

Limiting shapes for deterministic centrally seeded growth models

Anne Fey-den Boer* Frank Redig†

October 4, 2007

Abstract

We study the rotor router model and two deterministic sandpile models. For the rotor router model in \mathbb{Z}^d , Levine and Peres proved that the limiting shape of the growth cluster is a sphere. For the other two models, only bounds in dimension 2 are known. A unified approach for these models with a new parameter h (the initial number of particles at each site), allows to prove a number of new limiting shape results in any dimension $d \geq 1$.

For the rotor router model, the limiting shape is a sphere for all values of h . For one of the sandpile models, and $h = 2d - 2$ (the maximal value), the limiting shape is a cube. For both sandpile models, the limiting shape is a sphere in the limit $h \rightarrow -\infty$. Finally, we prove that the rotor router shape contains a diamond.

1 Introduction

In a growth model, there is a dynamical rule by which vertices of a graph are added to an initial collection. The existing literature on growth models deals mainly with stochastic growth models. Among many stochastic growth models, there are for example the Eden model, the Richardson model, first and last passage percolation and diffusion limited aggregation. An introduction to stochastic growth models and limiting shape theorems can be found in [6].

An example related to the models in this paper, is the internal diffusion limited aggregation (IDLA) model. One adds particles one by one to the origin, letting each particle perform a random walk that stops when it hits an empty site. The growth cluster is then the collection of sites that contain a particle.

The IDLA model on \mathbb{Z}^d has been studied by Lawler, Bramson and Griffeath. In their paper from 1992 [9], they prove that the limiting shape of the growth cluster for this model in any dimension is a sphere. Lawler [10] estimated the speed of convergence.

The three deterministic models to be discussed in this paper can be viewed as deterministic analogues of IDLA, and are all examples of a general height-arrow model studied in [3]. The closest analogue is the rotor router model, proposed by Propp (see [8]). Once more, particles are added to the origin of the d -dimensional lattice \mathbb{Z}^d , but now they perform a deterministic walk as follows: at each site, there is an arrow present, pointing at one of the $2d$ neighboring sites. If a particle at this site finds it already occupied, then it takes a step in the direction of the arrow, and the arrow rotates to the next position.

*Vrije Universiteit, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands, fey@eurandom.tue.nl

†Mathematisch Instituut Universiteit Leiden, Niels Bohrweg 1, 2333 CA Leiden, The Netherlands, redig@math.leidenuniv.nl

In our other models, occupied sites hold the particles until they can be sent out in pairs, in the directions of a two-pointed arrow (double router model), or groups of $2d$, one in every direction (sandpile model). In the last model, exactly the same amount of particles is sent in every direction. In the other two models, the difference is at most one.

The last two models are deterministic versions of sandpile models. A stochastic sandpile model, where addition occurs at a random site, has originally been proposed to study self-organized criticality [1]. Soon after, Dhar [4] introduced the abelian sandpile model, that has the symmetric toppling rule (“toppling” of a site means that particles are sent out), which is now the most widely studied variant. Manna’s sandpile model [16] is doubly stochastic: the addition site is random, and moreover, in a toppling two particles are sent to randomly chosen neighbors.

The rotor router model on \mathbb{Z}^d has been studied by Peres and Levine [11, 12]. They have found that the limiting shape of this model is also a sphere, and give bounds for the rate of convergence. This result is at the basis of a number of new results in the present paper. Recently, a new paper of these authors appeared [13], further extending their results.

The deterministic abelian sandpile model has been studied on a finite square grid [15, 17, 20], with emphasis on the dynamics, but hardly as a growth model. Le Borgne and Rossin [2] found some bounds for the limiting shape for $d = 2$.

From the above introductory description, it should be clear that all the deterministic models we introduced here are closely related. In Section 2, we define the above models on \mathbb{Z}^d in a common framework, to allow comparisons of the models. We introduce a parameter h to parametrize the initial configuration. For all models, a (large) number n of particles starts at the origin, and spreads out in a deterministic way through topplings. The rest of the grid initially contains a number h of particles at each site, where h can be negative. One can imagine a negative amount of particles as a hole that needs to be filled up; admitting negative particle numbers will be helpful in comparing the different models. Once every site is stable, (i.e., has a number of particles at most some maximal allowed number), the growth cluster is formed.

We present pictures, obtained by programming the models in Matlab, of growth clusters for finite n and $d = 2$. For all models, we obtain beautiful, self-similarly patterned shapes. The appeal of these pictures has been noted before, e.g., the sandpile pictures for $h = 0, -1$ can be found on the internet as “sandpile mandala”, but the patterns are so far unexplained. Similar patterns are found in the so-called sandpile identity configuration [2, 14].

Finally, we explain our main limiting shape results for each model. For the sandpile model, we obtain that the limiting shape is a cube for $h = 2d - 2$, (Theorem 4.1), and a sphere as $h \rightarrow -\infty$ (Theorem 4.7). This last result also applies for the double router model. We remark that the sandpile model with $h \rightarrow -\infty$, strongly resembles the divisible sandpile introduced by Levine and Peres [13], of which the limiting shape is also a sphere.

For the rotor router model, we find that the limiting shape is a sphere for all h (Theorem 4.3). Finally, we generalize the bounds of Le Borgne and Rossin for growth cluster of the sandpile model, to all h and d . As a corollary, we find that the rotor router shape contains a diamond, and is contained in a cube.

The rest of the paper contains proofs of these results. In Section 3, we derive inequalities for the growth clusters. In Section 4, we prove the various limiting shape theorems.

2 Model definitions and results

Before introducing each of the three models, we present some general definitions that apply to each model. All models are defined on \mathbb{Z}^d . We define a configuration $\eta = (H, T, D)$ as consisting of the

following three components: the particle function $H : \mathbb{Z}^d \rightarrow \mathbb{Z}$, the toppling function $T : \mathbb{Z}^d \rightarrow \mathbb{Z}$ and the direction function $D : \mathbb{Z}^d \rightarrow \mathcal{D} = \{0, 1, \dots, 2d - 1\}$. We will use the $2d$ possible values in \mathcal{D} to indicate the $2d$ unit vectors $\mathbf{e}_0 \dots \mathbf{e}_{2d-1}$. The results will not depend on which value is assigned to which vector, as long as the assignment is fixed.

We call a configuration *allowed* if

1. For all \mathbf{x} , $T(\mathbf{x}) \geq 0$,
2. For all \mathbf{x} with $T(\mathbf{x}) > 0$, $H(\mathbf{x}) \geq 0$,
3. For all \mathbf{x} with $T(\mathbf{x}) = 0$, $H(\mathbf{x}) \geq h$,
4. For all \mathbf{x} , $D(\mathbf{x}) \in \mathcal{D}$.

Here, $h \in \mathbb{Z}$, the “background height”, is a model parameter.

Each model starts with initial configuration $\eta_0 = \eta_0(n, h)$, which is as follows: $H_0(\mathbf{x}) = n$ if $\mathbf{x} = \mathbf{0}$, and $H_0(\mathbf{x}) = h$ otherwise, $T_0(\mathbf{x}) = 0$ for all \mathbf{x} , and $D_0(\mathbf{x}) \in \mathcal{D}$ for all \mathbf{x} . Observe that η_0 is allowed.

We now define a toppling as follows:

Definition 2.1 A toppling of site \mathbf{x} in configuration η consists of the following operations:

- $T(\mathbf{x}) \rightarrow T(\mathbf{x}) + 1$,
- $H(\mathbf{x}) \rightarrow H(\mathbf{x}) - c$, with $c \leq 2d$ some value specific for the model,
- $H(\mathbf{x} + \mathbf{e}_i) \rightarrow H(\mathbf{x} + \mathbf{e}_i) + 1$, with $i = (D(\mathbf{x}) + 1) \bmod 2d, \dots, (D(\mathbf{x}) + c) \bmod 2d$,
- $D(\mathbf{x}) \rightarrow (D(\mathbf{x}) + c) \bmod 2d$.

In words, in a toppling c particles from site \mathbf{x} move to c different neighbors of \mathbf{x} , chosen according to the value of $D(\mathbf{x})$ in cyclic order. We call a toppling of site \mathbf{x} *legal* if after the toppling $H(\mathbf{x}) \geq 0$. We call a site \mathbf{x} *stable* if $H(\mathbf{x}) \leq h_{\max}$, with $h_{\max} = c - 1$, so that a site can legally topple only if it is unstable.

We can now define stabilization of a configuration as performing legal topplings, such that a stable configuration is reached, that is, a configuration where all sites are stable. We call this configuration the final configuration $\eta_n = \eta_n(n, h)$, which will of course depend on the model. To ensure that the final configuration is reached in a finite number of topplings, we impose for each model $h < h_{\max}$. For each model, it is then known that η_n does not depend on the order of the topplings. This is called the abelian property. The abelian property has been proved for centrally seeded growth models in general ([5], Section 4), and for the sandpile model in particular [18].

In the remainder of this paper, we will choose various orders of topplings. In Section 4.2, we will even admit illegal topplings. In Section 4.4, we will organize the topplings into discrete time steps, obtaining after time step t a configuration η^t , consisting of H^t , T^t and D^t .

We define two growth clusters that are formed during stabilization, denoting by \mathbf{x}^\square the unit cube centered at \mathbf{x} (i.e., $\mathbf{x}^\square = \{\mathbf{y} : \mathbf{y} = \mathbf{z} + \mathbf{x}, \mathbf{z} \in [-\frac{1}{2}, \frac{1}{2}]^d\}$):

Definition 2.2 The toppling cluster is the cluster of all sites that have toppled, that is,

$$\mathcal{T}_n = \bigcup_{\mathbf{x}: T_n(\mathbf{x}) > 0} \mathbf{x}^\square,$$

and the particle cluster is the cluster of all sites that have been visited by particles from the origin, that is,

$$\mathcal{V}_n = \mathcal{T}_n \cup \bigcup_{\mathbf{x}: H_n(\mathbf{x}) > h} \mathbf{x}^\square.$$

A cluster as a function of n has a limiting shape if, appropriately scaled, it tends to a certain shape as $n \rightarrow \infty$, in some sense to be specified later. By the model definition, all clusters are path connected. Observe that from Definitions 2.1 and 2.2, it follows that

$$\mathcal{T}_n \subset \mathcal{V}_n \subseteq \mathcal{T}_n \cup \partial\mathcal{T}_n, \quad (1)$$

with $\partial\mathcal{T}_n$ the exterior boundary of \mathcal{T}_n . We also define the “lattice ball”:

$$\mathcal{B}_n = \bigcup_{i=1}^n \mathbf{x}_i^\square, \quad (2)$$

where the lattice sites $\mathbf{x}_1, \mathbf{x}_2, \dots$ of \mathbb{Z}^d are ordered in such a way that the Euclidean distance from the origin is non-decreasing.

2.1 The rotor-router (RR, h) model

For the rotor router model, $c = 1$, so that in each toppling only one particle moves to a neighbor. Therefore, $h_{\max} = 0$, and the model is defined for $h < 0$ only. In words, every site holds the first $|h|$ particles that it receives, and sends every next particle to a neighbor, choosing the neighbors in a cyclic order. This means that after $2d$ topplings, every neighbor received a particle. Instead of a direction function with numerical values, we can think of an arrow being present at every site. In a toppling, the arrow is rotated to a new direction, and the particle is sent in this new direction.

Propp, Levine and Peres studied the case $h = -1$. Levine and Peres have proven that for $h = -1$, the limiting shape of the rotor router model is a Euclidean sphere. More precisely, they showed ([12], Thm. 1.1):

$$\lim_{n \rightarrow \infty} \lambda(n^{-1/d} \mathcal{V}_n \triangle B) = 0, \quad (3)$$

where λ denotes d -dimensional Lebesgue measure, \triangle denotes symmetric difference, and B is the Euclidean sphere of unit volume centered at the origin in \mathbb{R}^d . This result is independent of D_0 , and of any assignment of different unit vectors to possible values of $D(\mathbf{x})$. Note that the scaling function is necessarily $n^{-1/d}$, since $|\mathcal{V}_n| = n$ by model definition.

In [8], a picture of D_n is given for $h = -1$ and $n = 3$ million and D_0 constant. The picture shows a circular shape with an intriguing seemingly self-similar pattern. Curiously enough, this picture actually shows the shape of \mathcal{T}_n rather than \mathcal{V}_n , indicating that the limiting shape of \mathcal{T}_n is also a sphere. This would be a stronger statement than (3), but it remains as yet unproven.

It has been noted that the shape of \mathcal{V}_n for the rotor-router model is remarkably circular, that is, as close to a circle as a lattice set can get, for every n . However, the shape of \mathcal{T}_n has not been studied before. We programmed the model for several values of h and n , and observed for all these values that $\mathcal{V}_n \setminus \mathcal{T}_n$ is concentrated on the inner boundary of \mathcal{V}_n .

Our first main result of the rotor router model is the generalization of (3) to all $h \leq -1$, stated in Theorem 4.3. The proof in fact uses (3) as main ingredient.

By an entirely different method, in fact as a corollary of Theorem 4.8, we moreover obtain the result that the rotor router shape contains a diamond of radius proportional to $(\frac{n}{2d-1-h})^{1/d}$. It is surprising that, in spite of the limiting shape of the RR, h model being known, this is still a new result. But, in the words of Levine and Peres, the convergence in (3) does not preclude the formation of, e.g., holes close to the origin, as long as their volume is negligible compared to n . The only comparable previous result is that for $d = 2$ and $h = -1$, the particle cluster contains a disk of radius proportional to $n^{1/4}$ [11].

2.2 The abelian sandpile (SP, h) model

For the abelian sandpile model, $c = 2d$. Therefore, $h_{\max} = 2d - 1$, and the model is defined for $h < 2d - 1$. In each toppling, one particle moves to every neighbor, so that in fact the value of $D(\mathbf{x})$ is irrelevant.

It follows that for the SP model, we can specify (1) to

$$\mathcal{V}_n = \mathcal{T}_n \cup \partial\mathcal{T}_n. \tag{4}$$

The SP model as a growth model has received some attention in the cases $h = -1$ (“greedy” sandpile) and $h = 0$ (“non-greedy”), for which pictures of H_n can be found [8]. It is noted that the shape does not seem to be circular. In Figure 1, we show a family of sandpile pictures for a range of values of h , obtained by programming the model in Matlab. We see the number of symmetry axes increasing as h decreases. The shape appears to become more circular as h decreases. The shape for $h = 2$ is observed to be a square, but for the other values of h it seems to tend to a more complicated shape.

Again, we find that $\mathcal{V}_n \setminus \mathcal{T}_n$ is concentrated on the inner boundary of \mathcal{V}_n , for all observed values of h .

Figure 1: H_n for the sandpile model with different h values. The number of particles ranges from $n = 60,000$ ($h = 2$) to $1,000,000$ ($h = -20$). The gray-scale colors are such that a lighter color corresponds to a higher value of $H(\mathbf{x})$.

For the sandpile model, we have the following results. Theorem 4.1 states that indeed the toppling cluster for the SP, $2d - 2$ model is a (d -dimensional) cube, and the particle cluster tends to a cube as $n \rightarrow \infty$.

Theorem 4.7 states that for $h \rightarrow -\infty$, the limiting shape is a sphere.

Finally, we generalize some bounds for the scaling function that have been obtained by Le Borgne and Rossin [2] in the case $d = 2$, $0 \leq h \leq 2$, to all d and h . This result is formulated in Theorem 4.8. Our proof moreover allows to deduce that \mathcal{V}_n for the sandpile model is simply connected for all n .

2.3 The double router (DR, h) model

In the double router model, $c = 2$. Therefore, $h_{\max} = 1$, and the model is defined for $h \leq 0$. In a toppling of this model, two particles are sent out in two different directions, such that after d topplings, every neighbor received a particle.

Many variants of this model are possible, e.g., for $d = 3$ one could choose $c = 3$, which amounts to dividing the 6 neighbors in 2 groups of 3. However, to avoid confusion, we only discuss the above explained variant.

Figure 2: H_n for the DR model with $h = 0$ ($n = 110,000$), $h = -1$ ($n = 100,000$), and $h = -5$ ($n = 400,000$), and $D_0(\mathbf{x}) = 0$ for all \mathbf{x} . Sites with height 1 are colored white, sites with height 0 or negative are colored black.

Figure 2 shows H_n for $h = 0$, $d = 2$ and $n = 110,000$. We ordered the unit vectors as $\mathbf{e}_0 = \text{left}$, $\mathbf{e}_1 = \text{right}$, $\mathbf{e}_2 = \text{up}$, $\mathbf{e}_3 = \text{down}$. Initially, $D(\mathbf{x}) = 0$ for all \mathbf{x} . Figure 3 shows D_n for the same case, to indicate that for this model we find complex patterns both for H_n and D_n .

As for the sandpile model, we see a symmetric, yet not circular shape, with notable flat edges. In Figure 2, also some other values of h are shown, again indicating that the shape seems to become

Figure 3: D_n for the DR model with $h = 0$ ($n = 110,000$), and $D_0(\mathbf{x}) = 0$ for all \mathbf{x} . White corresponds to $D(\mathbf{x}) = 2$, black to $D(\mathbf{x}) = 0$.

more and more circular with decreasing h . Indeed, Theorem 4.7 states this fact for both the sandpile model and the double router model.

Again, we observe that $\mathcal{V}_n \setminus \mathcal{T}_n$ is concentrated on the inner boundary of \mathcal{V}_n , for all observed values of h .

3 Comparing the models

From the model descriptions, it is clear that one SP toppling always equals d DR topplings, as well as $2d$ RR topplings. Furthermore, our toppling definition (Definition 2.1) ensures that one DR toppling always equals two consecutive RR topplings. We exploit these relations to compare the clusters for different models. We will with a subindex RR, DR or SP indicate the model that was used to obtain the growth cluster, and also use a subindex h because we compare different h values.

Proposition 3.1 For every fixed h (for which the model is defined), d and n ,

1. $\mathcal{V}_{n,SP,h} \subseteq \mathcal{V}_{n,DR,h} \subseteq \mathcal{V}_{n,RR,h}$,
2. For all models $i = \text{RR, DR and SP}$: $\mathcal{V}_{n,i,h-1} \subseteq \mathcal{V}_{n,i,h}$,
3. $\mathcal{V}_{n,RR,h-1} \subseteq \mathcal{V}_{n,DR,h}$, and $\mathcal{V}_{n,RR,h-(2d-1)} \subseteq \mathcal{V}_{n,SP,h}$.

Proof. The proof makes use of the abelian property of all models, that is, the property that the stabilized configuration η_n does not depend on the order of topplings. We are therefore free to choose a convenient order. Furthermore, a site \mathbf{x} that at some instant during stabilization has either $T(\mathbf{x}) > 0$, or $T(\mathbf{x}) = 0$ and $H(\mathbf{x}) > h$, must belong to \mathcal{V}_n in the final configuration, because by further topplings either $H(\mathbf{x})$ or $T(\mathbf{x})$, or both, increase.

part 1. The initial configurations for all these three models are the same. We first compare $\mathcal{V}_{n,SP,h}$ with $\mathcal{V}_{n,DR,h}$. We choose, for both the SP and the DR model, to first perform all legal SP-topplings. Since every SP-toppling consists of d DR-topplings, these are legal topplings for both models. The configuration is now stabilized for the SP model, but there can be sites in $\mathcal{V}_{n,SP,h}$ that are not stable in the DR model, since h_{max} is $2d - 1$ for the SP model and 1 for the DR model. Therefore, in the MD model, possibly more topplings follow, so that $\mathcal{V}_{n,SP,h} \subseteq \mathcal{V}_{n,DR,h}$.

We compare $\mathcal{V}_{n,DR,h}$ with $\mathcal{V}_{n,RR,h}$ in an analogous manner, this time using that h_{max} is 1 for the DR model and 0 for the RR, h model, to get $\mathcal{V}_{n,DR,h} \subseteq \mathcal{V}_{n,RR,h}$.

part 2. We start with comparing the RR, $h - 1$ model with the RR, h model. We choose, for both initial configurations, to first perform all RR-topplings that stabilize the initial configuration of the RR, $h - 1$ model. These are legal topplings for both configurations. The configuration is now stable for the RR, $h - 1$ model, but all sites that have $H(\mathbf{x}) = 0$ for this model, have $H(\mathbf{x}) = 1$ for the RR, h model. Therefore, for this model more topplings would follow, so that $\mathcal{V}_{n,RR,h-1} \subseteq \mathcal{V}_{n,RR,h}$. The same reasoning can be applied for the DR and SP model.

part 3. We first compare $\mathcal{V}_{n,DR,h}$ with $\mathcal{V}_{n,RR,h-1}$. We choose, for both models, to first perform all DR-topplings that would stabilize the initial configuration with height $h - 1$. These are legal topplings for both models. Then for both models, more topplings are needed. For the RR, $h - 1$ model there can be sites with $H(\mathbf{x}) = 1$, that are unstable. For the DR, h model, the same sites have $H(\mathbf{x}) = 2$, and therefore are also unstable. But since one DR-toppling equals two RR-topplings, the

set of sites where the configuration changes by the extra topplings for the $RR, h-1$ model, is a subset of those for the DR, h model. Therefore, $\mathcal{V}_{n,RR,h-1} \subseteq \mathcal{V}_{n,DR,h}$. The argument for the sandpile model is similar. \square

4 Limiting shape results

4.1 The sandpile model in \mathbb{Z}^d , with $h = 2d - 2$

As noted in Section 2, Figure 1 indicates that the limiting shape for the sandpile model with $h = 2$ and $d = 2$ is a square. This section contains the proof of a more general statement for arbitrary dimension, that is, we prove that indeed the toppling cluster for the $SP, 2d - 2$ model is a (d -dimensional) cube, and the particle cluster tends to a cube as $n \rightarrow \infty$. However, we have no explicit expression for the scaling function $f(n)$, so that the scaled clusters $f(n)\mathcal{V}_n$ and $f(n)\mathcal{T}_n$ would tend to the unit cube, just that this scaling satisfies $n^{-1} \leq f(n) \leq \frac{1}{2}n^{-1/d} - \frac{3}{2}$. Based on the calculations, we believe that $f(n)$ is $O(n^{-1/d})$.

Theorem 4.1 Let $\mathfrak{C}(r)$ be the cube $\bigcup_{\mathbf{x}: \max_i |x_i| \leq r} \mathbf{x}^\square$. For every n in the d -dimensional $SP, 2d - 2$ model, there is an r_n such that

$$\mathcal{T}_n = \mathfrak{C}(r_n),$$

and


$$\mathcal{V}_n = \mathfrak{C}(r_n) \cup \partial\mathfrak{C}(r_n).$$

For all n , this r_n satisfies

$$\frac{1}{2}n^{1/d} - \frac{3}{2} \leq r_n \leq n.$$

We start by outlining the case $d = 2$ as an example, for the sake of clarity. When simulating the model for several small values of n - this can even be done by hand - one notices that the configuration is always such that it contains a central square with all boundary sites full, except for the corner sites, which have height $h_{max} - 1$. Outside this square, all sites have height $h_{max} - 1$. We will call such a rectangular boundary a critical boundary. This is all the information we need to make an inductive argument in n . Suppose, η_n has a critical boundary, we add one grain to the origin and in the course of stabilization (to obtain η_{n+1}), one boundary site topples. As a consequence, all neighboring boundary sites topple, because they are all full and all in turn receive a grain. The two adjacent corner sites receive one grain, therefore become full. Therefore, after these topplings a new rectangular critical boundary is created. If any more boundary sites topple, we can reiterate this argument. We conclude that the presence of a rectangular critical boundary is stable under additions inside this rectangle. In the special case where additions are made only to the origin, we conclude by symmetry that the shape of the critical boundary will always be square.

Below, we schematically show the argument for $n = 7$.



We start with η_7 plus one extra grain at the origin. Full boundary sites are colored lightgrey, unstable sites grey. First the origin topples, causing the full boundary sites to become unstable. Next, these boundary sites topple, causing the corner sites to be unstable. The first full sites of the new boundary are created.

2	3	3	3	2	→	2	3	3	3	2
3	0	2	0	3		3	0	3	3	3
3	2	4	2	3		3	3	0	3	3
3	0	2	0	3		3	0	3	3	3
2	3	3	3	2		2	3	3	3	2

Next, the corner sites topple. The configuration now has a new square critical boundary. A further toppling of the origin, after which we obtain η_8 , does not change this new boundary anymore.

In the proof below, we generalize this critical boundary to arbitrary d and make the details more precise, but the idea will remain the same.

In the course of the proof, we will need the following lemma.

Lemma 4.2 Let ϕ_r be the following configuration: $H(\mathbf{0}) = 2d$, and for all other $\mathbf{x} \in \mathfrak{C}_r$, $H(\mathbf{x}) = h_{max}$, otherwise $H(\mathbf{x}) = h_{max} - 1$. If the configuration ϕ_r is stabilized, then during stabilization, every site \mathbf{x} topples exactly $\delta_{\mathbf{x},r} = \max\{r + 1 - \max_i |x_i|, 0\}$ times.

Proof. We choose to order the topplings into *waves* [7], that is, in each wave, we topple the origin once and then all other sites that become unstable, except the origin again. No site can topple more than once during a wave.

In the first wave, all sites in \mathfrak{C}_r topple once, because it is maximally filled. The wave stops at the boundary, because the sites outside \mathfrak{C}_r have at most one toppling neighbor, therefore they can not become unstable. All sites in \mathfrak{C}_{r-1} then have $2d$ once-toppling neighbors, so their particle number does not change. The sites in $\mathfrak{C}_r \setminus \mathfrak{C}_{r-1}$ have less toppling neighbors, so their particle number becomes at most $2d - 2$.

Therefore, the effect of the first wave on ϕ_r , is to make it, for all sites in \mathfrak{C}_r , at most equal to ϕ_{r-1} , with equality for all sites with \mathfrak{C}_{r-1} . It follows that the next wave will topple all sites in \mathfrak{C}_{r-1} once. Continuing this argument, the result stated in the lemma follows. \square

Proof of Theorem 4.1. We will use induction in n . For that, we choose to obtain the final configuration as follows: the n particles are added to the origin one by one, each time first stabilizing the current configuration through topplings. Due to abelianness, this procedure will give the same final configuration as when all n particles are added simultaneously, before toppling starts. We will show that during this procedure, only configurations in \mathcal{C}_r are encountered, where \mathcal{C}_r is the set of configurations that are as follows:

1. $\mathcal{T}_n = \mathfrak{C}_r$,
2. For all $\mathbf{x} \in \mathcal{T}_n$ such that $\max_i |x_i| = r$, $T_n(\mathbf{x}) = 1$,
3. $\mathcal{V}_n = \mathfrak{C}_r \cup \partial\mathfrak{C}_r$,
4. For all $\mathbf{x} \in \mathcal{V}_n \setminus \mathcal{T}_n$, $H_n(\mathbf{x}) = h_{max} = 2d - 1$,
5. For all $\mathbf{x} \notin \mathcal{V}_n$, $H_n(\mathbf{x}) = 2d - 2$.

In words, a configuration in \mathcal{C}_r is such that \mathcal{T}_n is a cube, where all inner boundary sites toppled exactly once, and all outer boundary sites, i.e., sites that do not belong to \mathcal{T}_n but have a neighbor in \mathcal{T}_n , have particle number h_{max} . Thus, \mathcal{T}_n is surrounded by $2d$ maximally filled square “slabs” of size $(2r+1)^{d-1}$. All sites outside \mathcal{V}_n have particle number $2d-2$, by model definition.

After the first particle is added, $H_1(\mathbf{0}) = h_{max}$, and no site has toppled yet. After the second particle is added, the origin topples once, and all neighbors of the origin receive one particle. Therefore, $\eta_2 \in \mathcal{C}_0$.

We will now show that if we assume $\eta_n \in \mathcal{C}_r$ for some r , then either $\eta_{n+1} \in \mathcal{C}_r$ or $\eta_{n+1} \in \mathcal{C}_{r+1}$. To prove this, we will add a particle at the origin to configuration η_n and start stabilizing. First we show that if during stabilization the configuration leaves \mathcal{C}_r , then it enters \mathcal{C}_{r+1} . Then we show that in this last case, it does not leave \mathcal{C}_{r+1} during further stabilization.

If no site $\mathbf{x} \in \mathcal{V}_n$ such that $\max_i |x_i| = r$ topples during stabilization, then η_{n+1} remains in \mathcal{C}_r . But if one such site topples, then one site in $\mathcal{V}_n \setminus \mathcal{T}_n$ becomes unstable, and also topples. This site is in one of the $2d$ maximally filled slabs. If one site of such a slab topples, then the entire slab must topple, because all of its sites will in turn receive a particle from a toppling neighbor. After the entire slab toppled once, all neighbors of the slab received a particle, and the configuration can be described as follows: The toppled cluster is a rectangle of size $(2r+1)^{d-1}(2r+2)$, surrounded by maximally filled slabs, two of them cubic and the rest rectangular.

But due to symmetry, if one slab topples once then this must happen for all slabs. We can choose in what order to topple the slabs; suppose we divide them in d opposite pairs. After we toppled the first pair of slabs, the configuration is a central rectangle of size $(2r+1)^{d-1}(2r+3)$, surrounded by maximally filled slabs, two of them cubic and the rest rectangular. After we toppled the k^{th} pair of (meanwhile possibly rectangular) slabs, the configuration is a central rectangle of size $(2r+1)^{d-k}(2r+3)^k$, surrounded by maximally filled slabs, so that after all slabs toppled, the toppled cluster is again a cube, now of size $(2r+3)^d$, centered at the origin and surrounded by $2d$ maximally filled square slabs. In other words, after these topplings the configuration is in \mathcal{C}_{r+1} . It follows that if $\eta_n \in \mathcal{C}_r$ for some r , then $\eta_{n+1} \in \mathcal{C}_{r'}$, for some $r' \geq r$. It now remains to show that r' can only have the values r or $r+1$.

We use lemma 4.2. A configuration $\eta \in \mathcal{C}_r$, plus an addition at the origin, has at every site at most the number of particles given by ϕ_{r+1} . Therefore, if we suppose $\eta_n \in \mathcal{C}_r$, then upon addition of a particle at the origin, by abelianness, the number of topplings for every \mathbf{x} will at most be $\delta_{\mathbf{x}, r+1}$. In particular, the outer boundary sites of \mathcal{T}_n will topple at most once.

The induction proof is now completed, therefore we now know that, for all n , $\eta_n \in \mathcal{C}_r$ for some r . We have also shown that $r_n \leq n$. From the description of \mathcal{C}_r , we see that \mathcal{T}_n is always a cube, and that \mathcal{V}_n is more and more like a cube as r_n increases.

To conclude that \mathcal{V}_n tends to a cube as $n \rightarrow \infty$, we finally need to show that r_n increases with n . For this, note that in η_n every site needs to be stable. Since $h = h_{max} - 1$, every site can accommodate at most one extra particle. Therefore, if $\eta_n \in \mathcal{C}_r$, then $r \geq \frac{1}{2}n^{1/d} - \frac{3}{2}$. \square

4.2 The rotor-router model with $h < -1$

In this section, we will prove the following result:

Theorem 4.3 The limiting shape of the particle cluster for the rotor-router model is a sphere for every $h \leq -1$. More precisely,

$$\lim_{n \rightarrow \infty} \lambda \left(\left(\frac{n}{|h|} \right)^{-1/d} \mathcal{V}_n \triangle B \right) = 0.$$

The proof will also reveal the following about the rotor-router toppling cluster:

Corollary 4.4 The toppling cluster of the RR,h model contains a cluster \mathcal{W}_n , with

$$\lim_{n \rightarrow \infty} \lambda \left(\left(\frac{n}{|h|+1} \right)^{-1/d} \mathcal{W}_n \triangle B \right) = 0.$$

Theorem 4.3 has been proven for the case $h = -1$, by Levine and Peres, see (3). We will use their result to prove the theorem for other values of h . The strategy of the proof will be as follows. First, we introduce a slightly different version of the RR,h model, which we will call the k -color rotor router, or RR',k model. We show that the limiting shape of this model is a sphere for all k , with scaling function $(\frac{n}{k})^{-1/d}$. Informally, the k -color model can be viewed as k iterations of the $\text{RR},-1$ model. The n particles at the origin are equally divided in k different-colored groups, and for each color the particles spread out until there is at most one particle of each color at each site. From (3), we then know that each color region is approximately spherical, so we get a final configuration that looks like k almost completely overlapping, approximately spherical, different-colored regions.

Then we show that $\mathcal{V}_{n,\text{RR},h}$ and $\mathcal{V}_{n,\text{RR}',k}$, with $h = -k$, differ only in a number of sites that is $o(n)$. This number is at most the number of sites in $\mathcal{V}_{n,\text{RR}',k}$ where not every color is present. In the proof of this second point, we will stabilize η_0 for the RR,h model by first performing all topplings needed to stabilize the $\text{RR}',-h$ model. Some of these topplings may be illegal for the RR,h model, so we will reverse them by performing untopplings. The main difficulty in the proof is to show that we do reach η_n by this procedure. First we define untopplings.

Definition 4.5 An untoppling of site \mathbf{x} in configuration η consists of the following operations:

- $T(\mathbf{x}) \rightarrow T(\mathbf{x}) - 1$,
- $H(\mathbf{x}) \rightarrow H(\mathbf{x}) + c$, with c according to the model,
- $D(\mathbf{x}) \rightarrow (D(\mathbf{x}) - c) \bmod 2d$,
- $H(\mathbf{x} + \mathbf{e}_i) \rightarrow H(\mathbf{x} + \mathbf{e}_i) - 1$, with $i = D(\mathbf{x}), (D(\mathbf{x}) + 1) \bmod 2d, \dots, (D(\mathbf{x}) + c - 1) \bmod 2d$.

We call an untoppling of site \mathbf{x} *legal* if $T(\mathbf{x}) > 0$ before the untoppling. We see that a legal untoppling of site \mathbf{x} is precisely the undoing of a toppling of site \mathbf{x} , in the sense that the particles that were sent out of site \mathbf{x} in the last toppling, return to this site, and the value of $D(\mathbf{x})$ returns to its previous value.

The following proposition provides a more elaborate description of η_n in the case where we allow illegal topplings, that we will need in the proof of Theorem 4.3 (we remark that the abelian property still holds for combinations of legal and illegal topplings). We introduce the notion of optimality to distinguish η_n from configurations that are stable and allowed, but that cannot be obtained from η_0 by legal topplings. An example for the rotor router model is as follows. Suppose that, after reaching η_n by legal topplings, there is somewhere a closed loop of sites with height 0 such that “the arrows form a cycle”, i.e., if we topple all of them, they would all send and receive one particle. After these topplings (of which at least the first would be illegal), the height function is still H_n , but the toppling number of these sites has increased by 1, so we obtained a different, stable and allowed configuration. We remark that a set of sites forming such a “cycle of arrows” cannot be found in \mathcal{T}_n . Since all sites that toppled are full, in the last toppling in this set a particle must have left the set, so that there is an arrow pointing out of the set (see also [19], section III). Therefore, if there is a set of sites forming a cycle of arrows, then at least one of these sites did not topple.

Proposition 4.6 Call a stable configuration *optimal* if every sequence of legal untopplings leads to an unstable configuration. For every model, η_n is the unique optimal, stable and allowed configuration that can be reached from η_0 by topplings, either legal or illegal, and legal untopplings.

Proof. We defined η_n before as the unique stable configuration that can be reached from η_0 by legal topplings. It follows from this definition that η_n is allowed, since if sites can only topple when they are unstable, then we automatically obtain for all \mathbf{x} where $T_n(\mathbf{x}) > 0$, that $H_n(\mathbf{x}) \geq 0$.

To prove that η_n is optimal, we proceed by contradiction. Suppose that, starting from η_n , there is a sequence of legal untopplings such that a stable configuration ξ is obtained. Then ξ can be obtained from η_0 by a sequence of legal or illegal topplings. Call T' the toppling function for this sequence, then $0 \leq T'(\mathbf{x}) \leq T_n(\mathbf{x})$ for all \mathbf{x} (because we undid some topplings of T_n), but $T' \neq T_n$. By abelianness, ξ depends only on η_0 and T' . Thus, we can choose the order of the topplings according to T' , to obtain ξ . There cannot be an order such that all topplings are legal, otherwise ξ cannot be different from η_n . We will therefore choose the order such that first all possible legal topplings are performed, and then the rest (as an example, suppose $T'(\mathbf{0}) = 0$. Then we must start with an illegal toppling, since in η_0 all sites but the origin are stable). After all possible legal topplings according to T' , we have a configuration with at least one unstable site, since we did not yet reach η_n . In the remainder, this site is not toppled, because otherwise another legal toppling could have been added to the first legal topplings. Therefore, ξ cannot be stable.

We so far established that η_n is optimal, stable and allowed. Now we prove that it is unique by deriving another contradiction. Suppose there is an other optimal, stable and allowed configuration ζ that can be reached from η_0 by topplings and legal untopplings. Call T'' the toppling function for this sequence. Then $T''(\mathbf{x}) \geq T_n(\mathbf{x})$ for all \mathbf{x} . We can see this as follows: we choose the toppling order such that we first perform all topplings that are also in T_n . If there were some sites \mathbf{y} where $T''(\mathbf{y}) < T_n(\mathbf{y})$, then at this point at least one of these sites would be unstable. But in the remainder of T'' , this site would not topple again, so in that case ζ could not be stable.

Call $\tau(\mathbf{x}) = T''(\mathbf{x}) - T_n(\mathbf{x})$. Since $\zeta \neq \eta_n$, $T'' \neq T_n$. Suppose we first perform the topplings according to T_n , and then according to τ . Then we have that ζ can be reached from η_n by performing $\tau(\mathbf{x})$ topplings for every site \mathbf{x} , and vice versa, that η_n can be reached from ζ by performing $\tau(\mathbf{x})$ legal untopplings for every site \mathbf{x} . But then ζ cannot be optimal. \square

Proof of Theorem 4.3 The RR',k model is defined as follows: As in the RR,h model, $c = 1$. Initially, there are n particles at the origin; at every other site there are $-k$ particles. We choose n a multiple of k , and divide the n particles into k groups of n/k particles, each group with a different color. The height $H(\mathbf{x})$ of a site \mathbf{x} is now defined as $-k$ plus the total number of particles present at \mathbf{x} , of either color. However, now we call a site stable if it contains at most one particle of each color. If a particle arrives at a site where its color is already present, then that particle will be sent to a neighbor in the subsequent toppling. Note that in this model, a site \mathbf{x} can be unstable even if $H(\mathbf{x}) < 0$. We furthermore restrict the order of legal topplings for this model. We say that in this model, first all legal topplings for the first color should be performed, then for the next, etc. Then for each color, the model behaves just as the $RR, -1$ model, with for each color a new initial T and D . Therefore, with this toppling order the final configuration consisting of H_n , T_n and D_n , is well-defined.

We now show that the limiting shape of the RR',k model is a sphere. We perform first all the topplings with particles of the first color, say, red. From [12], Thm. 2.1 it follows that these particles will form a cluster $\mathcal{V}_{n/k}^1$ close to the lattice ball $\mathcal{B}_{n/k}$. In fact, the number of sites in $\mathcal{V}_{n/k}^1 \triangle \mathcal{B}_{n/k}$ is $o(n)$. We will denote this as

$$|\mathcal{V}_{n/k}^1 \triangle \mathcal{B}_{n/k}| \leq f(n/k) = o(n).$$

D is now different from D_0 . However, if next we stabilize for the blue particles, these will again form a cluster close to the lattice ball $\mathcal{B}_{n/k}$, since the result (3) does not depend on the initial D . Of course, this cluster $\mathcal{V}_{n/k}^2$ need not be the same as $\mathcal{V}_{n/k}^1$.

When all sites are stable, we have $\mathcal{V}_{n,RR',k} = \bigcup_{i=1}^k \mathcal{V}_{n/k}^i$. We also define a cluster $\mathcal{W}_{n,k} = \bigcap_{i=1}^k \mathcal{V}_{n/k}^i$. From the above, the number of sites in the difference of both these clusters with $\mathcal{B}_{n/k}$, is at most $kf(n/k)$. Thus

$$|\mathcal{V}_{n,RR',k} \triangle \mathcal{B}_{n/k}| \leq kf(n/k) = o(n). \quad (5)$$

We will call $\mathcal{X}_n = \mathcal{V}_{n,RR',k} \setminus \mathcal{W}_{n,k}$, so that \mathcal{X}_n contains $|\mathcal{X}_n| \leq 2kf(n/k)$ sites. Sites in $\mathcal{W}_{n,k}$ contain k particles, sites in \mathcal{X}_n contain less than k particles.

Now we compare this model with the RR,h model, with $h = -k$. Disregarding the colors, the initial configuration is the same for both models. Suppose we perform in the RR,h model all the topplings that are performed as above in the $RR',-h$ model. The configuration is then as follows: for all $\mathbf{x} \in \mathcal{W}_{n,k}$, $H(\mathbf{x}) = 0$ and $T(\mathbf{x}) \geq 0$, and for all $\mathbf{x}' \in \mathcal{X}_n$, $-k \leq H(\mathbf{x}') < 0$ and $T(\mathbf{x}') \geq 0$.

This configuration is possibly not η_n for the RR,h model, since it is possibly not allowed, if there are sites \mathbf{x}' in \mathcal{X}_n with $T(\mathbf{x}') > 0$. We can reach an allowed configuration by performing legal untoppings; it will appear that we then in fact reach η_n .

Suppose first that only one untopping is needed. It might be that the neighbor \mathbf{y} that returned a particle, now has less than k particles, so that $H(\mathbf{y}) < 0$, while $T(\mathbf{y}) > 0$. In that case, \mathbf{y} would now also have to topple. We call this process an untopping avalanche. The avalanche stops if an allowed (stable) configuration is encountered, that is, if the last particle came from a site that did not topple. An untopping avalanche consists of untopping neighbors, each one passing a particle to the previous one. Therefore, an untopping avalanche of arbitrary length, changes the particle number of only two sites in the configuration.

In our case, at most $(k-1)|\mathcal{X}_n|$ untopping avalanches are required. Each avalanche stops if the last particle came from a site that did not topple. This will change the configuration in at most $2(k-1)|\mathcal{X}_n|$ sites. After all these untoppings, the configuration is allowed and stable, i.e., we have for all sites with $T(\mathbf{x}) > 0$, that $H(\mathbf{x}) = 0$. We can only perform legal untoppings at sites with $T(\mathbf{x}) > 0$, so that we cannot perform a sequence of legal untoppings in a closed loop of sites. Therefore, any further sequence of legal untoppings would increase $H(\mathbf{x})$ for at least one of these sites, and thus lead to an unstable configuration.

We have thus reached an optimal configuration, so we have reached η_n . Therefore, at this point we can conclude

$$|\mathcal{V}_{n,RR,h} \triangle \mathcal{V}_{n,RR',k}| \leq 2(k-1)|\mathcal{X}_n| \leq 4k(k-1)f(n/k) = o(n).$$

Combined with (5), this leads to Theorem 4.3. \square

To prove the corollary, we define $\mathcal{W}_{n,h} = \bigcup_{\mathbf{x}: \mathbf{x} \in \mathcal{V}_{n,RR,h}, H_n(\mathbf{x}) = |h|} \mathbf{x}^\square$. By the above argument, we have, as for $\mathcal{W}_{n,k}$, that $|\mathcal{W}_{n,h} \triangle \mathcal{B}_{n/|h|}| = o(n)$.

Suppose, in the RR,h model, we first perform all topplings that are needed to stabilize for the $RR,h-1$ model, as in the proof of Proposition 3.1 part 2. The particle cluster then contains a cluster $\mathcal{W}_{n,h-1}$, with $\lim_{n \rightarrow \infty} \lambda \left(\left(\frac{n}{|h-1|} \right)^{-1/d} \mathcal{W}_n \triangle B \right) = 0$.

Because the sites in $\mathcal{W}_{n,h-1}$ contain $|h-1| = |h| + 1$ particles, every site in $\mathcal{W}_{n,h-1}$ is unstable in the RR,h model. Therefore, $\mathcal{W}_{n,h-1} \subseteq \mathcal{T}_{n,RR,h}$. \square

4.3 The double router model, and the sandpile model with $h \rightarrow -\infty$

In discussing Figures 1 and 2, we observed that the DR and SP shapes seem to become more circular as h decreases. Proposition 3.1 and Theorem 4.3 can indeed be combined to give the following result:

Theorem 4.7 The limiting shape of the SP, h and the DR, h model, for $h \rightarrow -\infty$, is a sphere. More precisely,

$$\lim_{h \rightarrow -\infty} \limsup_{n \rightarrow \infty} \lambda \left(\left(\frac{n}{|h|} \right)^{-1/d} \mathcal{V}_n \triangle B \right) = 0.$$

Proof. From Proposition 3.1, for all $h \leq -1$,

$$\mathcal{V}_{n,RR,h-(2d-1)} \subseteq \mathcal{V}_{n,SP,h} \subseteq \mathcal{V}_{n,RR,h}, \quad (6)$$

and

$$\mathcal{V}_{n,RR,h-1} \subseteq \mathcal{V}_{n,DR,h} \subseteq \mathcal{V}_{n,RR,h}. \quad (7)$$

We will discuss the DR, h model first.

From Theorem 4.3, we know that the particle cluster of the RR, h model, scaled by $(\frac{n}{|h|})^{-1/d}$, tends for every fixed value of h to the unit volume sphere as $n \rightarrow \infty$, i.e.,

$$\lim_{n \rightarrow \infty} \lambda \left(\left(\frac{n}{|h|} \right)^{-1/d} \mathcal{V}_{n,RR,h} \triangle B \right) = 0,$$

and thus we also know (note that since h is negative, $|h-1| = |h|+1$)

$$\lim_{n \rightarrow \infty} \lambda \left(\left(\frac{n}{|h|} \right)^{-1/d} \mathcal{V}_{n,RR,h-1} \triangle B \right) = \lambda \left(\left[\left(\frac{|h|}{|h|+1} \right)^{1/d} B \right] \triangle B \right) = 1 - \frac{|h|}{|h|+1}.$$

Because of (7), it follows that

$$0 \leq \limsup_{n \rightarrow \infty} \lambda \left(\left(\frac{n}{|h|} \right)^{-1/d} \mathcal{V}_{n,DR,h} \triangle B \right) \leq 1 - \frac{|h|}{|h|+1},$$

so that

$$\lim_{h \rightarrow -\infty} \limsup_{n \rightarrow \infty} \lambda \left(\left(\frac{n}{|h|} \right)^{-1/d} \mathcal{V}_{n,DR,h} \triangle B \right) = 0.$$

The argument is similar for the SP, h model. \square

4.4 The toppling cluster for the sandpile model

In this subsection, we generalize Theorem 4 of Le Borgne and Rossin [2], who studied the particle cluster of the SP, h model for $d = 2$ and $h = 0, 1$ and 2 , to arbitrary d and h .

Theorem 4.8 Let $\mathfrak{D}(r)$ be the diamond $\bigcup_{\mathbf{x}: \sum_{i=1}^d |x_i| \leq r} \mathbf{x}^\square$, and $\mathfrak{C}(r)$ the cube $\bigcup_{\mathbf{x}: \max_i |x_i| \leq r} \mathbf{x}^\square$.

1. In the SP, h model, for every n and $h \leq 2d - 2$, there is an r_n such that

$$\mathfrak{D}(r_n - 1) \subseteq \mathcal{T}_n \subseteq \mathfrak{C}(r_n).$$

and

$$\mathfrak{D}(r_n) \subseteq \mathcal{V}_n \subseteq \mathfrak{C}(r_n + 1).$$

2. This r_n satisfies

$$\begin{aligned} \left(\frac{n}{2d-1-h} \right)^{1/d} - 3 &\leq 2r_n && \text{for all } n \text{ and } h \leq 2d - 2, \\ 2r_n &\leq \left(\frac{dn}{d-h} \right)^{1/d} + o(n^{1/d}) && \text{for } n \rightarrow \infty \text{ and } h < d. \end{aligned}$$

Corollary 4.9 $\mathcal{V}_{n,RR,h}$ contains the diamond $\mathfrak{D}\left(\frac{1}{2}\left(\frac{n}{2d-1-h}\right)^{1/d} - \frac{3}{2}\right)$ for all $h < 0$, and is contained in the cube $\mathfrak{C}\left(\frac{1}{2}\left(\frac{dn}{d-h}\right)^{1/d}\right)$ for all $h < -d$.

The corollary follows from combining the theorem with Proposition 3.1 part 1 and 3.

Remark 4.10 Note that the first inequality for r_n agrees with that in Theorem 4.1, for the sandpile model with $h = 2d - 2$. We can not use the second inequality in this case, because that is only valid for $h < d$. From combining Theorem 4.1 with Proposition 3.1 part 3, we find that for $-d \leq h < 0$, $\mathcal{V}_{n,RR,h}$ is contained in the cube $\mathfrak{C}(n + 1)$.

Proof of Theorem 4.8, part (1) We prove the inequality for \mathcal{T}_n . The inequality for \mathcal{V}_n then follows from (4).

We follow the method of Le Borgne and Rossin, who introduce the following stabilization procedure: Starting at $t = 0$ in the initial configuration with n particles at the origin, each time step every unstable site of the current configuration is stabilized. For example, to obtain η^1 we topple only the origin, as many times as legally possible. At $t = 2$, we topple only the neighbors of the origin, et cetera. We will call this the Le Borgne-Rossin procedure. Thus, η^t contains stable sites that toppled at time t , and unstable sites that received grains. T^t contains for every site the total number of topplings up to time t . For every site \mathbf{x} , we then have

$$T^{t+1}(\mathbf{x}) = \max\left\{\left\lfloor \frac{1}{2d} \left(H^0(\mathbf{x}) + \sum_{\mathbf{y} \sim \mathbf{x}} T^t(\mathbf{y}) \right) \right\rfloor, 0\right\}, \quad (8)$$

where $\mathbf{y} \sim \mathbf{x}$ means that \mathbf{y} is a neighbor of \mathbf{x} . We remark that this formula should correspond to formula (2) of [2], but, apparently due to a misprint, there the term $H^0(\mathbf{x})$ was omitted.

If we divide the grid \mathbb{Z}^d in two (“checked”) subgrids $\{\mathbf{x} : \sum_{i=1}^d x_i \text{ is even}\}$ and $\{\mathbf{x} : \sum_{i=1}^d x_i \text{ is odd}\}$, then at every time t , at most one subgrid contains unstable sites, because the neighbors of sites of one subgrid are all in the other subgrid, and we started with only one unstable site. We will consider next-nearest neighbor pairs \mathbf{x} and \mathbf{z} , i.e., \mathbf{x} and \mathbf{z} are in the same subgrid, at most 2 coordinates differ, and $\sum_i ||x_i| - |z_i|| = 2$, or in other words, \mathbf{z} is a neighbor of a neighbor of \mathbf{x} , \mathbf{x} itself being excluded.

First, we will prove by induction in t that for every next-nearest neighbor pair \mathbf{x} and \mathbf{z} , with $d(\mathbf{x}) \leq d(\mathbf{z})$, where $d(\mathbf{x})$ is the Euclidean distance of \mathbf{x} to the origin, $T^t(\mathbf{x}) \geq T^t(\mathbf{z})$. We remark that for a next-nearest neighbor pair \mathbf{x}' and \mathbf{z}' with $d(\mathbf{x}') = d(\mathbf{z}')$, we must have that \mathbf{x}' is equal to \mathbf{z}' up to a permutation of coordinates. The model should be invariant under such a permutation, therefore, in that case we have $T^t(\mathbf{x}') = T^t(\mathbf{z}')$.

The statement is true at $t = 0$, since at $t = 0$, we have $T^0(\mathbf{x}) = 0$ for all \mathbf{x} . It is also true at $t = 1$, since at $t = 1$, only the origin topples.

Now suppose that the statement is true at time t , and let \mathbf{x}_1 and \mathbf{x}_2 be two next-nearest neighbor sites that do not topple at time t . Suppose $d(\mathbf{x}_1) < d(\mathbf{x}_2)$. Then by assumption, $\sum_{\mathbf{y}_1 \sim \mathbf{x}_1} T^t(\mathbf{y}_1) \geq \sum_{\mathbf{y}_2 \sim \mathbf{x}_2} T^t(\mathbf{y}_2)$, because the neighbors of \mathbf{x}_1 and \mathbf{x}_2 can be grouped into next-nearest neighbor pairs with $d(\mathbf{y}_1) < d(\mathbf{y}_2)$. Furthermore, we always have $H^0(\mathbf{x}_1) \geq H^0(\mathbf{x}_2)$, since $H^0(\mathbf{0}) = n$, and $H^0(\mathbf{x}) = h$ if $\mathbf{x} \neq \mathbf{0}$.

Inserting this in (8), we obtain $T^{t+1}(\mathbf{x}_1) \geq T^{t+1}(\mathbf{x}_2)$, so that the statement remains true at $t + 1$ for one subgrid. The other subgrid must contain only sites that do not topple at $t + 1$, so that for all \mathbf{x} in this subgrid, $T^{t+1}(\mathbf{x}) = T^t(\mathbf{x})$. Therefore, the statement remains true at $t + 1$ for both subgrids.

Now let \mathbf{x} be a site with maximal $r = \max_i |x_i|$ where $T(\mathbf{x}) > 0$ (by symmetry, \mathbf{x} cannot be unique). Then no sites outside $\mathfrak{C}(r)$ can have toppled. Furthermore, all next-nearest neighbors of \mathbf{x}

that are closer to the origin have also toppled, and subsequently all next-nearest neighbors of those next-nearest neighbors, that are still closer to the origin, etc. Then we use symmetry to find that all sites in $\mathfrak{D}(r)$ and in the same subgrid as \mathbf{x} must also have toppled. An example for $d = 2$ is depicted in Figure 4.4.

·	·	z	·	z	·	·
·	y	·	z	·	z	·
x	·	y	·	z	·	z
·	y	·	0	·	z	·
z	·	z	·	z	·	z
·	z	·	z	·	z	·
·	·	z	·	z	·	·

Figure 4.4. Let x be a site with maximal $r = \max_i |x_i|$ that toppled. In this example, $r = 3$. By the fact that all next-nearest neighbors that are close to the origin have also toppled, we find that sites y , and the origin, have also toppled. By symmetry, we find that sites z have also toppled. Then we conclude that all sites in $\mathfrak{D}(3)$ (colored lightgrey) and in the same subgrid as x have toppled, and no sites outside $\mathfrak{C}(3)$ in the same subgrid can have toppled.

Furthermore, for at least one of the neighbors \mathbf{y} of \mathbf{x} , we must have $T(\mathbf{y}) > 0$, because \mathcal{T}_n is path connected. This neighbor has $\max_i |y_i| \geq r - 1$. For this neighbor, we can make the same observation. It follows that all sites in $\mathfrak{D}(r - 1)$ have toppled, so that the first part of the theorem follows. \square

To prove the inequalities for r_n , we need the following lemma:

Lemma 4.11 Let ρ_n be the average value of $H_n(\mathbf{x})$ in \mathcal{T}_n , σ_n the number of sites and β_n the number of internal bonds in \mathcal{T}_n . Then

$$\frac{\beta_n}{\sigma_n} \leq \rho_n \leq 2d - 1.$$

Proof. It suffices to prove that the configuration restricted to \mathcal{T}_n is recurrent [4, 18]. This can be checked in a nonambiguous way with the burning algorithm [4]. The inequality then follows from the fact that in a recurrent, stable configuration restricted to some set, the total number of particles in the set is at least the number of internal bonds in the set, and at most h_{max} times the number of sites in the set. In the case of an unstable configuration, the term “recurrent” is somewhat inappropriate, but we still define it to indicate configurations that pass the burning algorithm.

We use the following properties of recurrent configurations:

1. recurrence of a configuration restricted to a set is conserved under addition of a particle to the set, and under a legal toppling in the set.
2. if the configuration restricted to a set is recurrent, and we add an unstable site \mathbf{x} to the set, then the configuration is recurrent restricted to the extended set.

We will use induction in n . We will show that if η_n restricted to \mathcal{T}_n is recurrent, then η_{n+1} restricted to \mathcal{T}_{n+1} is recurrent. For a starting point of the induction, we choose $n' = 2d - h$, so that $|\mathcal{T}_{n'}| = 1$, because a configuration restricted to a single site is always recurrent.

We can obtain η_{n+1} by starting from η_n , adding a particle to the origin, and stabilizing. If we first stabilize restricted to \mathcal{T}_n , then by the first property of recurrence, the configuration restricted to

\mathcal{T}_n is still recurrent. If $\mathcal{T}_n = \mathcal{T}_{n+1}$, then we are now done. But it is also possible that at this point, there are unstable sites outside \mathcal{T}_n .

By both properties of recurrence, if we add one of these sites to \mathcal{T}_n and stabilize with respect to the new set, the configuration restricted to the new set is still recurrent. We can repeat this step until there are no more unstable sites to be added. Then the new set is \mathcal{T}_{n+1} , and the new configuration is η_{n+1} . \square

Proof of Theorem 4.8, part (2) To calculate the inequalities for r_n , we observe that the minimum r_n would be obtained if \mathcal{V}_n would equal the cube $\mathfrak{C}(r_n + 1)$, with each site containing the maximal number of particles. The number of sites in this cube is $(2r_n + 3)^d$, and the number of particles per site is then $2d - 1 - h$, so that

$$n \leq (2r_n + 3)^d(2d - 1 - h),$$

for every n .

The maximal r_n would be obtained if \mathcal{V}_n would equal the diamond $\mathfrak{D}(r_n)$, containing the minimum number of particles. As we deduced that the configuration restricted to \mathcal{T}_n is recurrent, we calculate the minimum average value of $H(\mathbf{x})$ on $\mathfrak{D}(r_n - 1)$, using Lemma 4.11. The number of internal bonds in $\mathfrak{D}(r_n - 1)$ is $(2r_n - 2)^d$. The number of sites in $\mathfrak{D}(r_n - 1)$ is $\frac{1}{d}(2r_n - 1)^d + o(r_n^d)$ as $r_n \rightarrow \infty$. Thus,

$$n \geq (2r_n - 2)^d - \frac{h}{d}(2r_n - 1)^d + o(r_n^d), \quad r_n \rightarrow \infty.$$

As in the limit $n \rightarrow \infty$ we also have $r_n \rightarrow \infty$, this leads to the second inequality for r_n in the theorem. \square

The proof of Theorem 4.8 allows to prove an interesting characteristic of \mathcal{V}_n :

Proposition 4.12 For the sandpile model, \mathcal{V}_n is simply connected for all n .

Proof. We first prove that \mathcal{T}_n is simply connected, by contradiction.

Suppose that $\mathcal{T}_{n,SP,h}$ is not simply connected, that is, $\mathcal{T}_{n,SP,h}$ contains holes. This means there must be a site \mathbf{y} in $\mathcal{T}_{n,SP,h}$ ‘beyond a hole’, that is, a site y with $T(\mathbf{y}) > 0$, and at least one neighbor \mathbf{x} of \mathbf{y} with $T(\mathbf{x}) = 0$, that is closer to the origin than y itself and all neighbors x' of y for which $T(x') > 0$. At least one neighbor \mathbf{x}' of \mathbf{y} must have $T(\mathbf{x}') > 0$, because \mathcal{T}_n is path connected. Thus, among the neighbors of \mathbf{y} , there must be a next-nearest neighbor pair \mathbf{z} and \mathbf{z}' , with $d(\mathbf{z}') \geq d(\mathbf{z})$, $T(\mathbf{z}) = 0$ and $T(\mathbf{z}') > 0$. But this contradicts the above derived property that for all t , for every next-nearest neighbor \mathbf{z} and \mathbf{z}' with $d(\mathbf{z}) \geq d(\mathbf{z}')$, $T^t(\mathbf{z}) \leq T^t(\mathbf{z}')$, as this property should also hold for the final configuration.

The above does not suffice to conclude that \mathcal{V}_n is also simply connected. Since $\mathcal{V}_n = \mathcal{T}_n \cup \delta\mathcal{T}_n$, we must show that \mathcal{T}_n does not contain so-called *fjords*, i.e., places where the boundary of \mathcal{T}_n nearly touches itself, such that $\mathcal{T}_n \cup \delta\mathcal{T}_n$ would contain a hole. But when one supposes that \mathcal{T}_n does contain a fjord, then one can derive the same contradiction as above. Therefore, \mathcal{V}_n is simply connected. \square

References

- [1] P. BAK, K. TANG AND K. WIESENFELD (1988) Self-organized criticality, *Phys. Rev. A* **38**, 364-374.
- [2] Y. le Borgne and D. Rossin (2002) On the identity of the sandpile group, *Discrete Math.* **256**, 775-790.
- [3] A. Dartois and D. Rossin (2004) Height Arrow Model, *Formal Power Series and Algebraic Combinatorics, 16th conference, Vancouver*.

- [4] D. Dhar (1990) Self-Organized Critical State of Sandpile Automaton Models, *Phys. Rev. Lett.* **64**(14), 1613-1616.
- [5] P. Diaconis and W. Fulton (1991) A growth model a game, an algebra, Lagrange inversion, and characteristic classes, *Rend. Sem. Mat. Univ. Pol. Torino* **49**, 95-119.
- [6] R. Durrett (1988) Lecture notes on particle systems and percolation. *Wadsworth and Brooks/Cole, Pacific Grove, Calif.*
- [7] E.V. Ivashkevich and V.B. Priezzhev(1998) Introduction to the sandpile model, *Physica A* **254**, 97-116.
- [8] M. Kleber (2005) Goldbug variations, *Math. Intelligencer* **27** (1), 55-63.
- [9] G.F. Lawler, M. Bramson and D. Griffeath (1992) Internal diffusion limited aggregation, *Ann. Probab.* **20**(4), 2117-2140.
- [10] G.F. Lawler (1995) Subdiffusive fluctuations for internal diffusion limited aggregation, *Ann. Probab.* **23**(1), 71-86.
- [11] L. Levine (2002) The rotor-router model, *Harvard University senior thesis.*
- [12] L. Levine, Y. Peres (2005) Spherical Asymptotics for the Rotor-Router Model in \mathbb{Z}^d , to appear in *Indiana University Math Journal.*
- [13] L. Levine, Y. Peres (2007) Strong Spherical Asymptotics for Rotor-Router Aggregation and the Divisible Sandpile, Preprint available at <http://www.arxiv.org/abs/0704.0688>.
- [14] S.H. Liu, T. Kaplan and L.J. Gray (1990) Geometry and dynamics of deterministic sand piles, *Phys. Rev. A* **42**(6), 3207-3212.
- [15] S. Lübeck, N. Rajewsky and D.E. Wolf (2000) A deterministic sandpile revisited, *Eur. Phys. Journ. B* **13**, 715-721.
- [16] S.S. Manna (1991) Two-state model of self-organized criticality, *J. Phys. A: Math. Gen.* **24**, L363-369.
- [17] M. Markořová and P. Markoř (1992) Analytic calculation of the attractor periods of deterministic sandpiles, *Phys. Rev. A* **46**(6), 3531-3534.
- [18] R. Meester, F. Redig and D. Znamenski (2001) The abelian sandpile model, a mathematical introduction, *Markov Proc. Rel. Fields* **7**, 509-523.
- [19] A.M. Povolotsky, V.B. Priezzhev and R.R. Shcherbakov (1998) Dynamics of Eulerian walkers, *Phys Rev. E* **58**, 5449-5454.
- [20] K. Wiesenfeld, J. Theiler and B. McNamara(1990) Self-organized Criticality in a Deterministic Automaton, *Phys. Rev. A* **65**(8), 949-952.

This figure "mannaD.gif" is available in "gif" format from:

<http://arxiv.org/ps/math/0702450v2>

This figure "mannashapes.gif" is available in "gif" format from:

<http://arxiv.org/ps/math/0702450v2>

This figure "sandpilesshapesanders.gif" is available in "gif" format from:

<http://arxiv.org/ps/math/0702450v2>