Ladder Sandpiles

A.A. Járai[1,*] and R. Lyons[2,†]

[1] Carleton University, School of Mathematics and Statistics, 1125 Colonel By Drive, Ottawa, ON K1S 5B6, Canada. E-mail: jarai@math.carleton.ca
[2] Dept. of Mathematics, Indiana University, Bloomington, IN 47405-5701 USA. E-mail: rdlyons@indiana.edu

Abstract. We study Abelian sandpiles on graphs of the form $G \times I$, where G is an arbitrary finite connected graph, and $I \subset \mathbf{Z}$ is a finite interval. We show that for any fixed G with at least two vertices, the stationary measures $\mu_I = \mu_{G \times I}$ have two extremal weak limit points as $I \uparrow \mathbf{Z}$. The extremal limits are the only ergodic measures of maximum entropy on the set of infinite recurrent configurations. We show that under any of the limiting measures, one can add finitely many grains in such a way that almost surely all sites topple infinitely often. We also show that the extremal limiting measures admit a Markovian coding.

1. Introduction

The sandpile model was introduced by Bak, Tang and Wiesenfeld [5, 6], who used it to illustrate the idea of self-organized criticality [3], a concept that became influential in theoretical physics [20]. The name Abelian sandpile model (ASM) was coined by Dhar [10], who discovered its Abelian property. The model also appeared independently in the combinatorics literature, where it is known as the chip-firing game, introduced in [7]. One of the remarkable features of the ASM is that its simple local rules give rise to complex long-range dynamics. See [13] for an overview.

Recently, a number of papers were devoted to sandpiles on infinite graphs, obtained as limits of sandpiles on finite subgraphs [2, 19, 23, 24, 26]. See the reviews [18, 25, 29] for the main ideas of these developments. A natural approach to studying sandpiles in infinite volume is the following. Start with the stationary measure μ_V of the model on a finite subgraph V, and characterize the set of weak limit points of $\{\mu_V : V \text{ finite}\}$. Then study avalanches on infinite configurations under the limiting measures, and construct a dynamics for the infinite system, if at all possible. The papers mentioned carry out this program to various degrees for the following infinite graphs: $\mathbf{Z}$; an infinite regular tree; $\mathbf{Z}^d$, $d \geq 2$ with or without dissipation; and finite-width strips with dissipation. Already the first step, determining the limiting behaviour of μ_V, is usually non-trivial. For each infinite graph mentioned above, there is a unique limit point, and a number of different techniques have been employed to show this. Currently there seems to be no unified method that applies to a general infinite graph.

In this paper, we study the sandpile measures on graphs that are the product of a finite connected graph G with a finite interval $I \subset \mathbf{Z}$, with particular view towards the limiting behaviour as $|I| \to \infty$. We call these "ladder graphs", where the "rungs" of the ladder consist of copies of G. The only dissipative sites in our model are the ones at the two ends of the ladder, that is, the sites in $G \times \{\text{endpoints of } I\}$. Hence, even when G is a finite interval in $\mathbf{Z}$, our models are different from the dissipative strips studied in [24].

When G is a single vertex, the model lives on an interval $I \subset \mathbf{Z}$. In this well-known case, μ_I can be found explicitly, and its limiting behaviour is trivial: the limit is concentrated on a single configuration with constant height 2. However, as was observed in [1], the behaviour of this model is atypical of general one-dimensional sandpiles. Already the simplest modifications, such as the "decorated chains" studied in [1], give rise to non-trivial limits with positive entropy.

Not surprisingly, the limiting behaviour is also non-trivial in our case, provided G has more than one vertex. As we show in Theorem 3.2 in Section 3.2, when $|G| \geq 2$, the set of weak limit points consists of all convex combinations of two different extremal measures μ^{L} and μ^{R}, which are related by a reflection of $\mathbf{Z}$. These measures arise from restricting the burning algorithm to act exclusively from the left or the right, and we call them the left- and right-burnable measures. In the case of $|G| = 1$, the left- and right-burnable measures happen to coincide.

In Theorem 3.1 in Section 3.2 we show that μ^{L} and μ^{R} are the only two ergodic measures of maximum entropy on the set of infinite recurrent configurations, and that there is a unique measure μ^{S} of maximum entropy that is invariant under reflection. It would be interesting to see whether there is a unique measure of maximum entropy for the infinite graphs studied earlier (where the weak limits are unique).

For most of our arguments on weak limits, some quite general properties of the model are sufficient. For example, existence of the limit of left- and right-burnable measures follows from the existence of renewals: if all sites in $G \times \{i\}$ have the maximum possible height for a fixed $i \in I$, then the subconfigurations to the left and right of i are conditionally independent.

Given an infinite configuration, we can ask what happens if particles are added and then the configuration is relaxed. First, in the case of $\mathbf{Z}$, it is easy to see that if we add a single grain to the system (having constant height 2), then every site topples infinitely often. On the ladder $G \times \mathbf{Z}$ with $|G| \geq 2$, finite avalanches do occur with positive probability. However, as we show in Section 4, it is possible to add a fixed number of grains in such a way that almost surely every site topples infinitely often, with respect to any of the limiting measures. Hence, there is no sensible dynamics for $G \times \mathbf{Z}$ in general.

Open question. *Does the probability of infinitely many topples at $(0,0)$, when adding 1 grain to $(0,0)$, tend to 0 for $G = \mathbf{Z}_n$ and $n \to \infty$? Here $\mathbf{Z}_n$ is the cycle of length n.*

The measure μ^{L} (and μ^{R}) can be regarded as a subshift on a finite alphabet (that depends on G), by grouping sites on each copy of G together. The set of recurrent configurations are characterized in terms of finite forbidden words, and there is an infinite number of constraints. As our results show (see Lemma 3.4), the number of constraints grows at a rate strictly smaller than the topological entropy. Hence the set of left-burnable configurations is a subshift of quasi-finite type, in the terminology of [8]. In fact, our subshift turns out to be more special. As we show in Section 5, it admits a Markovian coding, and hence it is a sofic shift [22, Theorem 3.2.1]. We note that the set of all recurrent configurations is also a sofic shift (by arguments similar to those for Lemma 3.4). However, since recurrent configurations lead to mixtures of μ^{L} and μ^{R}, we study only the latter in detail. An alternative approach to our results in Section 3 would be to analyze the Markovian coding obtained in Section 5. However, we prefer to present more direct arguments.

We will assume throughout that the reader is familiar with the basic properties of the ASM that can be found in [11–13, 17, 28].

2. Models considered

Throughout, G will be an arbitrary fixed finite connected graph. For $n \leq m$, let $I_{n,m}$ denote the graph on the vertex set $\{n, \ldots, m\}$ with nearest neighbour edges. Let $\deg_G(x)$ denote the degree of a vertex x in G. We consider Abelian sandpiles [28] defined on the product graph $\Lambda_{n,m} := G \times I_{n,m}$ (whose edges join vertices (x, k) and (y, ℓ) when either $x \sim y$ in G and $k = \ell$ or $x = y$ and $|k - \ell| = 1$). We are primarily interested in the limit sandpiles as $n \to -\infty$ and $m \to \infty$ that live on the graph $\Lambda := G \times \mathbf{Z}$. For convenience, we also introduce

$\Lambda_{-\infty,m}$ and $\Lambda_{n,\infty}$ with the obvious meaning. We refer to $G \times \{k\}$ as the *rung* at k.

We let Δ denote the graph Laplacian on Λ, that is, the following matrix indexed by vertices in Λ:

$$\Delta_{uv} := \begin{cases} \deg_G(x) + 2 & \text{if } u = v = (x, k); \\ -1 & \text{if } u \text{ and } v \text{ are neighbours}; \\ 0 & \text{otherwise.} \end{cases}$$

For finite vertex-subsets $V \subset \Lambda$ that induce a connected subgraph, we let Δ_V denote the restriction of Δ to the pairs $(u, v) \in V \times V$. In other words, a sink site is added to V, and each $u \in V$ is connected to the sink by $\Delta_{uu} - \deg_V(u)$ edges.

We are interested in the sandpile with toppling matrix $\Delta_{\Lambda_{n,m}}$. We will study the case when $G = I_{0,1}$ quite explicitly for illustration.

The space of stable configurations on a set $V \subset \Lambda$ is

$$S_V := \prod_{u \in V} \{1, \ldots, \Delta_{uu}\}.$$

We write $S := S_\Lambda$. For a convenient notation, we define $m(x) := \deg_G(x) + 2$, $x \in G$, which is the maximum allowed height at a site $u = (x, k)$. We write Ω_V for the set of recurrent configurations [28] on V when V is finite. We define

$$\Omega := \Omega_\Lambda := \{\text{recurrent configurations on } \Lambda\}$$
$$:= \{\eta \in S : \eta_W \in \Omega_W \text{ for all finite } W \subset \Lambda\}.$$

For $V \subset \Lambda$, if $\Lambda \setminus V$ has a connected component fully infinite to the left (that is, containing $\Lambda_{-\infty,m}$ for some m), we denote that connected component V^-. We similarly define V^+ to the right (which may coincide with V^-). We define the *left (interior) boundary* of V as

$$\partial_0^{\text{L}} V := \left\{ v \in V : \ v \text{ has a neighbour in } V^- \right\}.$$

We define $\partial_0^{\text{R}} V$ analogously.

3. Description of recurrent configurations

3.1. Left- and right-burnable measures

We define a one-sided version of the burning algorithm [28].

Definition 3.1. Let $V \subset \Lambda$ be finite. A configuration $\eta \in S_V$ is called *left-burnable* if there is an enumeration $v_1, \ldots, v_{|V|}$ of V such that

   (i) $v_i \in \partial_0^{\mathrm{L}}(V \setminus \{v_1, \ldots, v_{i-1}\}), 1 \le i \le |V|$;

  (ii) $\eta(v_i) > \Delta_{v_i v_i} - |\{u \in (V \setminus \{v_1, \ldots, v_{i-1}\})^- : u \sim v_i\}|$.

Note that this is the usual burning rule with the restriction that only sites in the left boundary can be burnt. When $\Lambda \setminus V$ is connected, the rule becomes identical to the usual burning rule. We denote by Ω_V^{L} the set of left-burnable configurations on V. We define right-burnable configurations and Ω_V^{R} analogously.

Lemma 3.1. *Let $V \subset \Lambda$ be finite. We have $\Omega_V^{\mathrm{L}} \subset \Omega_V$. If $\eta \in \Omega_V^{\mathrm{L}}$ and $W \subset V$, then $\eta_W \in \Omega_W^{\mathrm{L}}$. The same holds for Ω_V^{R}.*

Proof. The sequence $v_1, \ldots, v_{|V|}$ required by Definition 3.1 is a valid burning sequence in the ordinary burning algorithm, since

$$|\{u \in (V \setminus \{v_1, \ldots, v_{i-1}\})^- : u \sim v_i\}| \le |\{u \in (V \setminus \{v_1, \ldots, v_{i-1}\})^c : u \sim v_i\}|.$$

Therefore, $\Omega_V^{\mathrm{L}} \subset \Omega_V$. For $W \subset V$, let $w_1, \ldots, w_{|W|}$ be the enumeration of W in the order inherited from the enumeration of V. Since

$$|\{u \in (V \setminus \{v_1, \ldots, v_{i-1}\})^- : u \sim v_i\}| \le |\{u \in (W \setminus \{v_1, \ldots, v_{i-1}\})^- : u \sim v_i\}|,$$

this is a valid left-burning sequence for η_W. $\qquad\square$

Definition 3.2. For $V = \Lambda$, a configuration $\eta \in S$ is called *left-burnable* if η_W is left-burnable for every finite $W \subset \Lambda$. *Right-burnable* configurations are defined analogously. We write Ω^{L} and Ω^{R} for the sets of these configurations. We write $\Omega^{\mathrm{S}} := \Omega^{\mathrm{L}} \cap \Omega^{\mathrm{R}}$.

Definition 3.3. Let $\mu_{n,m}$ denote the uniform measure on the set of recurrent configurations on $\Lambda_{n,m}$. We denote by $\mu_{n,m}^{\mathrm{L}}$ the uniform measure on left-burnable configurations on $\Lambda_{n,m}$ and define $\mu_{n,m}^{\mathrm{R}}$ and $\mu_{n,m}^{\mathrm{S}}$ analogously.

In order to illustrate some of the results to come, we explicitly describe left-burnable configurations in the simplest non-trivial case $G = I_{0,1}$.

Lemma 3.2. *Assume $G = I_{0,1}$. A configuration $\eta \in \Omega_{n,m}$ is left-burnable if and only if the following three conditions hold:*

1 *Each rung contains a 3.*

2 *If the rung at k is $(3,1)$, then no rung other than $(3,2)$ can occur to the right of k before a $(3,3)$ occurs. That is, the rungs at $k, k+1, \ldots$ are of the form:*

$$
\begin{array}{ccccccc}
3 & 3 & \ldots & 3 & 3 & \ldots \\
1 & 2 & \ldots & 2 & 3 & \ldots
\end{array}
\tag{3.1}
$$

with the possibility that there is no $(3,2)$ rung at all, and the exception that the $(3,3)$ may be missing if the right end of $\Lambda_{n,m}$ was reached.

    3 *If the rung at k is $(1,3)$, then no rung other than $(2,3)$ can occur to the right of k before a $(3,3)$ occurs.*

The same holds for right-burnable configurations with left and right interchanged.

Proof. By symmetry, we may restrict to the left-burnable case. It is straightforward to verify that a configuration satisfying 1–3 in the Lemma is left-burnable. Namely, the configuration can be burnt rung-by-rung, except when a $(3,1)$ or a $(1,3)$ is encountered. In the latter case, observe that the configuration in (3.1) is left-burnable (as well as the one obtained by exchanging the rows).

    Assume now that we are given a left-burnable configuration, and we show that 1–3 hold. The proof is by induction on the number $N = m-n+1$ of rungs. The case $N = 1$ is trivial. Assume now that $N > 1$ and that the statement holds whenever the number of rungs is less than N. Observe that the leftmost rung has to contain a 3, otherwise the burning cannot start.

    <u>Case 1</u>. *The leftmost rung is $(3,3)$, $(3,2)$ or $(2,3)$.* Then without loss of generality, we may assume that the burning starts with removing the leftmost rung. Since the $N-1$ remaining rungs are left-burnable, the induction hypothesis implies the claim.

    <u>Case 2</u>. *The leftmost rung is $(3,1)$ or $(1,3)$.* We may assume the leftmost rung is $(3,1)$. Then, by the burning procedure, the next rung is of the form $(3,z)$. If $z = 3$, we can use the induction hypothesis for $N-2$. The value $z = 1$ leads to a forbidden subconfiguration 1 1. If $z = 2$, we can iterate the present argument until a rung of the form $(3,3)$ is reached, noting that configurations of the form 1 2 ... 2 1 are forbidden.   □

    It follows from the description in Lemma 3.2 that $(3,3)$ rungs are renewals, that is, given that rung k is $(3,3)$, the subconfigurations to the left and right are conditionally independent for the appropriate measure $\mu_{n,m}^{\bullet}$. The analogous statement holds for maximal rungs on a general graph, and we prove this next.

    The following terminology will be useful. Let $\mathcal{C} := \mathcal{C}(G) := \Omega_{\Lambda_{0,0}}^{\mathrm{L}}$ denote the set of left-burnable configurations on a single rung. We claim this is the same as the set of all recurrent configurations on $G \times \{0\}$ with at least one $x \in G$ such that $\eta(x,0) = m(x)$. Indeed, no burning will occur without such an x, and with such an x, we can left-burn $(x,0)$, and once this is done, left-burning becomes equivalent to ordinary burning (i.e., burning from both sides) since the left and right boundaries merge. By the same reasoning, $\mathcal{C}$ is also the set of right-burnable configurations at a single rung. By abuse of notation, we regard $\mathcal{C}$ as a set of configurations on any particular rung. For $\eta \in \Omega_{n,m}^{\mathrm{L}}$, let $C_k := C_k(\eta) := \eta_{\Lambda_{k,k}}$ denote the rung at k, which is in $\mathcal{C}$ by Lemma 3.1. Let $C^{\mathrm{max}} := C^{\mathrm{max}}(G) \in \mathcal{C}$ denote the configuration on G defined by $C^{\mathrm{max}}(x) := m(x)$, $x \in G$. The configuration C^{max} is the maximal configuration that can occur on a rung.

Lemma 3.3 (Renewals). *For the measures $\mu_{n,m}^{\mathrm{L}}$, $\mu_{n,m}^{\mathrm{R}}$ and $\mu_{n,m}^{\mathrm{S}}$, maximal rungs are renewals, that is, given $C_k = C^{\mathrm{max}}$, the subconfigurations to the left and right of rung k are conditionally independent.*

Proof. First consider the left-burnable measure. Let $\eta \in \Omega_{n,m}^{\mathrm{L}}$, and assume that $C_k(\eta) = C^{\mathrm{max}}$. By Lemma 3.1, both $\eta_{\Lambda_{n,k-1}}$ and $\eta_{\Lambda_{k+1,m}}$ are left-burnable. We need to show that the two vary independently, that is, for any $\xi \in \Omega_{n,k-1}^{\mathrm{L}}$ and $\zeta \in \Omega_{k+1,m}^{\mathrm{L}}$, we have $\eta' = \xi \vee C^{\mathrm{max}} \vee \zeta \in \Omega_{n,m}^{\mathrm{L}}$, where $\vee$ indicates concatenation. Start left-burning on η'. Since ξ is left-burnable, there will be a first time when a site $(x, k-1)$ is burnt. When this happens, we can fully burn rung k. After rung k is burnt, both the rest of ξ and all of ζ can be burnt, because they are left-burnable. Hence η' is left-burnable, and $\mu_{n,m}^{\mathrm{L}}$ has the renewal property since it is uniform on $\Omega_{n,m}^{\mathrm{L}}$.

The statement for $\mu_{n,m}^{\mathrm{R}}$ follows by symmetry. The statement for $\mu_{n,m}^{\mathrm{S}}$ can be proved by a very similar argument, now showing that burning from both left and right can be performed. $\qquad\square$

In order to investigate weak convergence of the finite-volume measures, we are going to use some comparisons between the growth rates (topological entropies) of certain sets of configurations. This is formulated in the lemma below.

Let $a_n := |\Omega_{1,n}^{\mathrm{L}}|$. By Lemma 3.1, we have

$$a_{n+m} \le |\Omega_{1,n}^{\mathrm{L}}||\Omega_{n+1,n+m}^{\mathrm{L}}| = a_n a_m.$$

Therefore,

$$h^{\mathrm{L}} := \lim_{n\to\infty} \frac{1}{n} \log |\Omega_{1,n}^{\mathrm{L}}| = \inf_{n\ge 1} \frac{1}{n} \log a_n \qquad (3.2)$$

exists. The limit h^{L} is the topological entropy $h_{\mathrm{top}}(\Omega^{\mathrm{L}})$ of Ω^{L} with respect to translations [21]. For any $\delta > 0$, there exists $C = C(\delta)$ such that

$$\exp\{h^{\mathrm{L}} n\} \le a_n \le C \exp\{h^{\mathrm{L}} n(1 + \delta)\}.$$

By symmetry, $h^{\mathrm{L}} = h^{\mathrm{R}} = h_{\mathrm{top}}(\Omega^{\mathrm{R}})$. We also define $s_n := |\Omega_{1,n}^{\mathrm{S}}|$, and again, by submultiplicativity, we have

$$h^{\mathrm{S}} := \lim_{n\to\infty} \frac{1}{n} \log s_n = h_{\mathrm{top}}(\Omega^{\mathrm{S}}).$$

We further define the spaces

$$\Omega^{\mathrm{L},0} := \{\eta \in \Omega^{\mathrm{L}} : C_k(\eta) \ne C^{\mathrm{max}}, -\infty < k < \infty\}$$
$$\Omega_{n,m}^{\mathrm{L},0} := \{\eta \in \Omega_{n,m}^{\mathrm{L}} : C_k(\eta) \ne C^{\mathrm{max}}, n \le k \le m\},$$

and we define $\Omega^{S,0}$ and $\Omega_{n,m}^{S,0}$ analogously in the symmetric case. Let $b_n := |\Omega_{1,n}^{\mathrm{L},0}|$, $r_n := |\Omega_{1,n}^{S,0}|$, $h^{\mathrm{L},0} := \lim_{n\to\infty}(1/n)\log b_n = h_{\mathrm{top}}(\Omega^{\mathrm{L},0})$ and $h^{S,0} := \lim_{n\to\infty}(1/n)\log r_n = h_{\mathrm{top}}(\Omega^{S,0})$.

Lemma 3.4. *Assume that G is not a single vertex. Then*

(i) $0 < h^{L,0} < h^L$;

(ii) $0 < h^{S,0} < h^S$;

(iii) $h^S < h^L$.

Proof. For $x \in G$, define the rung

$$C^x(z) := \begin{cases} m(x) - 1 & z = x, \\ m(z) & z \neq x. \end{cases}$$

It is straightforward to check that since G consists of more than one vertex, $C^x \in \mathcal{C}$. Now let $x, y \in G$, $x \neq y$. Any sequence consisting exclusively of rungs C^x and C^y is both left- and right-burnable. Hence $\log 2 \leq h^{S,0} \leq h^{L,0}$.

For $0 \leq k \leq n$ we select k of the rungs. Consider the function that changes these k rungs of an $\eta \in \Omega^{L,0}_{1,n}$ to $C^{\max}$. Any configuration so obtained is in $\Omega^L_{1,n}$ and has at most $|\mathcal{C}|^k$ preimages. Therefore, the number of different new configurations obtained is at least $b_n/|\mathcal{C}|^k$. Summing over k and all choices of k rungs, we have

$$a_n \geq b_n \sum_{k=0}^{n} \binom{n}{k} \frac{1}{|\mathcal{C}|^k} = b_n \left(1 + \frac{1}{|\mathcal{C}|}\right)^n.$$

Hence, $h^L \geq h^{L,0} + \log(1 + |\mathcal{C}|^{-1})$. This proves (i). By a similar argument, $h^S \geq h^{S,0} + \log(1 + |\mathcal{C}|^{-1})$, which proves (ii).

The argument to prove (iii) is also similar: Note that when $G = I_{0,1}$, Lemma 3.2 implies that the sequence of rungs $(3,3), (2,3), (3,1), (3,3)$ is left-burnable, but not right-burnable. We can adapt this observation to general G. Let $x \sim y \in G$, and define the rungs

$$C_1(z) := \begin{cases} m(z) - 1 & z \neq y, \\ m(y) & z = y, \end{cases} \qquad C_2(z) := \begin{cases} m(z) - 1 & z \neq x, y, \\ m(x) & z = x, \\ 1 & z = y. \end{cases}$$

We claim that the sequence $\xi = C^{\max}, C_1, C_2, C^{\max}$ is left-burnable, but not right-burnable. In case of left-burning, $C^{\max}$ burns first, then site y burns in C_1, and after that the rest of C_1 can be burnt. Now site x burns in C_2, and after this the $C^{\max}$ rung to the right burns. This makes the rest of the sites but y in C_2 burnable; finally, site y in C_2 can be burnt. In case of right-burning, the $C^{\max}$ rung on the right can be burnt. After this, site x in C_2 can be burnt. This may make other sites in C_2 burnable. However, crucially, y in C_2 cannot be burnt (since it has a neighbour in C_1), and no site in C_1 can be burnt, since burning could only start at y, which is "blocked" by the 1 in C_2.

Assume now that $\eta \in \Omega^{\mathrm{S}}_{1,4n}$, and subdivide $[1, 4n]$ into n intervals of length 4. Consider the mapping that replaces the rungs at a fixed set of k of these intervals by ξ. The configurations obtained are in $\Omega^{\mathrm{L}}_{1,4n}$, and since ξ is not right-burnable, they are not in $\Omega^{\mathrm{S}}_{1,4n}$. The number of preimages of a given element of $\Omega^{\mathrm{L}}_{1,4n}$ is at most $|\mathcal{C}|^{4k}$. Hence we get

$$a_{4n} \geq s_{4n} \sum_{k=0}^{n} \binom{n}{k} \frac{1}{|\mathcal{C}|^{4k}} = s_{4n} \left(1 + \frac{1}{|\mathcal{C}|^4}\right)^n.$$

This implies (iii).   □

3.2. Weak limits

Lemma 3.5. *The weak limits*

$$\mu^{\mathrm{L}} := \lim_{\substack{n \to -\infty \\ m \to \infty}} \mu^{\mathrm{L}}_{n,m}, \qquad \mu^{\mathrm{R}} := \lim_{\substack{n \to -\infty \\ m \to \infty}} \mu^{\mathrm{R}}_{n,m}, \qquad \mu^{\mathrm{S}} := \lim_{\substack{n \to -\infty \\ m \to \infty}} \mu^{\mathrm{S}}_{n,m}$$

exist. The limit measures μ^{L}, μ^{R} and μ^{S} are concentrated on Ω^{L}, Ω^{R} and Ω^{S} (respectively).

Proof. We first strengthen (3.2) to show that with $\lambda = \exp(-h^{\mathrm{L}})$, the limit $\lim_{n \to \infty} \lambda^n a_n$ exists and is positive. Lemma 3.3 implies the renewal equation:

$$a_n = b_n + \sum_{k=1}^{n} b_{k-1} a_{n-k}, \qquad n \geq 0, \tag{3.3}$$

where we set $a_0 = 1$, $b_0 = 1$. Let

$$F(z) := \sum_{n=0}^{\infty} a_n z^n \qquad \text{and} \qquad G(z) := \sum_{n=0}^{\infty} b_n z^n$$

be the generating functions of $\{a_n\}_{n\geq 0}$ and $\{b_n\}_{n\geq 0}$. The radius of convergence of F is $\lambda = \exp(-h^{\mathrm{L}})$ and that of G is $\exp(-h^{\mathrm{L},0}) > \lambda$. The relation (3.3) implies

$$F(z) = \frac{G(z)}{1 - zG(z)}, \qquad 0 \leq z < \lambda.$$

Since G is analytic in a disc of radius larger than λ, but F has a singularity on the circle $|z| = \lambda$, we need to have $1 = \lim_{z \nearrow \lambda} zG(z) = \lambda G(\lambda)$. It follows that $p_n := \lambda^n b_{n-1}$, $n \geq 1$, is a probability distribution, and with $c_n := \lambda^n a_{n-1}$, (3.3) has the probabilistic form

$$c_{n+1} = p_{n+1} + \sum_{k=1}^{n} p_k c_{n-k}, \qquad n \geq 0. \tag{3.4}$$

By the Renewal Theorem [15, page 330], we have $\lim_{n \to \infty} c_n = (\sum_{k \geq 1} k p_k)^{-1}$. Hence we have

$$\lim_{n \to \infty} \lambda^n a_n = \lim_{n \to \infty} c_n / \lambda = \frac{1}{\lambda \sum_{k \geq 1} k p_k} = \frac{1}{\lambda^2 [\lambda G(\lambda)]'} =: \alpha > 0. \qquad (3.5)$$

We are ready to establish the existence of μ^{L}. Fix $k \geq 1$ and an elementary cylinder event depending on the rungs $-k, \ldots, k$. That is, we fix $\eta_0 \in \Omega^{\mathrm{L}}_{-k,k}$, and let $E = E(\eta_0)$ denote the event that the subconfiguration in rungs $-k, \ldots, k$ equals η_0. We need to show that

$$\lim_{\substack{n \to -\infty \\ m \to \infty}} \mu^{\mathrm{L}}_{n,m}(E) =: \mu^{\mathrm{L}}(E)$$

exists. We first show that for $N > k$ large enough and $-n, m > N$, the event

$$A(N) = \{\exists \text{ renewal in } [-N, -k-1] \text{ and in } [k+1, N]\}$$

occurs with high $\mu^{\mathrm{L}}_{n,m}$-probability. Indeed, letting η denote a random variable with law $\mu^{\mathrm{L}}_{n,m}$ and using Lemma 3.4, we get

$$\mu^{\mathrm{L}}_{n,m}(\text{no renewal in } [k+1, N]) = \mu^{\mathrm{L}}_{n,m}\left(\eta_{\Lambda_{k+1,N}} \in \Omega^{\mathrm{L},0}_{k+1,N}\right)$$

$$\leq \frac{|\Omega^{\mathrm{L}}_{n,k}||\Omega^{\mathrm{L},0}_{k+1,N}||\Omega^{\mathrm{L}}_{N+1,m}|}{|\Omega^{\mathrm{L}}_{n,m}|}$$

$$\leq C e^{-\delta(N-k)}$$

for some $\delta > 0$ and $C = C(\delta)$ for all large N. This implies

$$\mu^{\mathrm{L}}_{n,m}(A(N)^c) \leq 2C e^{-\delta(N-k)}, \qquad -n, m > N. \qquad (3.6)$$

On the event $A(N)$, let

$$\tau := \text{leftmost renewal in } [k+1, N]$$

and

$$\sigma := \text{rightmost renewal in } [-N, -k-1].$$

We also define

$$u(s, t, E) := \Big|\big\{\xi \in \Omega^{\mathrm{L}}_{s+1,t-1} : \xi_{\Lambda_{-k,k}} = \eta_0 \text{ and no renewal}$$
$$\text{in } [s+1, -k-1] \cup [k+1, t-1]\big\}\Big|$$

for $-N \leq s \leq -k-1$ and $k+1 \leq t \leq N$. Considering the values of σ and τ and counting configurations, we can write

$$\mu_{n,m}^{\mathrm{L}}\big(E,\, A(N)\big) = \sum_{t=k+1}^{N} \sum_{s=-N}^{-k-1} \mu_{n,m}^{\mathrm{L}}(E,\, \tau = t,\, \sigma = s)$$

$$= \sum_{t=k+1}^{N} \sum_{s=-N}^{-k-1} \frac{a_{s-n} u(s,t,E) a_{m-t}}{a_{m-n+1}}.$$

Using (3.5), we have

$$\lim_{\substack{n \to -\infty \\ m \to \infty}} \mu_{n,m}^{\mathrm{L}}\big(E,\, A(N)\big) = \alpha \sum_{t=k+1}^{N} \sum_{s=-N}^{-k-1} \lambda^{t-s+1} u(s,t,E).$$

Letting $N \to \infty$ and applying (3.6), we deduce that

$$\lim_{\substack{n \to -\infty \\ m \to \infty}} \mu_{n,m}^{\mathrm{L}}(E) = \alpha \sum_{t=k+1}^{\infty} \sum_{s=-\infty}^{-k-1} \lambda^{t-s+1} u(s,t,E) =: \mu^{\mathrm{L}}(E).$$

The statement for μ^{R} follows by symmetry. In the case of μ^{S}, the proof follows a very similar line. □

Remark 3.1. It is not hard to extend the proof above to show that $\lim_{V \nearrow \Lambda} \mu_V^{\mathrm{L}} = \mu^{\mathrm{L}}$ (and similarly for μ^{R} and μ^{S}).

Lemma 3.6. *Maximal rungs are renewals for the measures μ^{L}, μ^{R} and μ^{S}, and these measures are ergodic. If G is not a single vertex, then μ^{L} and μ^{R} are not symmetric under reflection, while μ^{S} is. The measures μ^{L} and μ^{R} are reflections of each other.*

Proof. The renewal property follows from Lemma 3.3 by passing to the limit. Ergodicity follows from the existence of renewals. The configuration ξ given in the proof of Lemma 3.4 (iii) shows that $\mu^{\mathrm{L}} \neq \mu^{\mathrm{R}}$. □

Remark 3.2. It follows from general arguments that μ^L and μ^R have maximal entropy. For example, by a counting argument one can show that $|\Omega_{1,n}|$ has exponential growth rate h^L, and this allows one to adapt the argument of [4, Proposition 1.12 (ii)]. Below we show that there are no other measures of maximal entropy.

Theorem 3.1. *The only two ergodic measures of maximum entropy on Ω are μ^{L} and μ^{R}. The unique symmetric ergodic measure of maximum entropy is μ^{S}.*

Proof. There exists a measure μ of maximum entropy on Ω [21]. By ergodic decomposition, we may assume that μ is ergodic. We show that in this case either $\mu = \mu^{\mathrm{L}}$ or $\mu = \mu^{\mathrm{R}}$, which shows that these are the only two ergodic measures of maximum entropy.

We first show that $\mu\{C_0 = C^{\mathrm{max}}\} > 0$. To see this, note that increasing the height of any site will never create a forbidden subconfiguration. Suppose we had zero probability of seeing any C^{max} rungs. Consider the measure μ' obtained by changing each rung to C^{max} independently with some small probability $0 < \varepsilon < 1$. Then μ' is also ergodic, and a straightforward computation shows that its measure theoretic entropy is $h(\mu') = (1 - \varepsilon)h(\mu) + H(\varepsilon)$, where $H(\varepsilon) = -\varepsilon \log \varepsilon - (1 - \varepsilon) \log(1 - \varepsilon)$. Hence for ε sufficiently small, $h(\mu') > h(\mu)$, a contradiction.

Consider now the sequence of "blocks" between successive C^{max} rungs. These form a stationary sequence. Again, by maximum entropy, the blocks have to be independent. Indeed, if they were not, consider the measure μ', where the blocks are i.i.d. and each block has its μ-distribution. Let $E := \{C_0 = C^{\mathrm{max}}\}$. Since the expected length of a block is the same in μ' and μ, we have $\mu'(E) = \mu(E)$. The measure-preserving maps induced by E [9, Chapter 1, §5] are the translations of blocks, with invariant measures μ_E and μ'_E (the normalized restrictions of μ and μ' to E). Since μ'_E is i.i.d. and μ_E is not, we get

$$h(\mu') = \mu'(E)h(\mu'_E) > \mu(E)h(\mu_E) = h(\mu)$$

by [9, Chapter 10, §6, Theorem 2].

It follows that μ is determined by the joint distribution of renewal times (distance between C^{max} rungs) and the inter-renewal configuration. Suppose that a block has positive probability of being non-left-burnable and also positive probability of being non-right-burnable. Then with probability one, there will be a non-left-burnable block to the left of a non-right-burnable block. This creates a forbidden subconfiguration, and hence is impossible. Therefore, at most one of the above possibilities has positive probability. Assume without loss of generality that blocks are left-burnable with probability 1.

Consider now the configuration between two renewals (not necessarily consecutive) that are distance L apart. By maximum entropy, the conditional distribution of the configuration given L is uniform over all left-burnable configurations of length $L - 1$. Since this holds for arbitrarily large L, it implies that the finite-dimensional distributions of μ are given by the thermodynamic limit of $\mu^{\mathrm{L}}_{n,m}$, and hence $\mu = \mu^{\mathrm{L}}$. Analogously, we get $\mu = \mu^{\mathrm{R}}$ if blocks are right-burnable with probability 1.

The proof in the symmetric case is very similar. Adding C^{max} rungs in an i.i.d. fashion does not destroy the symmetry of the measure, and hence C^{max} rungs have to occur with positive probability. As before, they are renewals. Again, blocks have to be either left- or right-burnable, and by symmetry, they

have to be both with probability 1. As before, this implies that the measure coincides with μ^{S}.   □

Theorem 3.2. *If $-n, m \to \infty$ in such a way that $\lim -n/m = \rho/(1-\rho)$, $\rho \in [0,1]$, then*

$$\lim \mu_{n,m} = \rho \mu^{\mathrm{L}} + (1-\rho)\mu^{\mathrm{R}}.$$

Consequently, the set of weak limit points of $\{\mu_{n,m}\}$ consists of all convex combinations of μ^{L} and μ^{R}.

Proof. The idea of the proof is the following. A recurrent configuration has to burn if we burn from both the left and the right. We show that the left- and right-burnable portions of the configuration almost form a partition of $\Lambda_{n,m}$, up to an overlap or uncovered region of size $o(m-n)$ in probability, and that the location of the "boundary layer" between them is approximately uniform over $[n, m]$. This implies that in a fixed finite window we see a convex combination of μ^{L} and μ^{R}.

For $\eta \in \Omega_{n,m}$, let

$$\sigma^{\mathrm{L}} := \sigma^{\mathrm{L}}_{n,m} := \max\{k : C_k = C^{\mathrm{max}} \text{ and } \eta_{\Lambda_{n,k}} \text{ is left-burnable}\},$$

$$\sigma^{\mathrm{R}} := \sigma^{\mathrm{R}}_{n,m} := \min\{k : C_k = C^{\mathrm{max}} \text{ and } \eta_{\Lambda_{k,m}} \text{ is right-burnable}\},$$

where the values $n-1$ and $m+1$ are allowed in both cases if there is no k with the required property.

We show that

$$\frac{|\sigma^{\mathrm{L}} - \sigma^{\mathrm{R}}|}{m-n} \to 0 \tag{3.7}$$

in probability.

<u>Case 1: $\sigma^{\mathrm{L}} < \sigma^{\mathrm{R}}$</u>. We show that the number of possible configurations between $s := \sigma^{\mathrm{L}}$ and $t := \sigma^{\mathrm{R}}$ has exponential growth rate smaller than h^{L}.

Let $\eta^0 := \eta_{\Lambda_{s+1,t-1}}$. Consider left-burning on η^0. By the definition of σ^{L}, and since $C_t = C^{\mathrm{max}}$, the rightmost site of η^0 that will be left-burnt (when left-burning η^0) is in a rung k with $s \le k < t-1$. Here $k = s$ if left-burning cannot start. Similarly, the leftmost site of η^0 that can be right-burnt is in a rung l with $s+1 < l \le t$. Since η^0 is burnable, we need to have $l \le k+1$.

Let ξ^0 be the configuration obtained by replacing rung $k+1$ of η^0 by C^{max}. Since ξ^0 is also burnable, it follows easily that $\eta^0_{\Lambda_{s+1,k}}$ is left-burnable, and $\eta^0_{\Lambda_{k+2,t-1}}$ is right-burnable. Therefore, the number of possibilities for η^0 is bounded by

$$\sum_{k=s}^{t-2} b_{k-s} |\mathcal{C}| b_{t-k-2} \le C e^{(h^{\mathrm{L}}-\delta)(t-s)}$$

for some $\delta > 0$ and some $C < \infty$ by Lemma 3.4.

Summing over all possible values of s and t, it follows that for any $\varepsilon > 0$ there exist $C_1 = C_1(\varepsilon)$ and $c_1 = c_1(\varepsilon) > 0$ such that

$$\mu_{n,m}\{\sigma^{\mathrm{R}} - \sigma^{\mathrm{L}} \geq \varepsilon(m-n)\} \leq C_1 e^{-c_1(m-n)}. \tag{3.8}$$

Case 2: $\sigma^{\mathrm{L}} \geq \sigma^{\mathrm{R}}$. Observe that the configuration between σ^{R} and σ^{L} is both left-burnable and right-burnable, hence it belongs to $\Omega^{\mathrm{S}}_{\sigma^{\mathrm{R}},\sigma^{\mathrm{L}}}$. Also, the configuration to the left of σ^{R} is in $\Omega^{\mathrm{L}}_{n,\sigma^{\mathrm{R}}}$, and the configuration to the right of σ^{L} is in $\Omega^{\mathrm{R}}_{\sigma^{\mathrm{L}},m}$.

Since $h^{\mathrm{S}} < h^{\mathrm{L}}$, it follows that for any $\varepsilon > 0$ there exists C_2 and $c_2 = c_2(\varepsilon) > 0$ such that

$$\mu_{n,m}\{\sigma^{\mathrm{L}} - \sigma^{\mathrm{R}} \geq \varepsilon(m-n)\} \leq C_2 e^{-c_2(m-n)}. \tag{3.9}$$

The bounds (3.8) and (3.9) establish (3.7).

In the remainder of the proof we are going to need a minor variation on (3.9) when $\sigma^{\mathrm{R}} \leq \sigma^{\mathrm{L}}$. The reason is that the value of σ^{R} gives some information on the left-burnable configuration to the left of σ^{R} (namely, that it is not right-burnable if it contains a rung C^{max}), whereas we would like to achieve independence. Let $\hat{\sigma}^{\mathrm{R}}$ denote the rightmost C^{max} rung to the left of σ^{R} (we set $\hat{\sigma}^{\mathrm{R}} = n - 1$ if such a rung does not exist). Then the configuration between $\hat{\sigma}^{\mathrm{R}}$ and σ^{R} is left-burnable but not right-burnable. In any case, it is in $\Omega^{\mathrm{L},0}_{\hat{\sigma}^{\mathrm{R}}+1,\sigma^{\mathrm{R}}-1}$. We define $\hat{\sigma}^{\mathrm{L}}$ analogously. By similar arguments as before, we have the bound

$$\mu_{n,m}\{\sigma^{\mathrm{L}} \geq \sigma^{\mathrm{R}} \text{ and } \hat{\sigma}^{\mathrm{L}} - \hat{\sigma}^{\mathrm{R}} \geq \varepsilon(m-n)\} \leq C_3 \exp\{-c_3(m-n)\}. \tag{3.10}$$

Next we prove that the location of the "boundary layer" between the left- and right-burnable parts is approximately uniform.

First condition on the value of $d := \sigma^{\mathrm{R}} - \sigma^{\mathrm{L}}$ in the case when d is positive. Observe that given $\sigma^{\mathrm{L}} = s$ and $\sigma^{\mathrm{R}} = t$, the configurations on $\Lambda_{n,s}$, $\Lambda_{s,t}$ and $\Lambda_{t,m}$ are conditionally independent. Also, the configuration on $\Lambda_{n,s-1}$ has law $\mu^{\mathrm{L}}_{n,s-1}$ and the configuration on $\Lambda_{t+1,m}$ has law $\mu^{\mathrm{R}}_{t+1,m}$. Noting that μ^{R} is the reflection of μ^{L}, we can uniquely represent the configuration in the following way. Draw a sample η from $\mu^{\mathrm{L}}_{n,m-d}$ conditioned on having at least one renewal. Select one of the C^{max} rungs uniformly at random: suppose it is rung S. Draw an independent sample ξ from the set of configurations η^0 described under Case 1 above having length $d - 1$. Concatenate the configurations $\eta_{\Lambda_{n,S}}$, ξ, C^{max}, and the reversal of $\eta_{\Lambda_{S+1,m-d}}$. This gives all configurations with $\sigma^{\mathrm{L}} = S$ and $\sigma^{\mathrm{R}} = S + d$, and the representation is unique.

Next we want to show that the random variable S defined above is roughly uniformly distributed in $[n, m-d]$. First note that by Lemma 3.3, under $\mu^{\mathrm{L}}_{n,m-d}$, the distribution of the sequence of inter-renewal times is exchangeable. Also, due to the inequality $h^{\mathrm{L},0} < h^{\mathrm{L}}$, the longest inter-renewal time is $o(m - n - d)$ in probability. These two together imply that for any $0 < u < 1$,

$$\mu^{\mathrm{L}}_{n,m-d}\{S - n < u(m - n - d)\} \to u \quad \text{as } m - n \to \infty$$

uniformly in $1 \le d < (m - n)/2$. This implies that

$$\mu_{n,m}\{\sigma^L - n < u(m - n - d) \mid \sigma^R - \sigma^L = d\} \to u \quad \text{as } m - n \to \infty$$

uniformly in $1 \le d < (m - n)/2$. Averaging over $1 \le d \le \varepsilon(m - n)$, we get

$$\mu_{n,m}\left\{ \frac{\sigma^L - n}{m - n} < u \mid 1 \le \sigma^R - \sigma^L \le \varepsilon(m - n) \right\} = u + O(\varepsilon) + o(1) \qquad (3.11)$$

as $m - n \to \infty$.

Now condition on $d := \hat{\sigma}^L - \hat{\sigma}^R$ in the case when $\sigma^L \ge \sigma^R$. Given $\hat{\sigma}^R = \hat{s}^R$ and $\hat{\sigma}^L = \hat{s}^L$, the configurations on $\Lambda_{n,\hat{s}^R-1}$, $\Lambda_{\hat{s}^R,\hat{s}^L}$ and $\Lambda_{\hat{s}^L+1,m}$ are conditionally independent, with the first and the third having laws $\mu^L_{n,\hat{s}^R-1}$ and $\mu^R_{\hat{s}^L+1,m}$ (respectively). Therefore, the configuration can be represented analogously to the case $\sigma^L < \sigma^R$, which gives rise to the estimate

$$\mu_{n,m}\left\{ \frac{\hat{\sigma}^R - n}{m - n} < u \mid \hat{\sigma}^L - \hat{\sigma}^R \le \varepsilon(m - n) \right\} = u + O(\varepsilon) + o(1) \qquad (3.12)$$

as $m - n \to \infty$.

We are ready to complete the proof of the theorem. Suppose we have a cylinder event E depending on the configuration in $\Lambda_{-k,k}$. Let

$$\tau^L := \begin{cases} \sigma^L & \text{if } \sigma^L < \sigma^R, \\ \hat{\sigma}^R & \text{if } \sigma^L \ge \sigma^R, \end{cases} \quad \text{and} \quad \tau^R := \begin{cases} \sigma^R & \text{if } \sigma^L < \sigma^R, \\ \hat{\sigma}^L & \text{if } \sigma^L \ge \sigma^R. \end{cases}$$

Let $A_\varepsilon := \{\tau^L > \varepsilon(m-n)\}$ and $B_\varepsilon := \{\tau^R < -\varepsilon(m-n)\}$. For $t > \varepsilon(m-n) > k$,

$$\mu_{n,m}\{E \mid A_\varepsilon, \tau^L = t\} = \mu^L_{n,t-1}\{E\} = \mu^L\{E\}\big(1 + o_\varepsilon(1)\big)$$

as $n \to -\infty$ and $m \to \infty$, where the $o_\varepsilon(1)$ depends on ε, but not on t. Similarly, for $t < -\varepsilon(m-n) < -k$,

$$\mu_{n,m}\{E \mid B_\varepsilon, \tau^R = t\} = \mu^R_{t+1,m}\{E\} = \mu^R\{E\}\big(1 + o_\varepsilon(1)\big).$$

Since by (3.7), (3.10), (3.11) and (3.12), $\mu_{n,m}\{A_\varepsilon\} = -n/(m - n) + O(\varepsilon)$ and $\mu_{n,m}\{B_\varepsilon\} = m/(m - n) + O(\varepsilon)$, the theorem follows by letting $\varepsilon \to 0$. $\qquad \square$

4. Avalanches

By toppling in an infinite graph we mean the following. Suppose we start from a configuration η with finitely many unstable sites. We simultaneously topple all unstable sites, and repeat this as long as there are unstable sites (possibly infinitely many times). After each step there are only finitely many unstable sites. This is equivalent to toppling sites one-by-one during each step, before moving on to toppling other sites. Let us call this the *standard toppling*.

Definition 4.1. A (possibly infinite) sequence of topplings is called *legal* if it has the properties: (i) only unstable sites are toppled in each step; (ii) any site that is unstable at some step will be toppled at some later step.

Lemma 4.1. *Any two legal sequences of topplings are equivalent in the sense that each site topples the same number of times in both sequences (which may be infinity). In particular, any legal sequence of topplings is equivalent to standard toppling.*

Proof. This can be proved the same way as for finite sequences of topplings [28]. Given two legal sequences of topplings at the sites

$$x_1, x_2, \ldots$$
$$y_1, y_2, \ldots$$

we can transform one into the other. Since x_1 is unstable at the beginning, it has to occur in the second sequence. Suppose it occurs first as y_{k_1}. Then the toppling of y_{k_1} can be commuted through the topplings of $y_1, y_2, \ldots, y_{k_1-1}$, so the y-sequence is equivalent to

$$x_1, y_2, \ldots, y_{k_1-1}, y_{k_1+1}, y_{k_1+2}, \ldots$$

We can now eliminate x_1 from both sequences, and the lemma follows.   □

Theorem 4.1. *Suppose we add one grain to each site in rung 0 in an infinite left-burnable configuration. Then each site will topple infinitely many times. The same holds for right-burnable configurations.*

Proof. Add one grain to each site in rung 0, and initially, do not topple in rungs to the left of zero. (For the moment, let us disregard that this may be an illegal sequence of topplings.) The topplings that occur on the right are equivalent to the burning procedure on $\Lambda_{0,\infty}$. Since the configuration is left-burnable, each site in $\Lambda_{0,\infty}$ will topple exactly once. In particular, each site in rung 0 will have toppled. Also, it is easy to verify that each site in $\Lambda_{0,\infty}$ will have received as many grains as it has lost, and hence has its original height.

The topplings in rung 0 give one grain to each site in rung -1. Therefore, the argument can be repeated as if we have added one grain to each site in rung -1, and hence topplings continue forever. This almost completes the argument, apart from the technicality that this is not a legal sequence of topplings. Instead, now we carry out the topplings on the right only to a large finite time until a rung $p_1 \gg 1$ is toppled. Then carry out topplings started from rung -1, until a rung $1 < p_2 < p_1$ is toppled, and so on. If $p_1 \geq 2K$, we can repeat this with rungs $1 < K < p_K < \cdots < p_1$, for any given large K. At this point, rung $-k$ has toppled $K - k$ times for $k = 1, \ldots, K$, and rungs $0, 1, \ldots, K$ have toppled K times. It follows that in any legal sequence of topplings, sites $-K/2, \ldots, K/2$ each topple at least $K/2$ times. Since K was arbitrary, the theorem follows.   □

Remark 4.1. As the following example shows, there can be infinite avalanches such that every site topples only finitely many times. Take $G = I_{0,1}$. Under μ^L, there is positive probability that the configuration at rungs 1–6 equals

$$
\begin{array}{cccccc}
3 & 2 & 3 & 3 & 1 & 3 \\
3 & 3 & 1 & 3 & 3 & 3
\end{array}
$$

Now adding a grain to the first row in rung 4 yields an avalanche with toppling numbers:

$$
\begin{array}{cccccccc}
\dots & 0 & 0 & 1 & 2 & 1 & 1 & \dots \\
\dots & 0 & 0 & 0 & 1 & 1 & 1 & \dots
\end{array}
$$

5. Coding by Markov chains

In this section we show that the measures $\mu^L_{n,m}$ and μ^L can be coded by a Markov chain with finitely many states. Before proving this for a general graph G, we sketch a proof in the special case $G = I_{0,1}$. Although for general G we will not have as explicit a description as for $I_{0,1}$, the approach will be similar.

Coding by a finite Markov chain for $G = I_{0,1}$. Based on Lemma 3.2, the following equivalent description of left-burnable configurations can be given. Consider the alphabet of symbols

$$
\mathcal{A} := \{(3,3), (3,2), (2,3), (3,1), (1,3), \overline{(3,2)}, \overline{(2,3)}\}. \tag{5.1}
$$

Let $\mathcal{A}_{n,m} := \mathcal{A}^{I_{n,m}}$. We think of $\overline{(3,2)}$ replacing a $(3,2)$ rung that is following a $(3,1)$ before the next $(3,3)$ occurs. It follows from the characterization in Lemma 3.2 that elements of $\Omega^L_{n,m}$ can be coded in a one-to-one fashion by a set $\overline{\Omega}^L_{n,m} \subset \mathcal{A}_{n,m}$ that is a topological Markov chain (subshift of finite type) [21, Section 1.9] with alphabet $\mathcal{A}$. Namely, the only restrictions on sequences in $\overline{\Omega}^L_{n,m}$ are that certain pairs of symbols cannot occur next to each other. For example: (a) $(3,3)$ has to be followed by $(3,3), (3,2), (2,3), (3,1)$ or $(1,3)$; (b) $(3,1)$ has to be followed by $(3,3)$ or $\overline{(3,2)}$; (c) $\overline{(3,2)}$ has to be followed by $\overline{(3,2)}$ or $(3,3)$; etc. The full transition matrix is

$$
T := \begin{pmatrix}
1 & 1 & 1 & 1 & 1 & 0 & 0 \\
1 & 1 & 1 & 1 & 1 & 0 & 0 \\
1 & 1 & 1 & 1 & 1 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 1 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 1 \\
1 & 0 & 0 & 0 & 0 & 1 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 1
\end{pmatrix}, \tag{5.2}
$$

where the rows and columns correspond to the symbols in the order displayed in (5.1). Due to the special role of the symbols $\overline{(3,2)}$ and $\overline{(2,3)}$, we need to add the boundary condition that the rung at n is not one of these.

It is not hard to check that the topological Markov chain is transitive [21, Definition 1.9.6]; in fact, all entries of T^3 are positive. Let $\bar{\Omega}^{\mathrm{L}} \subset \mathcal{A}^{\mathbb{Z}}$ denote the subshift defined by T, and let $\bar{\mu}^{\mathrm{L}}$ be its Parry measure, which is a Markov chain. By [21, Section 4.4], $\bar{\mu}^{\mathrm{L}}$ is the unique measure of maximum entropy on $\bar{\Omega}^{\mathrm{L}}$. Let $\mathcal{P} : \bar{\Omega}^{\mathrm{L}} \to \Omega^{\mathrm{L}}$ denote the map that replaces each $\overline{(3,2)}$ by $(3,2)$ and each $\overline{(2,3)}$ by $(2,3)$. Since μ^{L} has maximal entropy by Theorem 3.1, and $\mathcal{P}^{-1}$ is well defined μ^{L}-almost surely, $\mathcal{P}$ is a metric isomorphism between $\bar{\mu}^{\mathrm{L}}$ and μ^{L}.

Now we generalize the coding to an arbitrary graph G. First note that it is not very surprising that such a coding should exist. Using Majumdar and Dhar's tree construction [27], recurrent configurations in $\Lambda_{n,m}$ are in one-to-one correspondence with spanning trees of $\Lambda_{n,m}$ with wired boundary conditions. It has been shown in [16] that spanning trees have a Markovian coding. However, since the correspondence is non-local, it does not seem easy to deduce a Markovian coding from the spanning-tree result.

We let $\mathcal{P} := \mathcal{P}(G)$ denote the set of all subsets of G.

Theorem 5.1. *There exists an alphabet $\mathcal{A} := \mathcal{A}(G) \subset \mathcal{C} \times \mathcal{P} \times \mathcal{P}^{\mathcal{P}}$, an inclusion $i : \mathcal{C} \to \mathcal{C} \times \mathcal{P} \times \mathcal{P}^{\mathcal{P}}$, and a transitive 0-1 matrix $T := T(G)$ indexed by $\mathcal{A}$ such that for each m, the set $\Omega_{1,m}^{\mathrm{L}}$ is in one-to-one correspondence with the set of sequences*

$$
\begin{aligned}
\overline{\Omega}_{1,m}^{\mathrm{L}} &:= \overline{\Omega}_{1,m}^{\mathrm{L}}(G) \\
&:= \{\omega \in \mathcal{A}_{1,m} : \omega_1 \in i(\mathcal{C}),\, T(\omega_k, \omega_{k+1}) = 1,\, k = 1, \ldots, m-1\}.
\end{aligned}
$$

The correspondence is given by the projection $P : \mathcal{C} \times \mathcal{P} \times \mathcal{P}^{\mathcal{P}} \to \mathcal{C}$ applied coordinatewise.

For the proof of Theorem 5.1, we will need to perform left-burning in a special way, as introduced below. This can be regarded as a generalization of the rung-by-rung argument from the proof of Lemma 3.2. Following the definition of the special burning rule, we use it to prove two lemmas that will lead to the proof of Theorem 5.1. Once Theorem 5.1 is established, the Markov chain that codes μ^{L} is the Parry measure, as for $G = I_{0,1}$.

Burning with leftmost rung rule. We perform burning one rung at a time, with the rule that whenever there are no more burnable sites in the rung currently being burnt, we move on to the leftmost rung that has burnable sites. We now describe the procedure in more detail.

We first burn sites in rung 1 that can be burnt consistent with the left-burning rule. When there are no more burnable sites in rung 1, we start burning sites in rung 2, and continue burning rung 2 until there are no more burnable sites in that rung. This may have created further burnable sites in rung 1. If there are such, we burn sites in rung 1, again until there are no more burnable

sites in that rung. At some point there will be no burnable sites in either rung 1 or 2. Now we burn sites in rung 3, and move between rungs 1, 2 and 3 until there are no more burnable sites in those rungs. In general, we move on to rung $k + 1$ when there are no more burnable sites in rungs $1, \ldots, k$.

If the configuration we started with is $C_1 \vee \cdots \vee C_m$, we adopt the following convention for burning the rightmost rung C_m. We add a "ghost" rung $C_{m+1} = C^{\max}$ that will remain unburnt until the first time there are no more burnable sites in rungs $1, \ldots, m$. At this time, we burn the ghost rung, and continue with the leftmost rule. It is easy to see that this yields an equivalent definition of left-burnability, that is, all rungs will burn if and only if the original configuration was left-burnable.

For $1 \leq k \leq m$ and $C_1 \vee \cdots \vee C_k$ left-burnable, let $T_k + 1$ be the first time we burn a site in rung $k + 1$. It is easy to see that all rungs are burnt at time T_{m+1} if and only if $\eta \in \Omega_{1,m}^{\mathrm{L}}$.

Before stating the two lemmas needed for Theorem 5.1, we need some notation. Let $\eta = C_1 \vee \cdots \vee C_m$ be a configuration with $C_k \in \mathcal{C}$, $1 \leq k \leq m$. Let $B_k \subset G$ denote the set of sites in rung k that have been burnt by time T_k. The sequence $(C_k, B_k)_{k=1}^m$ is non-Markovian in general. We note, however, in order to motivate the arguments to come, that if $G = I_{0,1}$, then B_{k+1} is a function of C_k, C_{k+1} and B_k only (it depends on $C_1, \ldots, C_{k-1}$ only through C_k, C_{k+1} and B_k). It is not hard to show that this implies that $(C_k, B_k)_{k=1}^m$ is Markovian. The proof is similar (and simpler) than that of Theorem 5.1 below, and is left to the reader. For general G, our strategy will be to augment the information contained in B_k so that we get a Markovian sequence.

Fix $(C_j, B_j)_{j=1}^k$, where $1 \leq k \leq m$. Depending on this sequence, we define a function $f_k : \mathcal{P} \to \mathcal{P}$ that will encode what the effect is of burning in rung $k + 1$ on the future of the burning process in rungs $1 \leq j \leq k$. We stress that the definition of f_k will ignore the actual value of C_{k+1}; in particular, it will also make sense for $k = m$. Fix $A \subset G$. Regardless of the value of C_{k+1}, let us declare all sites in $A \times \{k + 1\}$ to be burnt. This may create burnable sites in rung k after $\bigcup_{j=1}^k B_j$ has been burnt. Now let us perform burning with the leftmost rung rule until there are no more burnable sites in rungs $1 \leq j \leq k$. This process does not use information about rung $k + 1$ other than the specified set A. We define $f_k(A)$ to be the set of sites that are burnt in rung k at the end of this process. For example, we have $f_k(\emptyset) = B_k$, and more generally, $f_k(A) = B_k$ for $A \subset B_k$, since in this case no new burnable sites appear in rung k. Whenever $C_1 \vee \cdots \vee C_k$ is left-burnable, we have $f_k(G) = G$. In general, we have $B_k \subset f_k(A) \subset G$ for $A \in \mathcal{P}$.

We prove Theorem 5.1 by showing that $(C_k, B_k, f_k)_{k=1}^m$ is Markovian. We verify this in the two lemmas below that characterize the pairs that can occur next to each other for left-burnable $\eta = C_1 \vee \cdots \vee C_m$. To facilitate the proof, we define an auxiliary function $g : \mathcal{P} \times \mathcal{C} \times \mathcal{P} \to \mathcal{P}$. Given $A, A' \subset G$ and $C \in \mathcal{C}$, we

set the configuration in rung 1 to be C and declare all sites in $A \times \{0\} \cup A' \times \{2\}$ to be burnt. Now we perform left-burning in rung 1. By this we mean specifying a maximal sequence of vertices $v_1, \ldots, v_k \in G \times \{1\}$, such that the requirements of Definition 3.1 are satisfied with $V := (G \setminus A) \times \{0\} \cup G \times \{1\} \cup (G \setminus A') \times \{2\}$. We define $g(A, C, A')$ to be the set of sites that burn in rung 1.

Lemma 5.1. *For $\eta = C_1 \vee \cdots \vee C_m \in \Omega_{1,m}^{\mathrm{L}}$, the following properties hold:*

(a) $(B_1, f_1) = \psi(C_1)$ *for some function $\psi = \psi_G$, in fact, $B_1 = g(G, C_1, \emptyset)$ and $f_1(A) = g(G, C_1, A)$;*

(b) $g(B_k, C_{k+1}, \emptyset) \neq \emptyset$, $1 \le k < m$;

(c) $(B_{k+1}, f_{k+1}) = \phi(B_k, C_{k+1}, f_k)$ *for a function $\phi = \phi_G$ independent of k, $1 \le k < m$; and*

(d) $f_k(G) = G$, $1 \le k \le m$.

Proof. (a) follows directly from the definitions, and (d) has been observed before the statement of the lemma. If (b) failed for some $1 \le k < m$, that would mean that after time T_k there were no burnable sites in rungs $1, \ldots, k+1$, with rung $k+1$ completely unburnt. That means that there are no burnable sites at all after time T_k, which contradicts the burnability of η.

The proof of (c) is a bit lengthy. We first show that B_{k+1} is a function of B_k, C_{k+1} and f_k. For this, we look at the burning process between times $T_k + 1$ and T_{k+1} in more detail. We define the following intermediate times: letting $R_0 := T_k + 1$, we define

$$S_1 := \min \left\{ R_0 \le n \le T_{k+1} : \begin{array}{l} \text{there are no burnable sites in} \\ \text{rung } k+1 \text{ at time } n \end{array} \right\};$$

$$R_1 := \min \left\{ S_1 \le n \le T_{k+1} : \begin{array}{l} \text{there are no burnable sites in rungs} \\ 1 \le j \le k \text{ at time } n \end{array} \right\},$$

and for $i \ge 2$ we recursively set

$$S_i := \min \left\{ R_{i-1} \le n \le T_{k+1} : \begin{array}{l} \text{there are no burnable sites in} \\ \text{rung } k+1 \text{ at time } n \end{array} \right\};$$

$$R_i := \min \left\{ S_i \le n \le T_{k+1} : \begin{array}{l} \text{there are no burnable sites in rungs} \\ 1 \le j \le k \text{ at time } n \end{array} \right\}.$$

Set $B^{(0)} := B_k$. Between times $T_k + 1 = R_0$ and S_1, the subset $A^{(0)} := g(B^{(0)}, C_{k+1}, \emptyset)$ of rung $k+1$ is burnt. Let $B^{(1)} := f_k(A^{(0)}) \supset B^{(0)}$. By the definition of f_k, $B^{(1)}$ is the set of sites in rung k that is burnt by time R_1. We set $A^{(1)} := g(B^{(1)}, C_{k+1}, \emptyset)$ and $B^{(2)} := f_k(A^{(1)})$. By the definition of g, $A^{(1)}$ is the set of sites in rung $k+1$ burnt at time S_2. Although less obvious, $B^{(2)}$ is

the set of sites in rung k burnt at time R_2. The latter statement needs careful proof since f_k was defined in terms of the state of the burning process at time R_0 rather than at R_1. Consider the sequence of sites burnt in the computation of $f_k(A^{(1)})$ (following the definition). We merely get a rearrangement of this sequence if we first declare $A^{(0)} \times \{k+1\}$ to be burnt, let burning act on rungs $1, \ldots, k$, then declare $\big(A^{(1)} \setminus A^{(0)}\big) \times \{k+1\}$ to be burnt, and then let burning act on rungs $1, \ldots, k$. This observation proves our claim about $B^{(2)}$.

In general, for $i \geq 1$, after burning between times R_i and S_{i+1}, the set of sites burnt in rung $k+1$ is $A^{(i)} := g(B^{(i)}, C_{k+1}, \emptyset) \supset A^{(i-1)}$. We set $B^{(i+1)} := f_k(A^{(i)})$. Similarly to the case $i = 1$ spelled out above, by a decomposition of $A^{(i)}$, we get that $B^{(i+1)}$ is the set of sites burnt in rung k at time R_{i+1}. Since there is some j_0 for which $B^{(j+1)} = B^{(j)}$ and $A^{(j+1)} = A^{(j)}$ for $j \geq j_0$, we have $B_{k+1} = A^{(j_0)}$. To summarize, B_{k+1} is obtained as the stable result of applying the functions g and f_k according to

$$
\begin{aligned}
B^{(0)} &:= B_k, & A^{(0)} &:= g(B^{(0)}, C_{k+1}, \emptyset), \\
B^{(1)} &:= f_k(A^{(0)}), & A^{(1)} &:= g(B^{(1)}, C_{k+1}, \emptyset), \\
&\;\;\vdots & &\;\;\vdots \\
B^{(j_0)} &:= f_k(A^{(j_0-1)}), & B_{k+1} = A^{(j_0)} &:= g(B^{(j_0)}, C_{k+1}, \emptyset).
\end{aligned}
\tag{5.3}
$$

This shows that B_{k+1} is a function of B_k, C_{k+1} and f_k that does not depend on the value of k (is "k-independent").

Now we can prove the remainder of (c) by a similar argument. Consider the state of the burning process at time T_{k+1}, at which time the set of sites burnt in rung $k+1$ is B_{k+1}. By the definition of f_k, the set of sites in rung k burnt at time T_{k+1} is the set $\bar{B}^{(0)} := f_k(B_{k+1})$. Given $A \subset G$, declare all sites in $A \times \{k+2\}$ to be burnt (ignoring C_{k+2}). Now perform burning in rung $k+1$, which ends at some time $\bar{S}_1$. Then the set of sites in rung $k+1$ burnt at time $\bar{S}_1$ is

$$
\bar{A}^{(0)} := g(\bar{B}^{(0)}, C_{k+1}, A) \supset B_{k+1}.
$$

Now we can essentially apply the argument above starting with $\bar{A}^{(0)}$ in place of $A^{(0)}$. We perform burning on rungs $1, \ldots, k$, and let $\bar{R}_1 \geq \bar{S}_1$ be the first time when there are no burnable sites in these rungs. Then the set of sites in rung k burnt at time $\bar{R}_1$ is $\bar{B}^{(1)} := f_k(\bar{A}^{(0)})$. This is shown by the decomposition

$$
\bar{A}^{(0)} = A^{(0)} \cup \left[\bigcup_{j=1}^{\infty} \big(A^{(j)} \setminus A^{(j-1)}\big) \right] \cup \big(\bar{A}^{(0)} \setminus B_{k+1}\big).
$$

Next we perform burning in rung $k+1$ that stops at some time $\bar{S}_2 \geq \bar{R}_1$, and then on rungs $1, \ldots, k$, which stops at $\bar{R}_2 \geq \bar{S}_2$. The set of sites in rung $k+1$ burnt at time $\bar{S}_2$ is $\bar{A}^{(1)} := g(\bar{B}^{(1)}, C_{k+1}, A)$, and the set of sites in rung k burnt

at time $\bar{R}_2$ is $f_k(\bar{A}^{(1)})$. We continue to iterate g and f_k until the burnt sites in rung $k+1$ stabilize to some set $\bar{A}^{(\bar{j}_0)}$. We then have $f_{k+1}(A) = \bar{A}^{(\bar{j}_0)}$. We have

$$
\begin{aligned}
\bar{B}^{(0)} &:= f_k(B_{k+1})\,, & \bar{A}^{(0)} &:= g(\bar{B}^{(0)}, C_{k+1}, A)\,, \\
\bar{B}^{(1)} &:= f_k(\bar{A}^{(0)})\,, & \bar{A}^{(1)} &:= g(\bar{B}^{(1)}, C_{k+1}, A)\,, \\
&\ \ \vdots & &\ \ \vdots \\
\bar{B}^{(\bar{j}_0)} &:= f_k(\bar{A}^{(\bar{j}_0-1)})\,, & f_{k+1}(A) = \bar{A}^{(\bar{j}_0)} &:= g(\bar{B}^{(\bar{j}_0)}, C_{k+1}, A)\,.
\end{aligned}
\tag{5.4}
$$

This shows that f_{k+1} is a k-independent function of B_k, C_{k+1} and f_k, and hence (c) follows. $\square$

Lemma 5.2. *Let ψ and ϕ be as in Lemma 5.1. Suppose that the sequence $(C_k', B_k', f_k') \in \mathcal{C} \times \mathcal{P} \times \mathcal{P}^{\mathcal{P}}$, $k = 1, \dots, m$, satisfies the conditions:*

(A) $(B_1', f_1') = \psi(C_1')$;

(B) $g(B_k', C_{k+1}', \emptyset) \neq \emptyset$, $1 \leq k < m$;

(C) $(B_{k+1}', f_{k+1}') = \phi(B_k', C_{k+1}', f_k')$, $1 \leq k < m$; and

(D) $f_k'(G) = G$, $1 \leq k \leq m$.

Then $\eta := C_1' \vee \cdots \vee C_m' \in \Omega_{1,m}^{\mathrm{L}}$, and taking $C_k := C_k'$ in the definitions preceding Lemma 5.1, we have $B_k = B_k'$ and $f_k = f_k'$, $1 \leq k \leq m$.

Proof. We verify the statement by induction on m. When $m = 1$, $\eta = C_1' \in \Omega_{1,1}^{\mathrm{L}}$ since $C_1' \in \mathcal{C}$. Therefore by (A), $(B_1, f_1) = \psi(C_1) = \psi(C_1') = (B_1', f_1')$.

Now assume the statement of the lemma holds for some $m \geq 1$, and we prove it for $m + 1$. Hence assume that (A)–(D) hold with m replaced by $m + 1$. By the induction hypothesis, $\eta_m := C_1' \vee \cdots \vee C_m' \in \Omega_{1,m}^{\mathrm{L}}$. Since the definitions of B_k, f_k $(1 \leq k \leq m)$ do not depend on C_{m+1}', we also get $B_k = B_k'$ and $f_k = f_k'$ for $1 \leq k \leq m$. Note also that T_k, $1 \leq k \leq m$, has the same value whether we consider the burning of η_m or η_{m+1}. Since $B_m = B_m'$, we have

$$
B_{m+1} \supset g(B_m, C_{m+1}', \emptyset) = g(B_m', C_{m+1}', \emptyset) \neq \emptyset,
$$

by (B). Therefore, B_{m+1} is not empty. We show that this implies that η_{m+1} is left-burnable.

First, note that rung $m + 2$ (the ghost rung) can be burnt. Our assumption (C) says that $f_{m+1}'(G)$ is determined via the function ϕ by the data: $f_m' = f_m$, $B_m' = B_m$ and $C_{m+1}' = C_{m+1}$, and that its value can be obtained as the result of the computation in (5.4). After the ghost rung has been burnt, the burning of rungs m and $m + 1$ will follow the pattern of (5.4), with $A := G$. Since the computation will stabilize with result $f_{m+1}'(G) = G$, this means that eventually

everything in rung $m+1$ burns. By left-burnability of η_m, this means that also all the rungs $1, \ldots, m$ burn, and hence η_{m+1} is left-burnable.

By Lemma 5.1 (c) and (C), we now have

$$(B_{m+1}, f_{m+1}) = \phi(B_m, C_{m+1}, f_m) = \phi(B'_m, C'_{m+1}, f'_m) = (B'_{m+1}, f'_{m+1}).$$

This advances the induction, and the lemma follows.   □

Proof of Theorem 5.1. Let ψ and ϕ be as in Lemma 5.1. We define the inclusion $i(C) = \big(C, \psi(C)\big)$. We define the alphabet $\mathcal{A}$ as the set of (C, B, f) such that there exists $m \geq 1$ and a sequence $(C_k, B_k, f_k)_{k=1}^m$ with

 (i) $(C_m, B_m, f_m) = (C, B, f)$;

 (ii) $(C_1, B_1, f_1) = i(C_1)$;

 (iii) $g(B_k, C_{k+1}, \emptyset) \neq \emptyset$, $1 \leq k < m$;

 (iv) $(B_{k+1}, f_{k+1}) = \phi(B_k, C_{k+1}, f_k)$, $1 \leq k < m$; and

 (v) $f_k(G) = G$, $1 \leq k \leq m$.

We define the transition matrix T by

$$T\big((C, B, f), (C', B', f')\big) = \begin{cases} 1 & \text{if } g(B, C', \emptyset) \neq \emptyset,\ (B', f') = \phi(B, C', f), \\ 0 & \text{otherwise.} \end{cases}$$

It follows from these definitions that for any left-burnable η,

$$\omega := (\omega_k)_{k=1}^m := (C_k, B_k, f_k)_{k=1}^m \in \mathcal{A}_{1,m}.$$

By Lemma 5.1 we have in fact $\omega \in \overline{\Omega}_{1,m}^{\mathrm{L}}$. It follows from Lemma 5.2 that every element of $\overline{\Omega}_{1,m}^{\mathrm{L}}$ arises this way. The correspondence $\eta \mapsto \omega$ satisfies $P(\omega) = \eta$ by definition, and hence is one-to-one.

It remains to show that T is transitive. For this, we first show that for any $(C, B, f) \in \mathcal{A}$ we have $T\big((C, B, f), i(C^{\mathrm{max}})\big) = 1$. It is easy to verify that $i(C^{\mathrm{max}}) = (C^{\mathrm{max}}, G, f^{\mathrm{max}})$, where $f^{\mathrm{max}} \equiv G$. By the definition of $\mathcal{A}$, $B \neq \emptyset$, and hence $g(B, C^{\mathrm{max}}, \emptyset) = G \neq \emptyset$. Recalling the construction of ϕ in (5.3)–(5.4), we have, regardless of the values of B and f,

$$A^{(0)} = g(B, C^{\mathrm{max}}, \emptyset) = G,$$

and therefore $A^{(j_0)} = G$. This implies that $\bar{B}^{(0)} = f(G) = G$, and hence $\bar{A}^{(0)} = g(G, C^{\mathrm{max}}, A) = G$, regardless of what A is. It follows that

$$\phi(B, C^{\mathrm{max}}, f) = (G, f^{\mathrm{max}}),$$

as required.

Next we show that $T\big(i(C^{\max}), i(C)\big) = 1$ for every $C \in \mathcal{C}$. Since $i(C^{\max}) = (C^{\max}, G, f^{\max})$, the requirement that $g(G, C, \emptyset) \neq \emptyset$ is clearly satisfied. Recalling the construction of ϕ in (5.3)–(5.4), we have for any $A \subset G$,

$$A^{(0)} = g(G, C, \emptyset) = A^{(j_0)},$$
$$\bar{B}^{(0)} = f^{\max}(A^{(j_0)}) = G,$$
$$\bar{A}^{(1)} = g(G, C, A) = \bar{A}^{(\bar{j}_0)}.$$

This shows that $\phi(G, C, f^{\max}) = \big(g(G, C, \emptyset), g(G, C, \cdot)\big)$, as required.

We have shown that $i(C^{\max})$ can be reached from any state, and any $i(C)$ can be reached from $i(C^{\max})$. By the definition of $\mathcal{A}$, any state can be reached from some $i(C)$, and hence from $i(C^{\max})$. Using again that $i(C^{\max})$ can follow any state, we see that no periodicity issue can arise, and hence T is transitive. $\square$

Acknowledgements

References

[1] A.A. Ali and D. Dhar (1995) Structure of avalanches and breakdown of simple scaling in the Abelian sandpile model in one dimension. *Phys. Rev. E* **52**, 4804–4818.

[2] 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. Erratum: *Commun. Math. Phys.* **264**, 843 (2006).

[3] P. Bak (1996) *How Nature Works. The Science of Self-Organized Criticality.* Copernicus, New York.

[4] R. Burton and J.E. Steif (1994) Non-uniqueness of measures of maximal entropy for subshifts of finite type. *Ergodic Theory Dynam. Systems* **14**, 213–235.

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

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

[7] A. Björner, L. Lovász and P.W. Shor (1991) Chip-firing games on graphs. *European J. Combin.* **12**, 283–291.

[8] J. Buzzi (2005) Subshifts of quasi-finite type. *Invent. Math.* **159**, 369–406.

[9] I.P. Cornfeld, S.V. Fomin and Ya.G. Sinai (1982) *Ergodic Theory.* Grundlehren der mathematischen Wissenschaften **245**, Springer, New York.

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

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

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

[13] D. DHAR (2006) Theoretical studies of self-organized criticality. *Physica A* **369**, 29–70.

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

[15] W. FELLER (1968) *An Introduction to Probability Theory and Its Applications*, Volume I. Wiley, New York.

[16] O. HÄGGSTRÖM (1994) Aspects of spatial random processes. Ph. D. Thesis, Chalmers University of Technology and Göteborg University.

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

[18] A.A. JÁRAI (2005) Thermodynamic limit of the Abelian sandpile model on $\mathbf{Z}^d$. *Markov Process. Related Fields* **11**, 313–336.

[19] A.A. JÁRAI AND F. REDIG (2007) Infinite volume limit of the Abelian sandpile model in dimensions $d \geq 3$. To appear in *Probab. Theory Relat. Fields*. `http://arxiv.org/abs/math/0408060v3`

[20] H.J. JENSEN (2000) *Self-Organized Criticality. Emergent Complex Behavior in Physical and Biological Systems.* Cambridge Lecture Notes in Physics, **10**, Cambridge University Press.

[21] A. KATOK AND B. HASSELBLATT (1995) *Introduction to the Modern Theory of Dynamical Systems.* Encyclopedia of Mathematics and Its Applications **54**, Cambridge University Press, New York.

[22] D. LIND AND B. MARCUS (1995) *An Introduction to Symbolic Dynamics and Coding.* Cambridge University Press.

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

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

[25] C. MAES, F. REDIG AND E. SAADA (2005) Abelian sandpile models in infinite volume. *Sankhya, the Indian Journal of Statistics* **67**, 634–661.

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

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

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

[29] F. REDIG (2006) Mathematical aspects of the Abelian sandpile model. In: *Lecture Notes of Les Houches Summer School* 2005, *Mathematical Statistical Physics, Session* LXXXIII. A. Bovier, F. Dunlop, A. van Enter, F. den Hollander and J. Dalibard (eds.), Elsevier, Amsterdam, 657–730.