_t)_{t \geq 1}$ exhaust Λ. By the burning test of Section 3.1, this happens if and only if the configuration was recurrent. We regard δ as the root of the tree. A_t is precisely the set of sites at graph distance t from the root.

The procedure can be reversed to show that ϕ is a one-to-one mapping between recurrent configurations and spanning trees of $\widetilde{\Lambda}$. 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 δ for $t \geq 0$. 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 $m_{t-1}(j) - m_t(j)$. 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 $m_t(j) < z_j \leq m_{t-1}(j)$ are satisfied, and we pick that value which corresponds to the parent of j according to our fixed ordering of directions. It is clear that the resulting configuration z_Λ is such that in the burning test $A_t = B_t$, $n_t(j) = m_t(j)$ and $\phi(z_\Lambda) = T_\Lambda$.

Since ν_Λ gives equal weight to all recurrent states, the image of ν_Λ under ϕ is uniform on the spanning trees of $\widetilde{\Lambda}$. This distribution is called the uniform spanning tree (UST) on Λ with wired boundary conditions [4,23]. "Wired" refers to the fact that the connection between two sites in T_Λ may occur through the artificially added vertex δ (in contrast with "free" boundary conditions, when connections are required to occur within Λ). We denote the law of the wired UST by μ_Λ. The next section describes the limit of μ_Λ as $\Lambda \to \mathbf{Z}^d$.

3.4. The uniform spanning forest

It was shown by Pemantle [23] that as $\Lambda \to \mathbf{Z}^d$, the uniform spanning tree converges weakly to a limit called the uniform spanning forest (USF). The the-

orem below summarizes this result together with an extension (statement 5) proved in [4]. For more background on uniform spanning forests see [4].

To state the theorem, it will be convenient to think of μ_Λ as a measure on $\Omega' = \{0,1\}^{\mathrm{E}^d}$, where $\mathbf{E}^d$ denotes the set of edges in $\mathbf{Z}^d$. We equip Ω' with the Borel σ-field with respect to the product topology. Elements of Ω' can be viewed as subgraphs of $(\mathbf{Z}^d, \mathbf{E}^d)$, and any $\omega \in \Omega'$ defines a subgraph $T_\Lambda = T_\Lambda(\omega)$ of $\widetilde{\Lambda}$ by identifying all vertices in $\mathbf{Z}^d \setminus \Lambda$ to a single vertex and removing loops. Then μ_Λ is the unique measure under which T_Λ is a.s. a tree and uniformly distributed. We denote by $T = T(\omega)$ the set of edges present in $\omega \in \Omega'$. We say that an infinite tree has *one end*, if any two infinite paths have infinitely many vertices in common.

Theorem 3.1. *Let $d \geq 1$. For any finite sets $B \subset K$ of edges in $\mathbf{Z}^d$ the limit*

$$\mu(T \cap K = B) \stackrel{\mathrm{def}}{=} \lim_{\Lambda \to \mathrm{Z}^d} \mu_\Lambda(T_\Lambda \cap K = B) \tag{3.1}$$

exists, and uniquely defines a translation invariant measure μ on Ω'. We have $\mu = \lim_{\Lambda \to \mathrm{Z}^d} \mu_\Lambda$ in the sense of weak convergence, and μ has the following properties.

1) *T has no cycles μ-a.s.*

2) *If $d \leq 4$, T has a single component μ-a.s., that is, T is a tree μ-a.s.*

3) *If $2 \leq d \leq 4$, T has one end μ-a.s.*

4) *If $d > 4$, T has infinitely many components μ-a.s., that is, T is a forest. Each tree in T is infinite μ-a.s.*

5) *If $d > 4$, each component of T has a single end μ-a.s.*

We call the random graph T governed by μ the USF, omitting reference to the wired boundary condition. (In fact, on $\mathbf{Z}^d$ the wired and free spanning forests coincide [4].)

Given the existence of $\lim_{\Lambda \to \mathrm{Z}^d} \mu_\Lambda$, and the coding of sandpile configurations by trees, it will not come as surprise that $\lim_{\Lambda \to \mathrm{Z}^d} \nu_\Lambda$ also exists, and the answer to Question 1 in Section 2.6 is affirmative; see Section 4 below. Theorem 3.1 implies that in the case $2 \leq d \leq 4$, we can think of T as a "rooted tree" with "root at infinity". More precisely, statements 2 and 3 imply that for any $j \in \mathbf{Z}^d$ there is a unique infinite path $j = v_0, v_1, \ldots$ in T μ-a.s., and hence we can define the parent of j to be v_1. As we discuss in Section 5, this property of T allows us to extend the coding of sandpile configurations by trees to infinite volume. The situation is more subtle when $d > 4$, because the USF has multiple components. As we describe in Section 5, a coding is still possible in this case, if we add to the USF a random ordering of its components. However, before discussing these constructions in more detail, we place them into context by an overview of related infinite volume results.

3.5. One-dimensional sandpile

It is not difficult to determine the set of recurrent states on a subinterval of $\mathbf{Z}$. These are the configurations containing at most one 1, for example

$$2\ 2\ 2\ 1\ 2\ 2\ 2\ 2\ 2\ 2.$$

The thermodynamic limit $\lim_{\Lambda\to\mathbf{Z}} \nu_\Lambda$ is trivial, concentrating on the single configuration that is identically 2. Maes, Redig, Saada and van Moffaert [19] construct an infinite volume dynamics, and prove convergence to the trivial stationary state from an arbitrary initial configuration.

3.6. Bethe lattice and critical dimension

An analysis of the sandpile model on the Bethe lattice has been carried out by Dhar and Majumdar [8]. The Bethe lattice is a tree S where each vertex has degree $c \geq 3$. For our discussion, we assume $c = 3$, so heights satisfy $1 \leq z_i \leq 3$. Height probabilities and correlations of heights at different sites can be evaluated in the thermodynamic limit [8]. Also, it turns out that an avalanche started at some vertex i is entirely determined by the 3-cluster containing i (this is the largest connected set containing i where all heights are 3). The probability distribution of the 3-cluster containing a given site can be explicitly computed, and using this, critical exponents can be evaluated [8]. In particular,

$$\nu(\text{avalanche has size } s) \sim cs^{-3/2}, \quad \text{as } s \to \infty. \tag{3.2}$$

Maes, Redig and Saada [16] constructed infinite volume dynamics on the Bethe lattice that is an extension of the finite volume dynamics in a well-defined sense (in the sense of Theorem 4.2 below). Here it is convenient to define the finite volume dynamics in continuous time, by adding particles at the event times of independent Poisson processes with rates $\{\varphi(i)\}_{i\in\Lambda}$. This corresponds to the discrete time process with

$$q_\Lambda(i) = \varphi(i)\Big[\sum_{j\in\Lambda}\varphi(j)\Big]^{-1}.$$

When considering particle additions in infinite volume, one has to ensure that there are only finitely many topplings at any site in a finite time interval $[0, T]$. Due to the formula (2.5), one is led to impose the condition:

$$\sum_{i\in S}\varphi(i)G(i,0) < \infty, \tag{3.3}$$

where G is the inverse of Δ on the full lattice, and 0 is a fixed site, called the origin. The sum in (3.3) represents the expected number of topplings at 0, due

to particle additions in the time interval $[0, 1]$. Under condition (3.3), a Markov process with invariant measure ν can be constructed [16]. Here $\varphi(i)$ has to go to 0 sufficiently fast, as $|i| \to \infty$, to make the sum convergent, and hence the addition rate cannot be uniform.

It is expected that the behaviour on $\mathbf{Z}^d$ is similar to the behaviour on the Bethe lattice, for $d > d_u$, the upper critical dimension. In particular, critical exponents on $\mathbf{Z}^d$ for $d > d_u$ would take on the same values as on the Bethe lattice. Priezzhev [25] argued that the upper critical dimension of the sandpile model is $d_u = 4$. His arguments strongly support the conjecture that the asymptotics in (3.2) holds on $\mathbf{Z}^d$ when $d > 4$, and also in $d = 4$ with a logarithmic correction.

3.7. Dissipative models

Construction of the infinite volume limit has also been carried out for *dissipative models* on $\mathbf{Z}^d$ [17]. In our models so far, particles could only leave the system at boundary sites of Λ. If we modify Δ by letting $\Delta_{ii} = \gamma > 2d$, then particles will leave the system on toppling at any i. Stronger dissipation leads to faster correlation decay, which makes the infinite volume limit more tractable. All questions in Section 2.6 can be answered in the affirmative; see [17] for more details.

4. The thermodynamic limit on $\mathbf{Z}^d$

In this section we consider the questions listed in Section 2.6 in the case of $\mathbf{Z}^d$, $d \geq 2$. We are able to answer all questions in the case $d > 4$, and we address some of them when $2 \leq d \leq 4$. The main ideas of the proofs are summarized in Section 5.

Theorem 4.1. *Let $d \geq 2$. The limit $\nu = \lim_{\Lambda \to \mathbf{Z}^d} \nu_\Lambda$ exists in the sense of weak convergence. ν is translation invariant, and there are no FSC's ν-a.s.*

In [1], Theorem 4.1 is proved when $2 \leq d \leq 4$ and also when $d > 4$ with the restriction that Λ is sufficiently regular, for example, the limit exists along the sequence of cubes $\Lambda_n = [-n, n]^d$. The more general convergence result for $d > 4$ is proved in [13].

The proof is based on the correspondence with spanning trees given in Section 3.3 and the existence of the USF. When $2 \leq d \leq 4$, Theorem 4.1 follows rather easily from Pemantle's result, and a continuity argument. The coding of sandpile configurations by trees continues to hold in $\mathbf{Z}^d$ in the following sense. There is a measurable (in fact μ-a.s. continuous) transformation $\Phi : \Omega' \to \Omega$ that assigns a height configuration to any realization of the USF, such that $\nu = \mu \circ \Phi^{-1}$; see [13]. When $d > 4$, the correspondence between the sandpile model and the USF breaks down on $\mathbf{Z}^d$. However, one can still encode sandpile

configurations, if one adds to the USF a random ordering of its components. We explain these ideas in more detail in Section 5.1.

Remark 4.1.

(i) Let S_Λ denote one of the avalanche characteristics introduced in Section 2.6. Then it follows that the limit

$$p(s) = \lim_{\Lambda \to \mathbf{Z}^d} \nu_\Lambda(S_\Lambda = s)$$

exists for $s \geq 0$. We have

$$\sum_{s=0}^{\infty} p(s) = 1 \quad \text{if and only if} \quad \nu(\text{infinite avalanche}) = 0.$$

(ii) The limit $p(s)$ is non-trivial for all $d \geq 2$, by the method of Majumdar and Dhar described in Section 3.2. For example, in two dimensions the configuration

$$
\begin{array}{ccccccc}
 & 1 & 2 & \dots & 2 & & \\
1 & 4 & 4 & \dots & 4 & 2 \\
 & 1 & 2 & \dots & 2 & & \\
\end{array}
$$

$$\underbrace{\qquad\qquad\qquad}_{\text{length } s}$$

produces an avalanche of size s, and is minimally allowed.

Given the existence of a unique thermodynamic limit ν, the next natural step is to define addition operators. First, define $a_{i,\Lambda} \colon \Omega \to \Omega$ as the result of applying the addition operator in volume Λ and leaving the configuration unchanged in $\mathbf{Z}^d \setminus \Lambda$. Recall that $N_\Lambda(i,j)$, $i,j \in \Lambda$ denotes the number of topplings at j caused by addition at i under the dynamics in Λ. With the above definition of $a_{i,\Lambda}$, we can also view $N_\Lambda(i,j)$ as a random variable on Ω. With these conventions, the following is proved in [13].

Proposition 4.1. *For any $i,j \in \mathbf{Z}^d$, the limit $\lim_{\Lambda \to \mathbf{Z}^d} N_\Lambda(i,j) = N(i,j) \leq \infty$ exists on Ω, and*

$$\mathsf{E}_\nu \, N(i,j) \leq \lim_{\Lambda \to \mathbf{Z}^d} G_\Lambda(i,j) = G(i,j).$$

If $d > 2$, $N(i,j)$ is finite ν-a.s., and $\lim_{\Lambda \to \mathbf{Z}^d} a_{i,\Lambda} = a_i$ exists ν-a.s.

In the above proposition, the important ingredient is that when $d > 2$, the expected number of topplings (and therefore the number of topplings) at all sites is finite ν-a.s. This is sufficient to define a_i, but does not imply that avalanches are finite.

We can proceed further, if we know that avalanches are ν-a.s. finite[2]. We can prove this at least when $d > 4$; see Proposition 4.3 below.

---

[2]Russell Lyons informed us (private communication) that it is possible to prove item 1 of Proposition 4.2 based only on transience, that is when $d > 2$.

Proposition 4.2. *Assume that avalanches are ν-a.s. finite. Then the following statements hold.*

1) *ν is invariant under a_i, $i \in \mathbf{Z}^d$.*

2) *$\mathsf{E}_\nu \, N(i,j) = G(i,j)$ (extension of Dhar's formula).*

3) *$a_i a_j = a_j a_i$, ν-a.s. (Abelian property).*

4) *a_i^{-1} exists ν-a.s.*

5) *For any $i \in \mathbf{Z}^d$ and for ν-a.e. $z \in \Omega$, there exist finite sets $V_i(z), W_i(z)$, such that $a_i z = a_{i,\Lambda} z$ for $\Lambda \supset V_i(z)$, and $a_i^{-1} z = a_{i,\Lambda}^{-1} z$ for $\Lambda \supset W_i(z)$.*

The proof of the above proposition is fairly independent of the model, and is based on techniques developed in [16, 17]. The key step then is to prove the following proposition.

Proposition 4.3. *If $d > 4$, there are no infinite avalanches ν-a.s.*

Proposition 4.3 is proved in [13]. We describe the main ideas of its proof in Section 5.3. Given the results of Propositions 4.3 and 4.2, we can invoke the machinery developed in [16] to construct a Markov process with invariant measure ν. Using this machinery, one can prove the theorem below that appears in [13].

Theorem 4.2. *Assume $d > 4$. Let $\varphi : \mathbf{Z}^d \to (0,\infty)$ be an addition rate such that*

$$\sum_{i \in \mathbf{Z}^d} \varphi(i) G(0,i) < \infty.$$

Then the following statements hold.

1) *The closure of the operator on $L_2(\nu)$ defined on local functions by*

$$L_\varphi f = \sum_{i \in \mathbf{Z}^d} \varphi(i)(a_i - I)f$$

   is the generator of a stationary Markov process $\{\eta_t : t \geq 0\}$.

2) *Let $\{N_t^\varphi(i)\}_{i \in \mathbf{Z}^d}$ be independent Poisson processes with rates $\varphi(i)$, governed by the probability measure P. The limit*

$$\eta_t = \lim_{\Lambda \to \mathbf{Z}^d} \left[\prod_{i \in \Lambda} a_{i,\Lambda}^{N_t^\varphi(i)} \right] \eta$$

   exists $\mathsf{P} \times \nu$-a.s. Moreover, η_t is a cadlag version of the process with generator L_φ.

One would also like to investigate ergodic properties of both ν and the constructed Markov process. We recall the definition of tail triviality. The *tail σ-field* is defined to be

$$\mathcal{F}_\infty = \bigcap_\Lambda \mathcal{F}_{\mathbf{Z}^d \setminus \Lambda},$$

where the intersection is over all finite subsets of $\mathbf{Z}^d$. A measure μ is *tail trivial*, if $\mu(A) = 0$ or 1 for every $A \in \mathcal{F}_\infty$. Tail triviality implies mixing under spatial translations. The following theorem is proved in [13]; see Section 5.4 below for the main ideas of its proof.

Theorem 4.3. ν *is tail trivial for any $d \geq 2$.*

Finally, as a consequence of Theorem 4.3 we can prove that the sandpile process in infinite volume is mixing in time; see [13].

Theorem 4.4. *Let $d > 4$. The process $\{\eta_t : t \geq 0\}$ of Theorem 4.2 is mixing.*

5. Main ideas of the proofs

5.1. Existence of $\lim_{\Lambda \to \mathbf{Z}^d} \nu_\Lambda$

We start by outlining the ideas behind Theorem 4.1. Recall from Section 3.3 the procedure that recovers the sandpile configuration from T_Λ. The following observation plays an important role.

In order to reconstruct z_j from T_Λ, 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. We explain this in more detail. Recall that B_t is the set of sites at graph distance t from the root. Let $\mathcal{N}$ denote the set of neighbours of j. Define the indices t and $t(i)$, $i \in \mathcal{N}$, by the conditions $j \in B_t$ and $i \in B_{t(i)}$, $i \in \mathcal{N}$. Let i_0 denote the parent of j, so that $t(i_0) = t - 1$. By the construction of T_Λ, z_j is determined by i_0 and the following two sets:

$$
\begin{aligned}
\mathcal{A} &= \{i \in \mathcal{N} : t(i) \leq t - 1\} \\
\mathcal{B} &= \{i \in \mathcal{N} : t(i) = t - 1\} \ni i_0.
\end{aligned}
\tag{5.1}
$$

$\mathcal{A}$ consists of those neighbours of j that burned before j, and $\mathcal{B}$ of those that burned one step before j. In particular, i_0 and the set of integers $\{t(i) - t\}_{i \in \mathcal{N}}$ determines z_j, even without knowing the value of t.

The above observation allows us to reconstruct z_j knowing only a small portion of T_Λ. For example, let v_Λ denote the earliest common ancestor of the vertices in $\mathcal{N}$, that is the first vertex common to all paths leading from $\mathcal{N}$ to δ. Let F_Λ denote the subtree of T_Λ consisting of the paths leading from $\{j\} \cup \mathcal{N}$ to v_Λ. In other words, (F_Λ, v_Λ) is the smallest rooted subtree of T_Λ containing

$\{j\} \cup \mathcal{N}$. The pair (F_Λ, v_Λ) already determines z_j. This is because, letting dist denote graph distance,

$$t(i) - t = \text{dist}_{F_\Lambda}(i, v_\Lambda) - \text{dist}_{F_\Lambda}(j, v_\Lambda), \quad i \in \mathcal{N},$$

and i_0 is the parent of j in the rooted tree (F_Λ, v_Λ).

It turns out that when $2 \leq d \leq 4$, the distribution of v_Λ is tight, that is for any $\varepsilon > 0$ there exists a finite set $K = K(\varepsilon) \subset \mathbf{Z}^d$ such that $\mu_\Lambda(v_\Lambda \in K) \geq 1 - \varepsilon$ for all $\Lambda \supset \{j\} \cup \mathcal{N}$. Hence F_Λ remains "localized" near j. More precisely, statements 2 and 3 of Theorem 3.1 imply that for any $k \in \mathbf{Z}^d$ there exists a unique infinite path $k = v_0, v_1, v_2, \ldots$ in T μ-a.s. Defining v_1 to be the parent of k, we can regard T as a "rooted tree" with "root at infinity". Therefore we can define infinite volume analogues of v_Λ and F_Λ μ-a.s. Namely, the paths in T from $\{j\} \cup \mathcal{N}$ to infinity will have a first common vertex v μ-a.s., and we define F as the union of the paths in T from $\{j\} \cup \mathcal{N}$ to v. One can show that (F_Λ, v_Λ) converges weakly to (F, v). This leads to the fact that the limiting distribution of z_j as $\Lambda \to \mathbf{Z}^d$ is determined by (F, v). Generalizing this idea, one can show that the coding of sandpile configurations by trees extends to the infinite volume.

In the case $d > 4$, the USF has multiple components, and therefore v and F cannot be defined μ-a.s. In this case the probability that $v_\Lambda = \delta$ is bounded away from 0, and hence the distribution of v_Λ is not tight. We modify the definitions given for $2 \leq d \leq 4$ by subdividing T_Λ into "components". With slight abuse of language, we say that $x, y \in \Lambda$ are *connected*, if the path in T_Λ from x to y does not pass through δ (that is, if they are connected in the usual graph-theoretic sense in the graph obtained by removing δ). Let $A^{(1)}, \ldots, A^{(r)}$ be the (random) partition of $\{j\} \cup \mathcal{N}$ into mutually disconnected subsets. In other words, the $A^{(\alpha)}$ are the non-empty intersections of $\{j\} \cup \mathcal{N}$ with connected components of T_Λ; see Figure 4. We use a fixed deterministic rule to index elements of the partition. Analogously to the case $2 \leq d \leq 4$, we define $(F_\Lambda^{(\alpha)}, v_\Lambda^{(\alpha)})$, $\alpha = 1, \ldots, r$, as the smallest rooted subtree of T_Λ containing $A^{(\alpha)}$. Now z_j is a function of $(F_\Lambda^{(\alpha)}, v_\Lambda^{(\alpha)})_{\alpha=1}^r$ *and* the distances

$$X_\Lambda^{(\alpha)} = \text{dist}_{T_\Lambda}(v_\Lambda^{(\alpha)}, \delta), \quad \alpha = 1, \ldots, r.$$

This because

$$t(i) - t = \text{dist}_{F_\Lambda^{(\alpha)}}(i, v_\Lambda^{(\alpha)}) + X_\Lambda^{(\alpha)} - \text{dist}_{F_\Lambda^{(\beta)}}(j, v_\Lambda^{(\beta)}) - X_\Lambda^{(\beta)} \qquad (5.2)$$

if $i \in A^{(\alpha)}$ and $j \in A^{(\beta)}$.

Fortunately, we only need rough information about $X_\Lambda^{(\alpha)}$, $\alpha = 1, \ldots, r$. To see this, observe that $\mathcal{A}$ and $\mathcal{B}$ in (5.1) are already determined by the *linear ordering* of the integers $\{t(i) - t\}_{i \in \mathcal{N}}$ (if we know i_0), since

$$\mathcal{A} = \{i \in \mathcal{N} : t(i) - t \leq t(i_0) - t\}$$
$$\mathcal{B} = \{i \in \mathcal{N} : t(i) - t = t(i_0) - t\}.$$
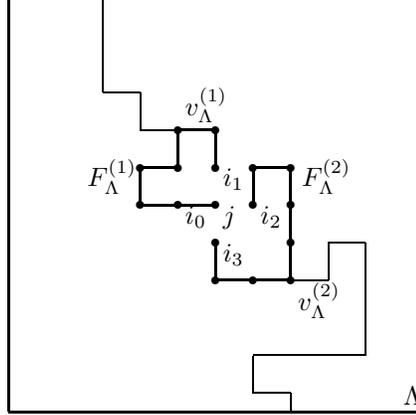
Figure 4. Imitation of the case $d > 4$ in a two-dimensional picture. i_0 is the parent of j, and i_1, i_2, i_3 are the other neighbours of j. There are two connected components intersecting $\{j\} \cup \mathcal{N} = \{j, i_0, \ldots, i_3\}$. The intersections are $A^{(1)} = \{j, i_0, i_i\}$ and $A^{(2)} = \{i_2, i_3\}$. Thick lines represent the subtrees $F_\Lambda^{(1)}$ and $F_\Lambda^{(2)}$. Thin lines indicate the (disjoint) connections from these trees to the boundary.

If $i_1, i_2 \in A^{(\beta)}$, then the order between $t(i_1) - t$ and $t(i_2) - t$ only depends on $(F_\Lambda^{(\beta)}, v_\Lambda^{(\beta)})$. When $i_1 \in A^{(\beta)}$ and $i_2 \in A^{(\gamma)}$ for $\beta \neq \gamma$, then due to (5.2), the order depends on $(F_\Lambda^{(\beta)}, v_\Lambda^{(\beta)})$, $(F_\Lambda^{(\gamma)}, v_\Lambda^{(\gamma)})$ and the difference $X_\Lambda^{(\beta)} - X_\Lambda^{(\gamma)}$. However, for large Λ, we expect that typically

$$\min_{1 \leq \beta < \gamma \leq r} \left| X_\Lambda^{(\beta)} - X_\Lambda^{(\gamma)} \right| \gg \max_{1 \leq \alpha \leq r} \operatorname{diam}(F_\Lambda^{(\alpha)}), \tag{5.3}$$

where diam denotes graph diameter. Therefore, when (5.3) holds, the order between $t(i_1) - t$ and $t(i_2) - t$ will only depend on whether $X_\Lambda^{(\beta)} > X_\Lambda^{(\gamma)}$ or $X_\Lambda^{(\beta)} < X_\Lambda^{(\gamma)}$. Therefore, only the *order* between $\{X_\Lambda^{(\alpha)}\}_{\alpha=1}^r$ is relevant in the limit $\Lambda \to \mathbf{Z}^d$. It turns out that the order is asymptotically uniform over all permutations of $\{1, \ldots, r\}$, given r and $(F_\Lambda^{(\alpha)}, v_\Lambda^{(\alpha)})_{\alpha=1}^r$. In proving that the picture in (5.3) is indeed correct, we rely on Wilson's algorithm, reviewed in Section 5.2 below; see [1] for details.

We can construct a coding of the sandpile model in terms of the USF based on the results described above; stated explicitly in [13]. Let $T^{(1)}, \ldots, T^{(r)}$ denote the components of the USF that have a non-empty intersection with $\{j\} \cup \mathcal{N}$. Define the infinite volume analogues of $(F_\Lambda^{(\alpha)}, v_\Lambda^{(\alpha)})_{\alpha=1}^r$ by letting $(F^{(\alpha)}, v^{(\alpha)})$ be the smallest subtree of the USF containing $(\{j\} \cup \mathcal{N}) \cap T^{(\alpha)}$. We consider

a random ordering $\prec$ on the components of the USF, defined by the following property. Let $T_1(\omega), T_2(\omega), \ldots$ be a list of the components for a given realization ω of the USF. We require that for any $s \geq 2$, and any permutation σ of $\{1, \ldots, s\}$ we have

$$\text{Prob}\big(T_{\sigma(1)}(\omega) \prec \cdots \prec T_{\sigma(s)}(\omega) \mid \omega\big) = \frac{1}{s!} \quad \text{almost surely.}$$

One can realize such a random ordering for example by assigning i.i.d. uniform random variables $U_1, U_2, \ldots$ to the components, and setting

$$T_s \prec T_p \quad \text{if} \quad U_s < U_p, \quad s, p \geq 1.$$

The arguments given above suggest that the limiting distribution of z_j as $\Lambda \to \mathbf{Z}^d$ is determined by $(F^{(\alpha)}, v^{(\alpha)})_{\alpha=1}^r$ *and* $\prec$. Generalizing this idea, one can define a coding of the sandpile configuration in $\mathbf{Z}^d$ in terms of the USF and the ordering $\prec$; see [13] for more details.

5.2. Wilson's algorithm

The following beautiful algorithm was given by Wilson [28] to generate a uniformly chosen spanning tree of an arbitrary finite graph G. Choose a vertex r of G, called the root. Enumerate the rest of the vertices as $v_1, \ldots, v_n$. Start a simple random walk

$$S(1) = \{S_n(1)\}_{n \geq 0}$$

on G from v_1, killed when it hits r. Let π_1 denote the *loop-erasure* of $S(1)$ obtained by chronologically erasing loops, as they are created. π_1 is a self-avoiding path from v_1 to r. Now start a second random walk $S(2)$ from v_2, killed when it hits π_1, and call its loop-erasure π_2 (if $v_2 \in \pi_1$, π_2 is the single point v_2). At the kth step we start a random walk $S(k)$ from v_k, killed when it hits $\cup_{1 \leq j < k} \pi_j$, and call its loop-erasure π_k. It is clear that $\cup_{1 \leq k \leq n} \pi_k$ is a spanning tree of G. It was shown by Wilson [28] that it is also uniformly distributed, regardless of the choices made in the process.

As an example, let $v^{(1)}, \ldots, v^{(r)}$ be fixed vertices, all elements of Λ. Using Wilson's algorithm with root δ, we obtain that the joint law of the paths in T_Λ from $v^{(1)}, \ldots, v^{(r)}$ to δ, conditioned on $v^{(1)}, \ldots, v^{(r)}$ lying in different components, is the same as the joint law of the loop-erased random walks $\pi_1, \ldots, \pi_r$ started at $v^{(1)}, \ldots, v^{(r)}$ conditioned on

$$S(k) \bigcap \Big(\bigcup_{1 \leq j < k} \pi_j \Big) = \emptyset$$

up to the time of hitting δ, $1 \leq k \leq r$.

Property (5.3) is established using the above description of the paths from $v^{(1)}, \ldots, v^{(r)}$ to δ. In this the following theorem about loop-erased random walk

plays an important role. Let $\{S_n\}_{n\geq 0}$ denote simple random walk. Due to transience, the walk visits any vertex finitely often, and hence it makes sense to talk about the loop-erasure of the infinite path. Let $\rho(n)$ denote the number of points remaining of the first n steps of the infinite path after loops are erased.

Theorem 5.1. *Let $d > 4$. There exists a constant $0 < a = a(d) < 1$ such that*

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

Theorem 5.1 is proved in [15]. The theorem says that asymptotically a constant, deterministic fraction of the path is erased. This has the consequence that the loop-erased path behaves qualitatively like random walk when $d > 4$ (it satisfies an invariance principle; see [15]). This allows one to reduce the proof of (5.3) to a random walk question. If $T_\Lambda^{(\beta)}$ denotes the hitting time of the random walk started at $v^{(\beta)}$, $1 \leq \beta \leq r$, then (without being precise about the meaning of the symbol $\approx$)

$$X_\Lambda^{(\beta)} \approx \rho(T_\Lambda^{(\beta)}) \approx aT_\Lambda^{(\beta)}.$$

Therefore one can analyze (5.3) by studying the hitting times $\{T_\Lambda^{(\beta)}\}_{\beta=1}^r$. See [1] for more details.

5.3. Finiteness of avalanches

To prove finiteness of avalanches, we use the decomposition of an avalanche into *waves*, an idea introduced in [11, 12]. Suppose we add a particle at site i, and i becomes unstable. Topple i once, and then topple all other unstable sites, except i (that is, if i becomes unstable again, do not topple it a second time). The set of sites that topple this way is called the first wave. One can verify that during the first wave, all sites topple at most once. If after the first wave, i is again unstable, we topple it a second time, starting the second wave, and so on. When i has toppled the maximum number of times, the avalanche ends (the number of waves is the duration of the avalanche).

Since the expected number of waves is $\mathsf{E}_\nu(i,i) = G(i,i) < \infty$, we have finitely many waves ν-a.s. Therefore, we need to show that each wave is finite a.s. The definition of a wave is reminiscent of the burning process defined through toppling, and in fact waves are equivalent to a certain burning process. Suppose we have a configuration in Λ in which site i is unstable. We start a burning process by burning i. This may cause some of its neighbours to burn, and so on. The set of sites that burn is the same as the set of sites toppled in the wave started at i [12]. Continuing with a second burning process started at the boundary, we obtain a two-component spanning forest representation of the wave [12].

To show that waves remain finite in the limit $\Lambda \to \mathbf{Z}^d$, we show that the component of the two-component spanning forest containing i remains finite in this limit, at least when $d > 4$. The proof of this again uses Wilson's algorithm.

5.4. Tail triviality

The proof of Theorem 4.3 in the case $2 \le d \le 4$ is a rather easy consequence of the tail triviality of the USF [4], and the fact that each height variable only depends on an a.s. finite portion of the USF. The main technical difficulty in the case $d > 4$ is the presence of the random ordering, which prevents us from using tail triviality of the USF directly. We show the following statement that is equivalent to tail triviality. For any local event A and $\varepsilon > 0$ there is a finite set $K = K(A, \varepsilon)$ such that for any $B \in \mathcal{F}_{\mathbf{Z}^d \setminus K}$:

$$|\nu(A \cap B) - \nu(A)\nu(B)| \le \varepsilon. \tag{5.4}$$

Translated into the language of the USF, establishing (5.4) roughly amounts to showing the following. Given the configuration of the USF outside some large set Λ, and given the order of the components, the conditional distribution of $(F_\Lambda^{(\alpha)}, v_\Lambda^{(\alpha)})_{\alpha=1}^r$ and the order of the components containing them, is close to its unconditional distribution. See [13] for details.

Acknowledgments

I thank Siva Athreya and Frank Redig for fruitful collaborations leading to the results discussed in this paper. I also thank Thierry Gobron and Ellen Saada for organizing a very interesting conference and providing the incentive to write these notes. I am grateful to Ellen Saada and the referee for helpful comments. The first version of this paper was written at CWI, Amsterdam.

References

[1] S.R. Athreya and A.A. Járai (2004) Infinite volume limit for the stationary distribution of Abelian sandpile models. *Commun. Math. Phys.* **249**, 197–213. An Erratum is to appear, correcting some mistakes discovered in the paper.

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

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

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

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

[6] D. DHAR (1999) The Abelian sandpile and related models. *Physica A* **263**, 4–25.

[7] D. DHAR (1999) Studying self-organized criticality with exactly solved models. Preprint. http://arXiv.org/abs/cond-mat/9909009.

[8] D. DHAR AND S.N. MAJUMDAR (1990) Abelian sandpile model on the Bethe lattice. *J. Phys. A* **23**, 4333–4350.

[9] D. DHAR, P. RUELLE, S. SEN AND D.N. VERMA (1995) Algebraic aspects of Abelian sandpile models. *J. Phys. A* **28**, 805–832.

[10] E.V. IVASHKEVICH AND V.B. PRIEZZHEV (1998) Introduction to the sandpile model. *Physica A* **254**, 97–116.

[11] E.V. IVASHKEVICH, D.V. KTITAREV AND V.B. PRIEZZHEV (1994) Critical exponents for boundary avalanches in two-dimensional Abelian sandpile. *J. Phys. A* **27**, L585–L590 (1994).

[12] E.V. IVASHKEVICH, D.V. KTITAREV AND V.B. PRIEZZHEV (1994) Waves of topplings in an Abelian sandpile. *Physica A* **209**, 347–360.

[13] A.A. JÁRAI AND F. REDIG (2004) Infinite volume limits of high-dimensional sandpile models. Preprint. To appear in *Probab. Theory and Relat. Fields*.

[14] H.J. JENSEN (2000) *Self-organized Criticality. Emergent Complex Behavior in Physical and Biological Systems.* Cambridge Lect. Notes Phys. **10**. Cambridge University Press.

[15] G.F. LAWLER (1996) *Intersections of Random Walks.* Birkhäuser.

[16] C. MAES, F. REDIG AND E. SAADA (2002) The Abelian sandpile model on an infinite tree. *Ann. Prob.* **30**, 2081–2107.

[17] C. MAES, F. REDIG AND E. SAADA (2004) The infinite volume limit of dissipative abelian sandpiles. *Commun. Math. Phys.* **244**, 395–417.

[18] C. MAES, F. REDIG AND E. SAADA (2004) Abelian sandpiles in infinite volume. Preprint. To appear in *Sankhyā, The Indian Journal of Statistics*.

[19] C. MAES, F. REDIG, E. SAADA AND A. VAN MOFFAERT (2000) On the thermodynamic limit for a one-dimensional sandpile process. *Markov Processes Relat. Fields* **6**, 1–22.

[20] S.N. MAJUMDAR AND D. DHAR (1991) Height correlations in the Abelian sandpile model. *J. Phys. A* **24**, L357–L362.

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

[22] R. MEESTER, F. REDIG AND D. ZNAMENSKI (2002) The Abelian sandpile; a mathematical introduction. *Markov Processes Relat. Fields* **7**, 509–523.

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

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

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

[26] E.R. Speer (1993) Asymmetric Abelian sandpile models. *J. Stat. Phys.* **71**, 61–74.

[27] F. Spitzer (1976) *Principles of Random Walk*. Graduate Texts in Mathematics **34**. Springer-Verlag.

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