Abelian Sandpiles and Self-Organized Criticality

Math 336 Final Paper

Mark Polyakov

June 15, 2021

Abstract

The abelian sandpile is a misleadingly simple construction with wide implications for real-world phenomena. We allow sand to reside at one of many sites. When too many grains of sand are added to one site, it topples over, spilling sand onto adjacent sites, which may in turn topple. In this paper we will explore the most essential facts about sandpiles, many of which are not immediate nor intuitive. We focus on facts related to the recurrent states, which are those configurations of grains that can always be reached by adding more grains to any existing sandpile. The structure of these recurrent states is surprisingly regular. Later, we consider Dhar's directed sandpile model and determine that it exhibits self-organized criticality, which, roughly speaking, is the property that adding a single grain causes very large avalanches of topplings to occur surprisingly frequently. Self-organized criticality is perhaps the most important feature of sandpiles with regards to applications, and at the end of the paper we will discuss how failures of an electric grid might be modeled in terms of a sandpile, and what might be implied about those electric grids by the comparison. In addition to facts about sandpiles, we prove some more general facts, one related to quotient groups of integer vectors and several related to random walks, which are usually taken for granted in sandpile literature.

1 Definition and Basic Properties of the Abelian Sandpile

A suitable precursor to the formal definition of a sandpile is a definition of the 2-dimensional grid/lattice sandpile. The present state of the sandpile is a $V \times V$ grid. Associated with each point on the grid is a number from 0 to 3 indicating how many grains of sand are on that point. We may add grains of sand to any point. When we add a grain of sand to a point with height 3, instead of incrementing the height to 4, we *topple* the site by resetting its height to 0 and adding a grain of sand to each neighboring point. If some of the neighbors were at height 3, additional topplings may be triggered, in what is called an *avalanche*. When a site at the edge of the grid is toppled, the grains that would normally be added to its neighbors are instead discarded, and thus it is possible for the number of grains in the system to decrease as the result of adding a grain somewhere. Now, we move on to the general definition of a sandpile, where we allow sites to be linked together arbitrarily and even asymmetrically.

Now we define sandpiles in full generality. The definition we will use is most similar to the one in [6]. Say there are V nodes/sites that sand can reside on. The state or configuration of the sandpile is a vector in $\mathbb{Z}^V$. If $\mathcal{S} \in \mathbb{Z}^V$ is a state vector, then $\mathcal{S}(i)$ denotes the i-th entry in that vector. To describe how the sandpile behaves we define a matrix Δ of size $V \times V$. A toppling at site i is denoted by T_i; when this occurs, the i-th row of Δ is subtracted from the state vector.

Definition 1.1. *The i-th site of $\mathcal{S}$ is* unstable *if $\mathcal{S}(i) \geq \Delta_{ii}$. A state as a whole is* unstable *if it contains any unstable sites.*

Definition 1.2. *A* toppling at site i*, denoted by T_i, is performed on a state $\mathcal{S}$ like so: For every node k in $\mathcal{S}$ (including $k = i$), subtract Δ_{ik} from the number of grains at that site.*

Definition 1.3. *A* single toppling *at no specific site, which I will denote as T' with no subscript, is an operation where all unstable sites are toppled once.*

For example, if sites i and j are unstable, then $T' = T_i T_j$. Since toppling at a single site is simply vector addition (adding the state and a row of Δ together), the ordering of the individual site topplings is immaterial, so T' is well-defined.

Definition 1.4. *A* toppling *at no specific site, which I will denote as T with no subscript, is a multi-step operation wherein T' is applied repeatedly until the state is stable.*

With only the facts presented thus far, it is clear that T will terminate, so the action of T is not yet well-defined. However, when it does terminate, its action is well-defined (i.e., T is single-valued). The main fact about T is that it maps all states to stable states.

We also consider A_i, the operation of adding a grain of sand to site i, and then, most importantly, the operation σ_i, which is equal to TA_i (adding followed by toppling).

It is helpful to think of the sandpile as a directed graph where the i-th node is annotated with Δ_{ii} and the edge between the i-th and j-th nodes, if such an edge exists, is annotated with $-\Delta_{ij}$; when a site topples, it sends grains to each node it has an outgoing edge to, according to the value that edge is annotated with. In the restricted set of sandpiles where, for all i and j, it is true that $\Delta_{ij} = \Delta_{ji}$ (the matrix is symmetric), $\Delta_{ij} \in (0, 1)$ (edges, when they exist, are always annotated with 1), the graph can be considered undirected and Δ is almost equal to the Laplacian matrix of the graph. However, we will have $\Delta_{ii} \neq -\sum_j \Delta_{ij}$, for some i, for reasons to be explained below.

We will now introduce restrictions on the matrix Δ which will ensure that T is well-defined, specifically, that T always terminates.

Definition 1.5. *A* dissipative site *is a site i such that $\sum_j \Delta_{ij} > 0$.*

When a toppling occurs at a dissipative site, the total number of grains of sand in the state decreases. Dissipative sites are necessary for T to terminate. Additionally, by counting topplings at dissipative sites, one can measure the number of grains that leave the sandpile, which is one way to quantify the size of an avalanche.

Proposition 1.6. *Under the following restrictions on the Δ matrix, the operation T always terminates:*

1. *$\Delta_{ii} > 0$ for all i. I.e., no site topples when empty.*
2. *$\Delta_{ij} \leq 0$ for all i and j. I.e., toppling never steals grains from neighbors.*
3. *$\sum_j \Delta_{ij} \geq 0$ for all i. I.e., toppling never results in a net increase in grains.*

4. *For every j, there is a list $x_1, \ldots, x_n$ with $x_1 = j$ and $x_n = i$, where i is a dissipative site and $\Delta_{x_k, x_{k+1}} \neq 0$ for all $1 \leq k < n$. I.e., the directed graph representing the sandpile has a path from any given site to a dissipative site.*

Proof. Suppose T does not terminate. Then there must be at least one site i that topples infinitely many times. Let M be the distance (number of nodes in the shortest directed path) from i to a dissipative site, which exists by restriction 4. Site i must have a neighbor j with distance $M-1$ to a dissipative site. After some number of topplings at i, each neighbor of i will overflow and topple as well; restriction 2 ensures there is no way to prevent this. Thus j must topple an infinite number of times as well. Repeating inductively, the dissipative site node i is linked to must topple infinitely many times.

By restriction 3, none of the toppling can increase the total number of grains in the state. However, the infinite toppling at the dissipative site will decrease the number of grains in the system. The definition of toppling prevents forces sites to have a nonnegative number of grains, so eventually every site will have zero grains. Then, by restriction 1, the toppling terminates, a contradiction.

QED

In light of this result, we will henceforth only consider sandpiles where the conditions set out in the statement of the proposition are satisfied.

Notice that the restrictions do not force the graph to be connected. Some important results in the field require connectivity, though not the ones in this paper.

Theorem 1.7. *If $T_{x_n} \cdots T_{x_1} \mathcal{S} = T\mathcal{S}$ and $T_{y_m} \cdots T_{y_1} \mathcal{S} = T\mathcal{S}$, and at no point during either list of operations is a stable site toppled, then $(x_1, \ldots, x_n)$ is a permutation of $(y_1, \ldots, y_m)$.*

Proof. It is true that site x_1 is unstable in $\mathcal{S}$, else it would not have been toppled first. Thus, $x_1 \in (y_1, \ldots, y_m)$, since the unstable site can only be made stable by toppling it at some point. Commuting T_{x_1} to the front, we get:

$$T_{y_m} \cdots T_{y_{k+1}} T_{y_{k-1}} \cdots T_{y_1} T_{x_1} \mathcal{S} = T\mathcal{S}.$$

The property that no stable site shall be toppled is maintained. We can proceed inductively on the remaining y_i (for example, x_2 must be unstable in $T_{x_1}\mathcal{S}$, so $x_2 \in (y_1, \ldots, y_{k-1}, y_{k+1}, \ldots, y_m)$, and we commute it to the front, etc). This shows that there is a one-to-one map from the x_i to equal y_i. Relabelling x and y then repeating the argument, there's also a one-to-one map from the y_i to equal x_i. Thus there exists a one-to-one correspondence between the two lists, which in this finite case is a permutation. QED

The following corollary justifies the term "Abelian Sandpile".

Corollary 1.7.1 (Abelianness)**.** *From any starting state $\mathcal{S}$, it is true that $\sigma_j \sigma_i \mathcal{S} = \sigma_i \sigma_j \mathcal{S}$.*

Proof. Choose some $(x_1, \ldots, x_n)$, $(y_1, \ldots, y_m)$, k, and q such that

$$\begin{aligned} \sigma_j \sigma_i \mathcal{S} &= T_{x_n} \cdots T_{x_{k+1}} A_j T_{x_k} \cdots T_{x_1} A_i \mathcal{S} \\ \sigma_i \sigma_j \mathcal{S} &= T_{y_m} \cdots T_{y_{q+1}} A_i T_{y_q} \cdots T_{y_1} A_j \mathcal{S}. \end{aligned} \tag{1}$$

We can commute the terms around, since they are all simply vector addition, yielding

$$\begin{aligned} \sigma_j \sigma_i \mathcal{S} &= T_{x_n} \cdots T_{x_1} A_j A_i \mathcal{S} \\ \sigma_i \sigma_j \mathcal{S} &= T_{y_m} \cdots T_{y_1} A_j A_i \mathcal{S}. \end{aligned} \tag{2}$$

A stable site is never toppled in (2); after p many topplings have been applied in one of the equations in (2), the state is the same as after p many topplings were applied in the corresponding equation of (1), except for an extra grain at site i or j, and thus all sites which were unstable at that point in (1) are still unstable, so if the toppling originally toppled an unstable site, it still does. Therefore it is valid to apply Theorem 1.7, which completes the proof. QED

2 Recurrent and Transient States

In this section we will explore the properties of the recurrent states, which are always reachable and therefore of great interest when analyzing the long-term behavior of the sandpile. The proofs in this section are mainly inspired by those in [6] and [7], with some inspiration from [5] as well.

2.1 Essential Facts

Definition 2.1. *A state $\mathcal{R}$ is* recurrent *if for any list $x_1, \ldots, x_n$, there is another list $y_1, \ldots, y_m$ such that $(\sigma_{y_n} \cdots \sigma_{y_1})(\sigma_{x_n} \cdots \sigma_{x_1})\mathcal{R} = \mathcal{R}$. A state that is not recurrent is* transient*. Define $\mathbf{R}$ as the set of all recurrent states.*

The maximal state where every site is "fully loaded" is certainly recurrent. As it happens, there are typically plenty of transient states too.

With what we've seen thus far, it may seem reasonable that the set of recurrent states is disconnected in the sense that one could break the set of recurrent states into separate parts, with no possible list of σ_ns that causes a state in one part to become a state in another part. However, this is not true; the recurrent states are globally accessible.

Proposition 2.2. *Given any state $\mathcal{S}$ and recurrent state $\mathcal{R}$, there is a list $x_1, \ldots, x_n$ such that $(\sigma_{x_n} \cdots \sigma_{x_1})\mathcal{S} = \mathcal{R}$.*

Proof. We can define the maximal state $\mathcal{M}$ where $\mathcal{M}(i) = \Delta_{ii} - 1$ for all i. This is the state where adding a grain anywhere will cause the state to become unstable. $\mathcal{M}$ is recurrent because from any state $(\sigma_n \cdots \sigma_1)\mathcal{M}$, we can reach $\mathcal{M}$ by adding grains to each site i until that site has $\Delta_{ii} - 1$ grains. Now suppose we are given an arbitrary state $\mathcal{S}$ and arbitrary recurrent state $\mathcal{R}$. Choose lists of x and y such that $(x_n \cdots x_1)\mathcal{S} = \mathcal{M}$ and $(y_m \cdots y_1)\mathcal{R} = M$. Since $\mathcal{R}$ is recurrent, there is another list $(z_q \cdots z_1)\mathcal{M} = \mathcal{R}$, and then $(z_q \cdots z_1)(x_n \cdots x_1)\mathcal{S} = \mathcal{R}$ as desired. QED

We now consider how recurrence interacts with individual σ_i operators, rather than all possible combinations of them.

Definition 2.3. *$\mathcal{R}$ is* recurrent with respect to i with period n *if $\sigma_i^n \mathcal{R} = \mathcal{R}$, and n is the lowest number satisfying this equation.*

Lemma 2.4. *If there are two lists $(x_1, \ldots, x_n)$ and $(y_1, \ldots, y_m)$ such that $(\sigma_{x_n} \cdots \sigma_{x_1})\mathcal{R}_1 = (\sigma_{y_m} \cdots \sigma_{y_1})\mathcal{R}_1$ for some recurrent state $\mathcal{R}_1$, then $(\sigma_{x_n} \cdots \sigma_{x_1})\mathcal{R}_2 = (\sigma_{y_m} \cdots \sigma_{y_1})\mathcal{R}_2$ for any recurrent state $\mathcal{R}_2$.*

Proof. Since the states are recurrent, there must exist a list $(z_1, \ldots, z_k)$ such that $(z_k \cdots z_1)\mathcal{R}_1 = \mathcal{R}_2$. Then

$$(x_n \cdots x_1)\mathcal{R}_2 = (x_n \cdots x_1)(z_k \cdots z_1)\mathcal{R}_1 = (y_m \cdots y_1)(z_k \cdots z_1)\mathcal{R}_1 = (y_m \cdots y_1)\mathcal{R}_2.$$

QED

Theorem 2.5.

1. *A state $\mathcal{R}$ is recurrent if and only if $\mathcal{R}$ is recurrent with respect to all i.*
2. *If recurrent state $\mathcal{R}_1$ is recurrent for i with period n, then any other recurrent state $\mathcal{R}_2$ is also recurrent for i with period n.*

Proof. First I will prove statement 1. Suppose we are given a state $\mathcal{T}$ that is not recurrent with respect to i. Since there are a finite number of states, the sequence $\{\sigma_i^j \mathcal{T}\}_{j=0}^{\infty}$ contains duplicates, i.e., there exist positive integers m and k such that $\sigma_i^k \mathcal{T} = \sigma_i^{k+m} \mathcal{T} = \sigma_i^m \sigma_i^k \mathcal{T}$. Thus the state $\sigma_i^k \mathcal{T}$ is recurrent with respect to i. We cannot bring a state that is recurrent for i to a state that is not recurrent for i using σs, so there is no list of σs that will bring σ_i^k back to $\mathcal{T}$, and therefore $\mathcal{T}$ is transient. Since i is arbitrary, the "only if" direction of the statement is proved.

Now suppose we are given a state $\mathcal{R}$ that is periodic for all i with periods n_i for each. My goal is to show that for any list $\sigma_{x_1}, \ldots, \sigma_{x_m}$, there is another list of σs that will bring $(\sigma_{x_m} \cdots \sigma_{x_1})\mathcal{R}$ back to $\mathcal{R}$. Indeed, the desired list is given by $\sigma_{x_1}^{n_1 - 1}, \ldots, \sigma_{x_m}^{n_m - 1}$, since

$$(\sigma_{x_m}^{n_m - 1}, \ldots, \sigma_{x_1}^{n_1 - 1})(\sigma_{x_m} \cdots \sigma_{x_1})\mathcal{R} = \sigma_{x_m}^{n_m} \cdots \sigma_{x_1}^{n_1} \mathcal{R} = \mathcal{R}.$$

This proves statement 1.

Statement 2. follows from Lemma 2.4 with the first operation being σ_i^n and the second operation being the empty/identity operation. QED

The preceding theorem implies that the period of recurrence for a certain site, among the recurrent states, is well-defined. I will use $\text{per}(i)$ to denote the period of i for recurrent states. In light of these two propositions, [2] describes the structure of the set of recurrent states as a multidimensional torus.

It becomes desirable to define σ_i^{-n} for $n \in \mathbb{N}$ (when $n = 0$, it adds no grains, so has no effect). We only define this action on recurrent states. Then $\sigma_i^{-1} = \sigma_i^{\text{per}(i)-1}$. Generalizing, $\sigma_i^{-n} = \sigma_i^{\text{per}(i)-n}$. (If this last $\text{per}(i) - n$ is negative, we can simply add a multiple of $\text{per}(i)$ until it is positive, which is equivalent to recursively expanding $\sigma_i^{\text{per}(i)-n}$ until the exponent is positive).

In light of the last paragraph, an inverse operation for any list of σs exists.

2.2 The Sandpile Group

Definition 2.6. *The group G (we shall soon see it is a group) on $\mathbf{R}$ consists of all functions $\mathbf{R} \to \mathbf{R}$ that can be realized by a list of σ operations. The group operation is function composition.*

To better understand this definition, recall Lemma 2.4, which states that any two operations made out of list of σ which are equivalent with respect to one recurrent state are equivalent with respect to all recurrent states. Thus, an $f \in G$ which is realized by some list of σs on a recurrent state $\mathcal{R}_1$ is realized by the same list of σs on all recurrent states.

Proposition 2.7. *The group G as defined is in fact a group.*

Proof. Function composition is associative.

The composition of any two elements of G can be realized as the concatenation of two lists of σs, which yields a list of σs, so G is closed.

The identity element, henceforth referred to as e, is realized by the empty operation where no σs are applied (this operation can be realized by nontrivial lists of σs as well).

Inverses exist by the definition of a recurrent state (and I have previously described how to realize them). QED

Lemma 2.8. *For any i, then $\sigma_V^{\Delta_{i,V}} \cdots \sigma_1^{\Delta_{i,1}}$ is a way to realize e (the identity element of G).*

Proof. The operation $\sigma_i^{\Delta_{i,i}}$ will appear in $\sigma_V^{\Delta_{i,V}} \cdots \sigma_1^{\Delta_{i,1}}$. Applying $\sigma_i^{\Delta_{i,i}}$ will always cause exactly one toppling at site i, and will ultimately leave site i at the same height. Thus, the only net effect of $\sigma_i^{\Delta_{i,i}}$ is its effect on neighboring sites, specifically, $\sigma_j^{\Delta_{i,j}}$ for all neighboring sites j. But $\sigma_V^{\Delta_{i,V}} \cdots \sigma_1^{\Delta_{i,1}}$ contains exactly the inverse of all those operations, and therefore acts as identity on recurrent states. QED

Lemma 2.9. *The group G has the same cardinality as $\mathbf{R}$. I.e., $|G| = |\mathbf{R}|$.*

Proof. Fix some $\mathcal{R}_0 \in \mathbf{R}$. Define $\mathbf{R}' = \{g\mathcal{R}_0 : g \in G\}$. See that $|\mathbf{R}'| = |G|$; certainly $|\mathbf{R}'| \leq |G|$, and the only way for $|\mathbf{R}'| < |G|$ to hold is if $g_1\mathcal{R}_0 = g_2\mathcal{R}_0$ for some $g_1, g_2 \in G$. But then g_1 and g_2, having the same effect on some recurrent state, would be the same function (see Lemma 2.4). Thus, our task is reduced to showing $|\mathbf{R}'| = |\mathbf{R}|$.

First, since all recurrent states are accessible, there is certainly a list of operations from $\mathcal{R}_0$ to any other recurrent state, and the function performing that list of operations is in G, and therefore $\mathbf{R}' \supset \mathbf{R}$.

As for $\mathbf{R}' \subset \mathbf{R}$, we argue by contradiction. Suppose there is a state $\mathcal{S} \in \mathbf{R}' \setminus \mathbf{R}$. Then there must be a $g \in G$ such that $\mathcal{S} = g\mathcal{R}_0$. But that g is realizable via σs, and no list of σs can bring a recurrent state to a non-recurrent state, so $\mathcal{S} \in \mathbf{R}$, contradicting our assumption. QED

Lemma 2.10. *The group G is isomorphic to the quotient group $\mathbb{Z}^V/\Delta^\top\mathbb{Z}^V$ (under vector addition), where $\Delta^\top\mathbb{Z}^V = \{\Delta^\top v : v \in \mathbb{Z}^V\}$.*

Proof. First, we'll construct a surjective homomorphism from $\mathbb{Z}^V \to G$. Then we'll know G is isomorphic to the quotient of $\mathbb{Z}^V$ and the preimage of e (the identity element of G) under the homomorphism.

The homomorphism is quite simple: Let $\phi(z) : \mathbb{Z}^V \to G$ simply map to the function in G that is realized by $\sigma_V^{z_V} \cdots \sigma_1^{z_1}$, where z_i is the i-th element of the vector z. The map ϕ is surjective, since every element of G is realizable. Furthermore, with the group operation on $\mathbb{Z}^V$ being vector addition, the map ϕ is a homomorphism:

$$\phi(z+w) = \sigma_V^{z_V+w_V} \cdots \sigma_1^{z_1+w_1} = \sigma_V^{z_V} \cdots \sigma_V^{z_1}\sigma_V^{w_V} \cdots \sigma_1^{w_1} = \phi(z)\phi(w).$$

It is well-known in algebra that, if one has a surjective homomorphism $f : A \to B$ and $fC = e$ (the identity in B), then $A/C \simeq B$ (where $\simeq$ denotes isomorphism). Thus, to proceed we seek out the preimage of e under ϕ, which I will call E. I claim $E = \Delta^\top\mathbb{Z}^V$.

First I will show $E \supset \Delta^\top\mathbb{Z}^V$. Take an arbitrary $v \in \Delta^\top\mathbb{Z}^V$. Denote the i-th row of Δ as Δ_i. Then $v = \Delta_{x_1} + \cdots + \Delta_{x_n}$ for an appropriate vector x. Since ϕ is a homomorphism, then $\phi(v) = \phi(\Delta_{x_1}) \cdots \phi(\Delta_{x_n})$. Each of these is identity per Lemma 2.8, and therefore $\phi(v) = e$ and $v \in E$.

Next, I show $E \subset \Delta^\top \mathbb{Z}^V$. Take an arbitrary element $z \in E$. We know $z \in \mathbb{Z}^V$. We can decompose z as $z^+ - z^-$ with $z_i^+ \geq 0$ and $z_i^- \geq 0$ for all i. Since the operation $\phi(z) = e$, then $\phi(z^+)\phi(-z^-) = e$ and $\phi(z^+) = \phi(-z^-)^{-1}$. The elements of $-z^-$ are the exponents of σs in $\phi(-z^-)^{-1}$, so we have $\phi(-z^-)^{-1} = \phi(z^-)$, since inverting the exponents on every σ in the realization of the function has the effect of inverting the function. Therefore, for any recurrent state $\mathcal{R}_0$, we have $\phi(z^+)\mathcal{R}_0 = \phi(z^-)\mathcal{R}_0$. Using the definition of ϕ and then equation (2), which indicates that the additions and topplings involved in a list of σ can be separated,

$$\phi(z^+)\mathcal{R}_0 = \mathcal{R}_0 + z^+ - \Delta^\top v$$

$$\phi(z^-)\mathcal{R}_0 = \mathcal{R}_0 + z^- - \Delta^\top w$$

For the appropriate vectors v and w, which record how many times each site toppled during application of $\phi(z^+)$ or $\phi(z^-)$ respectively. But the two right expressions are equal, so $z^+ - \Delta^\top v = z^- - \Delta^\top w$, and rearranging further, $z^+ - z^- = \Delta^\top(v - w)$, i.e., $z^+ - z^-$ is in the span of the rows of Δ and thus $z \in E$. Since z was arbitrary, we conclude $E \subset \Delta^\top \mathbb{Z}^V$. QED

The following lemma is fairly well-known, and is usually proved using the Smith Normal Form of a matrix. However, as suggested briefly in [7], it is possible to prove the lemma using almost no facts about matrices, as we will do here.

Lemma 2.11. *For any $N \times N$ nonsingular integer matrix A, the cardinality of the quotient group $\mathbb{Z}^N / A\mathbb{Z}^N$ is equal to $|\det A|$.*

If A is singular, then $|\mathbb{Z}^N / A\mathbb{Z}^N| = \infty$.

Proof. If A is singular, then $A^\perp$ is nonempty. We can choose infinitely many distinct cosets of $A\mathbb{Z}^N$ along any vector in $A^\perp$, proving the special case of the lemma.

The general idea of the proof is to take a large hypercube in $\mathbb{Z}^N$ and tile it as full as possible with copies of the parallelepiped formed by the columns of A, then notice that the sum of the volumes of the tiled parallelepipeds approaches the volume of the cube as the cube becomes large.

Let P_p be the parallelepiped formed from the columns of A and with its "bottom-left" point being p, which is a point in $\mathbb{Z}^N$; specifically, from p, we form P_p by creating edges from p where each edge, as a vector, is a column of A (the remaining edges and points are well-defined). A point $x \in \mathbb{Z}^N$ *belongs* to a P_p if x is in the interior of P_p or if x lies only on side(s) of P_p that is directly connected to p (i.e., if x lies on the "bottom left" of P_p). Figure 1 may clarify the definition of belonging.

Define the *lattice volume* of P_p, also latticeVolume(P_p), to be the number of points of $\mathbb{Z}^N$ belonging to P_P. It is quite intuitive that latticeVolume$(P_p) = |\mathbb{Z}^N / A\mathbb{Z}^N|$; it is possible to reach any point in $\mathbb{Z}^N$ by adding some linear combination of the columns of A to some point in P_p, and there is no nontrivial linear combination of columns of A which when added to one point belonging to P_p yields a different point also belonging to P_p. I take it for granted that volume$(P_p) = |\det A|$. Let

$$\Phi = \{P_{Av} : v \in \mathbb{Z}^N\}.$$

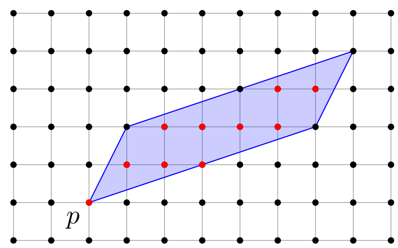

Figure 1: The red points belong to P_p. Note there are 10 red points, equal to the area of the parallelogram. Showing these values are always equal is the main challenge of the proof.

(the Av denotes matrix multiplication between A and v). I.e., Φ is the set of parallelepipeds that start at points in the column space of A. Notice that the elements of Φ cover $\mathbb{Z}^N$ and have non-overlapping interiors; i.e., Φ tiles $\mathbb{Z}^N$.

Let C_n be the cube $[0,n]^N$. Define lattice size for C_n in the same way as for the parallelepipeds, except using the origin as the point that determines which sides are "bottom-left" sides of C_n. Notice that the volume of C_n is equal to its lattice size, both n^N. This connection enables the rest of the proof. Let Φ_n be the subset of Φ such that all the parallelepipeds in Φ_n are in C_n (touching on the boundary is allowed). One more definition: Let m be a constant large enough so that any of the parallelepipeds could fit inside a cube with side length m.

Once $n > 2m$, then the hypercube $[m, n-m]^N$ will be entirely covered by Φ_n. This is because any parallelepiped in $[m, n-m]^N$ will have additional parallelepipeds attached to all of its sides, since there is certainly enough room left in C_n for such parallelepipeds. Figure 2 may clarify the geometry of C_n. Therefore,

$$\text{volume}(C_n) - \sum_{P\in\Phi_n} \text{volume}(P) \leq n^N - (n-2m)^N$$

$$\text{latticeVolume}(C_n) - \sum_{P\in\Phi_n} \text{latticeVolume}(P) \leq n^N - (n-2m)^N$$

The largest powers of $n^N - (n-2m)^N$ will cancel out, so $n^N - (n-2m)^N \in O(n^{N-1})$ (with n being the variable). Then, subtracting the above equations, we get

$$\sum_{P\in\Phi_n} \text{volume}(P) - \sum_{P\in\Phi_n} \text{latticeVolume}(P) \in O(n^{N-1}).$$

We can rewrite as

$$|\Phi_n|(|\det A| - |\mathbb{Z}^N/A\mathbb{Z}^N|) \in O(n^{N-1})$$

But $|\Phi_n|$ grows as n^N, the volume of C_n, so we conclude that

$$\lim_{n\to\infty} |\det A| - |\mathbb{Z}^N/A\mathbb{Z}^N| = 0$$

But the expression does not depend on n, so is simply zero, i.e., $|\mathbb{Z}^N/A\mathbb{Z}^N| = |\det A|$ as claimed. QED

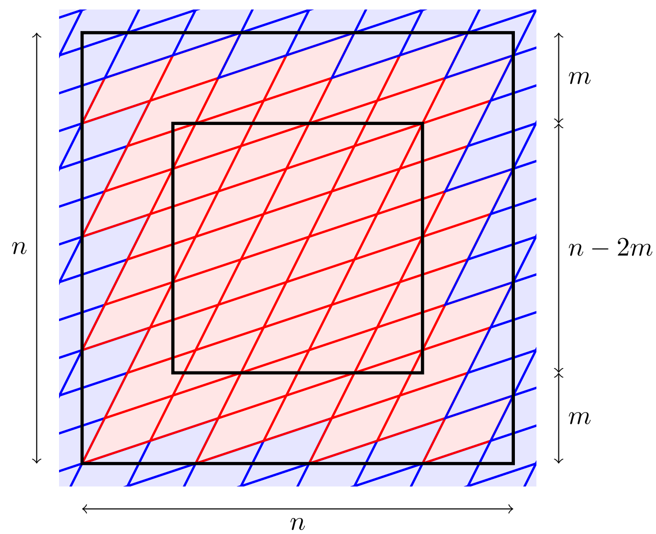

Figure 2: Geometry of the proof. The outer square is C_n. Notice that the central $(n - 2m) \times (n - 2m)$ square is covered by red parallelograms. As $n \to \infty$, the area in the margin becomes negligible as a fraction of the total area of C_n, and thus the difference between geometric area and lattice area *per parallelogram* approaches zero.

Theorem 2.12. *The number of distinct recurrent states is equal to* $|\det \Delta|$.

Proof. There is a one-to-one correspondence between recurrent states and elements of $\mathbb{Z}^V/\Delta^T\mathbb{Z}^V$, and the latter has cardinality $|\det \Delta|$. QED

We have arrived at this result by showing that the cardinality of $\mathbb{Z}^V/\Delta^\top\mathbb{Z}^V$ and the cardinality of $\mathbf{R}$ are equal, but have not shown an isomorphism between the two. An isomorphism does in fact exist. Specifically, $\mathbb{Z}^V/\Delta^\top\mathbb{Z}^V$ is isomorphic to $\mathbf{R}$ with the group operation on the latter being $\mathcal{R}_1 \circ \mathcal{R}_2 = T(\mathcal{R}_1 + \mathcal{R}_2)$. Furthermore, each element of $\mathbb{Z}^V/\Delta^\top\mathbb{Z}^V$ contains exactly one recurrent state, the one it maps to under the isomorphism. We will not prove these facts.

As an alternative to the way we developed Theorem 2.12, it is possible to prove that each coset in $\mathbb{Z}^V/\Delta^\top\mathbb{Z}^V$ contains exactly one recurrent state by directly analyzing the sandpile dynamics, as performed in [5].

Corollary 2.12.1. *The matrix* Δ *is nonsingular.*

3 Self-Organized Criticality and the Directed Sandpile Model

3.1 Preliminaries

One of the main motivations for studying abelian sandpiles is that they exhibit self-organized criticality. I will not attempt a thorough definition of self-organized criticality here. One attribute of critical systems is that the expected value of the size of events that result from random impulses diverges. In the case of sandpiles, a possible interpretation is that the expected number of sites toppled due to the addition of a grain of sand should diverge.

In a finite sandpile, the expected value of any metric of avalanche size will converges, since the sum that defines the expected value will be finite. However, as the size of the sandpile becomes large, so does the expected value (for most reasonable metrics of avalanche size). It is also possible to formalize an infinite sandpile where the expected value diverges more naturally.

The most common probability distributions which exhibit self-organized criticality feature powers of x between -1 and -2. In this range, the expression for the expected value will feature a sum of powers of x between 0 and -1, which will diverge.

Sandpiles are the canonical example of a system that exhibits self-organized criticality, so it may come as a surprise that the exact coefficients of the power laws describing avalanche sizes are still unknown. Additionally, not all sandpiles exhibit criticality. Researchers typically study sandpiles on lattices with dimension from 2 to 4 for this reason.

Shortly after sandpiles were formalized, a restricted version of the sandpile, the *directed sandpile*, was developed by Dhar in [3]. In the directed model, the power law exponents are much easier to calculate. This section is based on on ideas from [3] and [2], also by Dhar.

Definition 3.1. *The 2-dimensional directed sandpile is a restricted version of the sandpile. The nodes of the graph of the sandpile are* (x, y) *where* $0 \leq x + y \leq N$ *and* $0 \leq y \leq N$*. Every node* (x, y) *has directed edges to* $(x + 1, y)$ *and* $(x, y + 1)$ *(if those points are still in the sandpile), and each edge is annotated with 1 (the number of grains sent along that edge during a toppling). Each node* (x, y) *where* $y \neq 0$ *and* $y \neq N$ *is annotated with 2 (i.e.,* $\Delta_{ii} = 2$*), and each node where* $y = 0$ *or* $y = N$ *is annotated with 1.*

I.e., all sites topple when they reach height 2, and they send grains to their upper and right neighbors only. Notice that the nodes where $x + y = N$ are dissipative, since they have no outgoing edges, but still lose 2 grains when they topple. The exact locations of the dissipative sites are not of importance – in the analysis that follows, we will assume N is large enough so that the avalanche of interest does not reach any dissipative sites.

In the context of self-organized criticality, we are primarily concerned with the long-term behavior of the sandpile when grains are repeatedly added to the sandpile at sites that are chosen uniformly randomly. In the long-term, the sandpiles will always reach a recurrent state. I.e., if $s(t)$ denotes the state of the sandpile after t grains have been randomly added (from any starting state), then $\lim_{t\to\infty} \mathbb{P}(s(t) \in \mathbf{R}) = 1$. (This isn't strictly true for sandpiles as defined in part 1 – some additional connectivity requirements are necessary. One possibility is requiring that there exactly one dissipative site).

The following definition will be valuable for reasoning about avalanches.

Definition 3.2. *Let* L_i*, the* i*-th* level set *of a 2-dimensional directed sandpile, be the set of sites in the sandpile for which* $x + y = i$*. Let* L_{ij} *or* $L_{i,j}$ *be the unique element of* L_i *that has* $x = j$*.*

In the long-term behavior of a generic sandpile, we usually consider the behavior when grains are repeatedly added to uniformly randomly selected sites. In the long-term behavior of a directed sandpile, we restrict this and instead add grains uniformly randomly into L_0 only. Under this restriction, grains only move through the other level sets due to avalanches started in L_0, which will be important later.

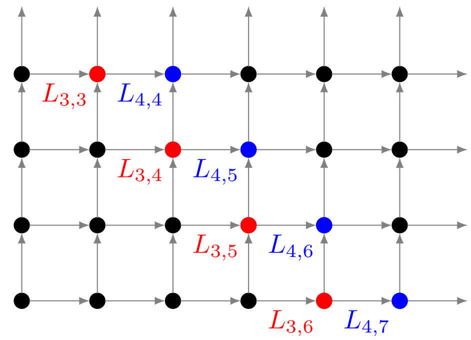

Figure 3: A subset of a 2-dimensional directed sandpile. The red dots belong to the level set L_3 and the blue dots belong to L_4.

Lemma 3.3. *In the long term behavior of the (directed or undirected) sandpile, it is equally likely to be in any of the recurrent states. I.e., if $s(t)$ is the state of the sandpile after t grains have been added randomly, and $s(0) \in \mathbf{R}$, then for any $\mathcal{R} \in \mathbf{R}$, we have $\mathbb{P}(s(t) = \mathcal{R}) = 1/|\mathbf{R}|$.*

I will not prove this lemma, despite its importance. It is related to Theorem 2.5. A proof can be found in [6].

Lemma 3.4. *The sets of recurrent and stable states are equal for 2-dimensional directed sandpiles.*

Proof. I will describe a procedure to reach the empty state from any state by applying a list of σs. This will be sufficient to prove the lemma, for we can reach any stable state from the empty state by adding grains.

Let m be the smallest integer such that L_m is not empty. Add grains to sites in L_m until all of them have toppled exactly once; then all the sites in L_m will be empty. Repeat for L_{m+1}, L_{m+2}, etc. for all non-empty level sets. Since a toppling in L_i does not affect other sites in any $L_j, j \leq i$, at the end of this procedure the state will be entirely empty. QED

Lemma 3.5. *Given a uniformly randomly selected recurrent state of a 2-dimensional directed sandpile, there is an equal probability for any given site to have zero grains or one grain, even if we have knowledge about the number of grains at some other sites.*

Proof. Given a site i, there are an equal number of stable states where that i is empty and where that i has one grain. This remains true if we restrict the set to stable states where some other sites have a fixed number of grains. Since the set of recurrent states is equal to the set of stable states, and all recurrent states are equally probable, we are done. QED

3.2 Random Walks

Next, we develop several facts about probability that will aid us greatly. The essence of the proofs in this subsection are based on Chapter III of [4].

A random walker, informally, is a person who starts off at the origin, then, at fixed intervals, randomly steps forwards or backwards.

Definition 3.6. *A* path *of a 1-dimensional random walker is a function $\gamma(t)$ defined on $t = 1, 2, \dots$. Paths are constrained so that $\gamma(t) = \gamma(t-1) \pm 1$. Unless otherwise noted, $\gamma(0) = 0$. The notation $\gamma_t = \pm 1$ indicates $\gamma(t+1) - \gamma(t)$, i.e., which direction the walker walked at time t.*

There are 2^n possible paths for a walker that takes n steps.

Definition 3.7. *Let $N(n, x)$ (defined for integers x and $n \geq 0$) be the number of possible paths γ for a 1-dimensional random walker satisfying $\gamma(n) = x$.*

Proposition 3.8.

$$N(n,x) = \begin{cases} \binom{n}{\frac{n+x}{2}} & n + x \text{ is even} \\ 0 & \text{otherwise} \end{cases}$$

Proof. If n is odd, then $\gamma(n)$ is odd, and if n is even, then $\gamma(n)$ is even. But if $n + x$ is odd, then n and $\gamma(n)$ would have different evenness for any valid choice of path, so $N(n, x) = 0$.

Now I assume $n + x$ is even. Let p denote the number of distinct integers t such that $\gamma_t = +1$ and q similarly for -1. Then $x = \gamma(n) = \sum \gamma_t = p - q$ and $n = p + q$. We get $N(n, x) = N(p + q, p - q) = \binom{p+q}{p} = \binom{n}{p}$, as this is the number of ways to distribute p many $+1$s into a path of length $p + q$. Finally, we have $\frac{n+x}{2} = \frac{(p+q)+(p-q)}{2} = p$, so $N(n, x) = \binom{n}{(n+x)/2}$. QED

Lemma 3.9 (The Reflection Principle)**.** *Let P be the set of paths γ satisfying $\gamma(0) = x_0 > 0$ and $\gamma(n) = x_1$, and $\gamma(i) = 0$ for some $1 \leq i \leq n$. Let Z be the set of paths γ satisfying $\gamma(0) = -x_0$ and $\gamma(n) = x_1$, for the same x_0 and x_1. Then $|P| = |Z| = N(n, x_0 + x_1)$. In words, the number of paths from x_0 to x_1 which touch zero is equal to the total number of paths from $-x_0$ to x_1.*

Proof. First, $|P| \leq |Z|$. For every $\gamma \in P$, there is a minimal y satisfying $\gamma(y) = 0$. Define a path ζ that starts at $-x_0$ and has $\zeta_i = -\gamma_i$ for $0 \leq i < y$ then $\zeta_i = \gamma_i$ for $i \geq y$. Certainly $\zeta(n) = x_1$, since $\gamma(t)$ and $\zeta(t)$ agree on $t \geq y$, so $\zeta \in Z$. Furthermore, each $\gamma \in P$ maps to a distinct $\zeta \in Z$, so $|P| \leq |Z|$.

Secondly, $|Z| \leq |P|$. For every $\zeta \in Z$, there is a minimal y satisfying $\zeta(y) = 0$. We can define γ the way we defined ζ in the preceding paragraph, and each ζ maps to a distinct γ, so $|Z| \leq |P|$.

The equality $|Z| = N(n, x_0 + x_1)$ is immediate from the definition of Z. QED

The probability of a random walk starting at the origin returning to the origin after n steps is given simply by $N(t, 0)/2^n$, the number of paths returning at time n divided by the total number of paths of length n. Somewhat more difficult is the following:

Theorem 3.10. *For nonnegative even integer n, there are*

$$\frac{1}{n-1}\binom{n}{\frac{n}{2}} \tag{3}$$

many paths that return to the origin for the first time after exactly n steps. The probability of a random path starting at the origin returning to the origin for the first time at the n-th step is

$$\frac{1}{n-1}\binom{n}{\frac{n}{2}}\frac{1}{2^n}. \tag{4}$$

Proof. Let γ be an arbitrary path that returns to the origin for the first time at n. For now we will assume $\gamma_0 = +1$, i.e., the walker's first step is upwards. Then it must be that $\gamma_{n-1} = -1$ (at time $n-1$, the walker cannot be at or below the origin, so much be just above it). In this way the problem is reduced to finding the number of paths of length $n-2$ from $+1$ to $+1$ that do not touch zero. Using Lemma 3.9, this is equal to $N(n-2,0)-N(n-2,2)$. The expression can be rewritten:

$$\begin{aligned}
N(n-2,0)-N(n-2,2) &= \frac{(n-2)!}{\left(\frac{n}{2}-1\right)!\left(\frac{n}{2}-1\right)!} - \frac{(n-2)!}{\left(\frac{n}{2}\right)!\left(\frac{n}{2}-2\right)} \\
&= \frac{\frac{n}{2}}{n}\cdot\frac{\frac{n}{2}}{n-1}\cdot\frac{n!}{\left(\frac{n}{2}\right)!\left(\frac{n}{2}\right)!} - \frac{\frac{n}{2}-1}{n-1}\cdot\frac{\frac{n}{2}}{n}\cdot\frac{n!}{\left(\frac{n}{2}\right)!\left(\frac{n}{2}\right)!} \\
&= \left(\frac{\frac{n}{2}}{n}\cdot\frac{\frac{n}{2}}{n-1} - \frac{\frac{n}{2}-1}{n-1}\cdot\frac{\frac{n}{2}}{n}\right)\binom{n}{\frac{n}{2}} \\
&= \frac{1}{2}\left(\frac{\frac{n}{2}}{n-1} - \frac{\frac{n}{2}-1}{n-1}\right)\binom{n}{\frac{n}{2}} \\
&= \frac{1}{2}\left(\frac{1}{n-1}\right)\binom{n}{\frac{n}{2}}
\end{aligned}$$

Recall, this is the number of paths that return to the origin for the first time at the n-th step and whose first step is upwards. We multiply by 2 to account for paths whose first step is downwards, yielding (3). Finally, we divide by 2^n, the size of the sample space, which yields (4). QED

To describe (4) asymptotically, we invoke *Stirling's formula*, a tool for approximating factorials. Specifically, we will use

$$n! \sim \sqrt{n}\left(\frac{n}{e}\right)^n = n^{n+1/2}e^{-n}$$

Where $\sim$ denotes that the limit of the ratio between the two sides approaches a constant (not necessarily 1) as $n \to \infty$. In the case of Stirling's formula, the constant is $\sqrt{2\pi}$, but is not of importance to us. Using Stirling's formula, we get

$$\binom{2n}{n}\frac{1}{2^n} = \frac{(2n)!}{(n!)^2} \sim \frac{(2n)^{2n+1/2}e^{-2n}}{(n^{n+1/2}e^{-n})^2}\frac{1}{2^{2n}} = \frac{2^{2n+1/2}n^{2n+1/2}e^{-2n}}{n^{2n+1}e^{-2n}2^{2n}} \sim \frac{1}{\sqrt{n}}$$

Accounting for the extra $1/(n-1)$ factor in (4), we get the following:

Proposition 3.11. *The equation (4) is asymptotically proportional to $t^{-3/2}$, in the sense that (4) divided by $t^{-3/2}$ approaches a constant as $t \to \infty$.*

3.3 Avalanche Exponents for the Directed Sandpile

Now we consider how avalanches progress in the directed sandpile.

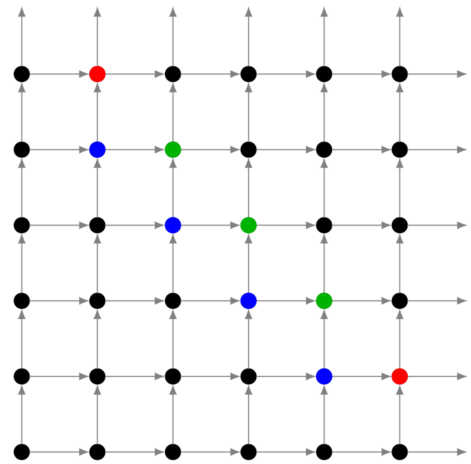

Figure 4: An example of how an avalanche might progress through a directed sandpile. After the blue dots topple, the green dots must also topple, and the red dots topple with probability 1/2. Notice that the set of toppled dots on any level set will always be contiguous.

Lemma 3.12. *In the toppling resulting from a single σ operation on the directed sandpile, for each i there exist m and M such that the sites $L_{i,j}$, $m \leq j \leq M$ are exactly the sites that topple in L_i. I.e., the toppled sites in any level set are contiguous.*

Proof. We proceed by induction. Certainly the statement is true for the level set L_i containing the site of the σ operation, since there is at most one toppling there. Now we consider an arbitrary level set L_j and assume the lemma statement holds for L_{j-1}. Let m_0 and M_0 be the values of m and M for L_{j-1}. All the sites $L_{j,k}$ for $m_0 < k \leq M_0$ will have incoming directed edges from $L_{j-1,k-1}$ and $L_{j-1,k}$, which both toppled, so $L_{j,k}$ will topple also. The sites L_{j,m_0} and L_{j,M_0+1} each have one incoming directed edge from a toppled site in L_{j-1}, so by Lemma 3.5 these two sites have a 1/2 chance of toppling. Any site in L_j not yet mentioned has no incoming link from a toppled site and thus won't topple.

Therefore, if L_{j,m_0} topples, let $m_1 = m_0$, else m_0+1. If L_{j,M_0+1} topples, let $M_1 = M_0+1$, else M_0. As described in the last paragraph, all sites in L_j between these two will topple, and none outside will topple. QED

Thanks to the preceding lemma, we can characterize the effect of a σ operator on each level set by a single integer, the number of sites toppled on that level set.

Definition 3.13. *Let $\delta(i)$ be the number of topplings on L_i as the result of some σ operator on some directed sandpile state.*

Lemma 3.14. *During a toppling due to a σ on a random directed sandpile, if $\delta(i) > 0$, then*

$$\mathbb{P}(\delta(i+1) = \delta(i)) = \frac{1}{2} \tag{5}$$

$$\mathbb{P}(\delta(i+1) = \delta(i)+1) = \mathbb{P}(\delta(i+1) = \delta(i)-1) = \frac{1}{4} \tag{6}$$

Proof. As we saw in the proof of Lemma 3.12, if sites L_{i,m_0} through L_{i,M_0} topple, then sites L_{i+1,m_0+1} through L_{i+1,M_0} also topple. There are $\delta(i) - 1$ such sites. The only other sites in L_{i+1} that may topple are L_{i+1,m_0} and L_{i+1,M_0+1}, each with 1/2 probability. If neither topples, then $\delta(i+1) = \delta(i) - 1$, and this happens with probability $1/2 \cdot 1/2 = 1/4$. If both topple, then $\delta(i+1) = \delta(i) + 1$, with the same probability. Finally, if one topples and the other does not, which occurs with probability 1/2, then $\delta(i+1) = \delta(i)$. QED

Now we show that the "width" $\delta(i+t)$ of the avalanche behaves like a random walker. This is intuitive if we think about each edge of the avalanche being controlled by a random walker who instructs each edge to widen or shrink as the avalanche progresses from one level set to the next. Rather than formalizing a notion of two walkers, we will model the width as a single walker who takes two steps at a time.

Definition 3.15. *An* annihilating walker *is a random walker who stops moving after reaching the origin.*

Lemma 3.16. *Let γ be a uniformly randomly selected path of an annihilating walker who started at $x = 2$. Suppose that σ was applied to a site in L_0 and caused that site to topple. Then $\mathbb{P}(\delta(t) = x) = \mathbb{P}(\gamma(2t)/2 = x)$.*

Proof. The statement is true for $t = 0$, since $\delta(0) = 1$ and $\gamma(0)/2 = 1$ certainly. Now let $t > 0$ be arbitrary and assume the statement for $t - 1$. Also assume $\delta(t-1) > 0$. Then $\mathbb{P}(\delta(t) = \delta(t-1)) = 1/2$ by Lemma 3.14, and $\mathbb{P}(\gamma(2t)/2 = \gamma(2t-2)/2) = 1/2$, since for $\gamma(2t) = \gamma(2t-2)$ the walker must have one step up then one step down, with probability 1/4, or one step down then one step up, with probability 1/4, summing to 1/2. Continuing, $\mathbb{P}(\delta(t) = \delta(t-1) + 1) = 1/4$, and $\mathbb{P}(\gamma(2t)/2 = \gamma(2t-2)/2 + 1) = 1/4$, because the walker must have made two steps up for $\gamma(2t) = \gamma(2t-2) + 2$ to hold. The same argument shows $\mathbb{P}(\delta(t) = \delta(t-1) - 1) = 1/4 = \mathbb{P}(\gamma(2t)/2 = \gamma(2t-2)/2 - 1)$.

Now we do away with the assumption $\delta(t-1) > 0$. If $\delta(t-1) = 0$, then $\delta(t) = 0$ certainly since topplings cannot occur spontaneously. Similarly, if $\gamma(2t-2) = 0$, then $\gamma(2t) = 0$ because the walker is annihilating.

Since $\mathbb{P}(\delta(t) = x)$ is entirely determined by the probabilities $\mathbb{P}(\delta(t-1) = y)$, which by the inductive hypothesis are equal to the corresponding probabilities on γ, and the probabilities $\mathbb{P}(\delta(t) = x | \delta(t-1) = y)$, which I showed in the previous paragraphs to be equal to the corresponding probabilities on γ for all relevant values. QED

We say the avalanche started by a σ in L_i ends at the lowest t satisfying $\delta(i+t) = 0$. With the preceding lemma in hand, we see that the probability of the avalanche ending at t is equal to the probability of a random walker starting at $x = 2$ first reaching the origin at $2t$.

Lemma 3.17. *If a σ in L_0 causes a toppling, then the probability of the avalanche ending at t is*

$$\frac{1}{2t+1}\binom{2t+2}{t+1}\frac{1}{2^{2t+2}} \tag{7}$$

Proof. Consider the possible paths of a walker who starts at $x = 0$ as usual and returns to the origin for the first time at the n-th step, with $n > 2$. In all such paths γ, it must be that $\gamma_1 = \gamma_2$, else the walker would return to the origin after exactly two steps. Hence half of the valid paths γ have $\gamma(2) = 2$ and the other half have $\gamma(2) = -2$.

That is, the number of paths starting from $x = 2$ which reach the origin for the first time at step n is equal to half the number of paths from the origin that return to the origin for the first time at $n + 2$ as given by (3). Specifically:

$$\frac{1}{n+1}\binom{n+2}{\frac{n+2}{2}}\frac{1}{2}.$$

The sample space is paths of length n, since n steps occur from reaching $x = 2$ through reaching the origin. Therefore, the probability of a walker starting at $x = 2$ first reaching the origin at time n is

$$\frac{1}{n+1}\binom{n+2}{\frac{n+2}{2}}\frac{1}{2^{n+1}}$$

The probability of the avalanche ending at $i + t$ is equal to half the probability of the walker starting at $x = 2$ first reaching the origin at step $2t$, per Lemma 3.16. Substituting $n = 2t$ then dividing the above equation by two yields (7). QED

The asymptotic estimate in Proposition 3.11 applies to equation (7) as well.

The discussion that follows is fairly close to what may be found in [2], and serves as the reasoning behind the final theorem. It is quite informal.

What's more interesting than the number of level sets that an avalanche affects is the total number of sites that topple in an avalanche. This is what we call the *size* of an avalanche. The ultimate goal of this section is to determine the probability that an avalanche will have a certain size.

We could calculate the average size of an avalanche based on its duration using more probability about random walks, but it's more elegant to use a law about sandpiles instead. Seeing as the sandpile is finite, there is some maximum number of grains it may contain. The only dissipative sites are at $x + y = N$, i.e., in the "last" level set. The only way for grains to get from L_0, where they are added, to L_N, where they exit, is to pass through every level set. Therefore, the total number of grains entering and exiting each level are asymptotically equal as more grains are added. Equivalently, the number of topplings in any given level set asymptotically approaches half the total number of grains added to the sandpile (since each toppling causes two grains to leave the level set).

We will now formalize the ideas of the preceding paragraph. Let D be the duration (number of level sets affected) of an avalanche. Let $\text{layerToppleCount}(i)$ be the expected number of topplings in L_i that an avalanche causes given that the avalanche reaches L_i (causes at least one toppling in L_i). Then, if t grains are added, it must be that:

$$2 \cdot \text{layerToppleCount}(i) \cdot \mathbb{P}(D > x) \cdot t \sim t \tag{8}$$

Specifically, $\text{layerToppleCount}(i) \cdot \mathbb{P}(D > x)$ must approach $1/2$ as many grains are added. To determine the asymptotic behavior of $\mathbb{P}(D > x)$, we recall from Proposition 3.11 that the probability of an avalanche ending at L_i grows like $i^{-3/2}$. As the sandpile is large, we can use

$$\mathbb{P}(D > i) \sim \int_i^\infty x^{-3/2} dx \sim i^{-1/2}$$

Then, by (8), it must be that $\text{layerToppleCount}(i) \sim 1/\mathbb{P}(D > i) \sim i^{1/2}$. Let us now define S, the previously discussed size of an avalanche, or more specifically, the total number of sites toppled during an avalanche. Let S_i be the expected value of S given that the avalanche ends at level set i. Then we have, roughly,

$$S_i \sim \text{layerToppleCount}(i) \cdot i \sim i^{3/2}$$

(The number of layers toppled times the number of sites toppled per layer). Our goal is to find $\mathbb{P}(S = x)$. First, we focus on $\mathbb{P}(S > x)$. Notice that when $D > x$, we can expect $S > x^{3/2}$, and therefore, if we want $\mathbb{P}(S > x)$, we should look for conditions on $\mathbb{P}(D > x^{2/3})$. We know that $\mathbb{P}(D > i) \sim i^{-1/2}$, so $\mathbb{P}(S > x) \sim \mathbb{P}(D > x^{2/3}) \sim x^{-1/3}$. Then, assuming the sandpile is large enough, $\mathbb{P}(S = x)$ is like $\frac{d}{dx}\mathbb{P}(S > x) \sim x^{-4/3}$.

Theorem 3.18. *In the two-dimensional directed sandpile model, in a uniformly randomly selected recurrent state, the probability of t many sites toppling due to a σ operation in L_0 varies like $t^{-4/3}$ as $t \to \infty$.*

4 Application to Electricity Grids

In 2003, the worst North American electricity blackout in recent history started in Ontario before quickly spreading to much of the eastern seaboard. The cause wasn't a nuclear meltdown or an unprecedented solar flare. Instead, a medium-size coal plant unexpectedly went offline. The plant wasn't small, but neither was it so large that its failure should have overwhelmed the rest of the grid. Many other components of the grid were operating near maximum capacity, so the increase in long-distance current caused some transmission lines to overheat, fall onto trees, and disconnect, putting even more load on the remaining transmission lines. In other words, the grid was near criticality, so a relatively small event was able to trigger a devastating avalanche of events. In this section, I'll provide a simple model of an electrical grid as a sandpile, as suggested in [1].

Each grain of sand represents some unit of "load" on the electricity grid. For example, a transmission line that is near capacity may merit several grains of sand. Pieces of infrastructure that are geographically near and are affected by one another have nonzero Δ_{ij}, representing that when a piece of infrastructure fails, nearby pieces of infrastructure are the ones that must immediately handle the load. When a site topples, it represents the failure of a node, for example a transmission line that contacts a tree and then disconnects.

It's a bit harder to explain the dissipative sites. After all, when a grid recovers from a blackout, it is not immediately significantly more robust than it was before the blackout, so one would expect there to be the same number of grains in the same places. However, we make the observation that after an electrical blackout, grid operators usually significantly upgrade the systems that caused the fault, install new systems to better balance the load, and set in place new protocols to help decrease the risk of future blackouts. The scale of these improvements is usually correlated to the seriousness of the blackout. Thus, after a large "toppling", the "number of grains" (effective load on the grid) does in fact decrease.

As sandpiles exhibit self-organized criticality, and we can roughly model an electrical grid with a sandpile, one might expect most electrical grids to naturally be near criticality, where blackouts ranging from very small to effectively infinitely large can result from very small impulses, such as the failure of a coal plant in the case of the 2003 Ontario blackout. The authors of [1] empirically discovered that the size of electrical blackouts between 1993 and 1998 varied approximately with a power law with the exponent in the critical range, which is typical of a system exhibiting self-organized criticality.

References

[1] Benjamin Carreras, D.E. Newman, and A Poole. "Initial Evidence for Self-Organized Criticality in Electric Power System Blackouts". In: International Conference on System Sciences. Maui, Hawaii, Jan. 2000.

[2] Deepak Dhar. "Theoretical studies of self-organized criticality". In: *Physica A: Statistical Mechanics and its Applications* 369.1 (Sept. 2006), pp. 29–70. ISSN: 03784371. DOI: 10.1016/j.physa.2006.04.004. URL: https://linkinghub.elsevier.com/retrieve/pii/S0378437106004006 (visited on 04/29/2021).

[3] Deepak Dhar and Ramakrishna Ramaswamy. "Exactly solved model of self-organized critical phenomena". In: *Physical Review Letters* 63.16 (Oct. 16, 1989), pp. 1659–1662. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.63.1659. URL: https://link.aps.org/doi/10.1103/PhysRevLett.63.1659 (visited on 06/03/2021).

[4] William Feller. *An introduction to probability theory and its applications. Vol. 1.* 3. ed. 1968. 509 pp.

[5] Akhil Jalan. "The Structure of the Sandpile Group". B.S. Thesis. University of California Berkeley, May 20, 2019. 45 pp.

[6] Ronald Meester, Frank Redig, and Dmitri Znamenski. "The Abelian sandpile; a mathematical introduction". In: *Markov Processes and Related Fields* (Apr. 25, 2001).

[7] F Redig. "Mathematical aspects of the abelian sandpile model". In: (June 21, 2005), p. 60.

The figures in this paper are original.