The Stationary Density and Threshold Density of Sandpiles

A Thesis
Presented to
The Division of Mathematics and Natural Sciences
Reed College

In Partial Fulfillment
of the Requirements for the Degree
Bachelor of Arts

Vivian Tylińska

May 2017

Approved for the Division
(Mathematics)

David Perkinson

Acknowledgements

Thank you, Sean, Charlie, Stefan, Zach, Kasie, Yelena, Mike, Carlo, Taylor, Anders, Jimmy, Jackson, Jake, Josh, Jenna, Jess, Emily, Violet, Emma, Grace, and the whole lot of you, for making college fun.

Thank you Dave, for teaching me about sand.

Table of Contents

List of Figures

Abstract

In this paper we examine the relationship between the stationary density ζ_{D_0} of a graph G and its stationary density ζ_{st}, and give an exposition of the Threshold Density Theorem due to Levine [2015], which states that the threshold density converges to the stationary density as the degree of the initial divisor D_0 approaches $-\infty$. We will then prove the Threshold Density Theorem directly on the banana graph B_n.

Introduction

Imagine a group of friends, Alice, Bob, Charlie, and Diane, all of whom live in houseboats, connected by bridges. Bob and Charlie can walk directly to any of their friends' houses, while Alice and Diane can walk directly to Bob and Charlie's. We can imagine the following map:

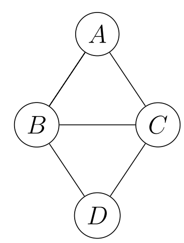

The orientation of where the houseboats lie and the length of the bridges does not concern us. A map like this is an example of a *graph*. The notion of a mathematical graph generalizes a set of relationships between objects. We say a graph consists of a set of vertices and edges which connect them. As another example, we have the following graph with four vertices and five edges.

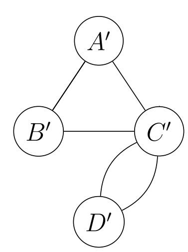

For the scope of this paper we will restrict ourselves to graphs that are *connected*, that is, each of our friends can reach each others' houses via some bridges and *undirected*, that is, if Bob can walk to Alice's, Alice can also walk to Bob's. Further, we allow multiple bridges between the same pair of destinations.

Let us return to our four friends and imagine that they are playing a game of poker. Alice has one dollar, Bob is two dollars in debt, Charlie has two dollars, and Diane has four dollars. We can assign these values to our previous graph:

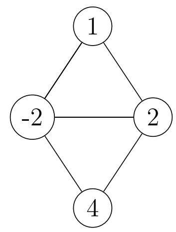

We now introduce the notion of *chip-firing*. In order to move around the money they have, we allow a person to either "lend" or "borrow," that is, for each neighbor, either send them one dollar or borrow a dollar from each. For example, on our graph, if Diane were to perform a lending move, we would see,

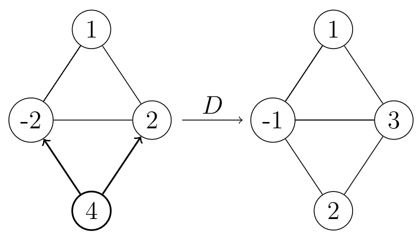

Likewise, if after this, Bob performed a borrowing move, we would see,

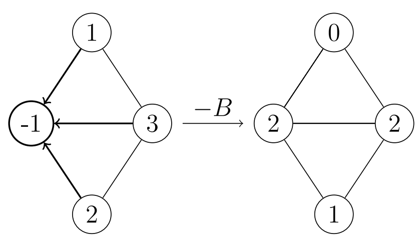

Note that if a vertex lends a dollar to one of its neighbors, it must simultaneously lend a dollar to *all* its neighbors, and vice versa with borrowing.

So that our running example will conform with later notation, we will relabel the vertices.

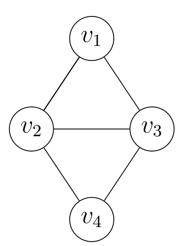

Figure 1: The Diamond Graph.

Next, we will want to introduce the notion of stability and instability. We will say a vertex is *unstable* if it has as least as much money as it has neighbors. Otherwise, it is *stable*. We say that a graph is *stable* if all its vertices are stable. An unstable vertex will "fire," that is, send a dollar to each of its neighbors. After some sequence of firings a graph will either reach a stable state, or return to a past unstable state and fire indefinitely. We will call this graph *alive*. For example, observe the behavior of the following:

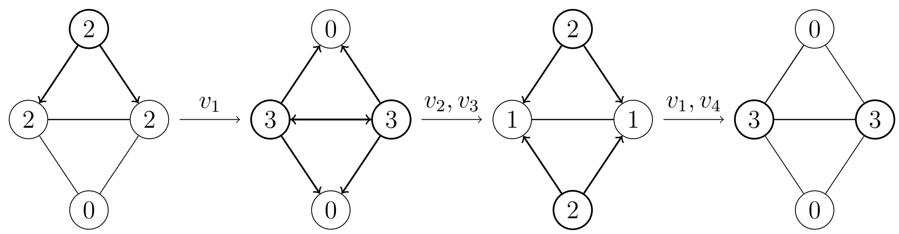

We can see that this system will not reach a stable state.

Let us suppose now that we start with zero dollars on each vertex and, one at a time, randomly drop a dollar on some vertex and allow the system to stabilize until it becomes alive. For example,

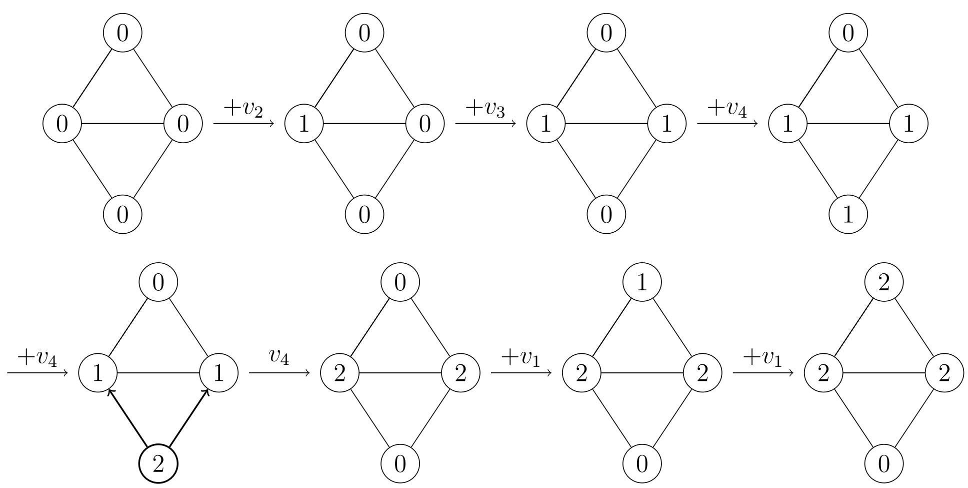

Since this final arrangement is unstabilizable, as we have seen before, the system becomes alive. We will call the first unstabilizable state that is reached a *threshold state*. The average amount of dollars on any given vertex is called the *threshold density*, and in this case it is $\frac{6}{4}$.

Depending on which vertices we drop the dollars on, we may reach any one of many possible threshold states, each with a different probability. We use the notation (a_1, a_2, a_3, a_4) to represent the amounts of money on v_1, v_2, v_3, and v_4, respectively. In our example there are 16 threshold states where the total number of dollars is 5:

$$(0,1,3,1),(0,1,2,2),(0,3,1,1),(0,2,1,2),(0,3,2,0),(0,2,3,0),(2,0,2,1),(1,0,3,1)$$

$$(1,0,2,2),(2,1,2,0),(1,1,3,0),(2,2,0,1),(1,3,0,1),(1,2,0,2),(2,2,1,0),(1,3,1,0).$$

There are 12 threshold states where the total amount is six,

$$(0,3,2,1),(0,2,3,1),(0,2,2,2),(2,1,2,1),(1,1,3,1),(1,1,2,2)$$

$$(2,2,1,1),(1,3,1,1),(1,2,1,2),(2,2,2,0),(1,3,2,0),(1,2,3,0),$$

and 4 where the total is 7:

$$(2,2,2,1),(1,3,2,1),(1,2,3,1),(1,2,2,2).$$

So, the threshold density will be somewhere between $\frac{5}{4}$ and $\frac{7}{4}$, depending on the probabilities of reaching these states. Calculating the probabilities of reaching these states tends to be difficult by hand, but with a computer we are able to find that the threshold density is $\frac{5775}{4096} \approx 1.41$.

In this essay we will be examining facts about the threshold density. In Chapter 1, we will first formalize our chip-firing model and then introduce the notion of a *sink*, a vertex down which sand can disappear. We will use a process called a Markov chain to examine the long-term behavior of a graph with a sink, which will be called our Abelian Sandpile model, and then apply those results to the sinkless model. In Chapter 2, we will compare these two models to prove the *Stationary Density Theorem* due to Levine [2015], which states that the threshold density approaches a related constant called the *stationary density*, as the initial state of the sinkless model becomes increasingly negative. Finally, in Chapter 3, we will examine in depth an example on a simple type of graph called the *banana graph*, on which we will prove directly that the threshold density approaches the stationary density.

Chapter 1

The Abelian Sandpile Model

In this chapter, we will properly define the chip-firing model we presented in the introduction. Then, we will introduce the Markov chain, which will allow us to examine long-term distributions of evolving systems. Finally, we will apply the Markov chain to our model. For detailed proofs of the theorems presented, refer to Perkinson and Corry [2017].

1.1 Formalizing the Chip-firing Model

1.1.1 Without Sink

In the introduction, we examined amounts of “money” on a graph. We will now formally define a graph. By a *multiset* we mean a set in which elements may occur multiple times.

Definition 1.1.1. A *graph*, $G = (V, E)$, consists of a set of *vertices* V and a multiset of *edges* which connect them, E. The elements of E will be unordered pairs of elements of V.

In this essay we will restrict ourselves to *connected* graphs, where any vertex is reachable from another by a series of consecutive edges. Since E is a multiset, multiple edges connecting two vertices will also be allowed.

Definition 1.1.2. A *divisor* D on a graph $G = (V, E)$ is an element of the free Abelian group on the vertices,

$$\mathrm{Div}(G) = \mathbb{Z}V = \sum_{v \in V} D(v)v.$$

A divisor represents the distribution of wealth on a graph by letting $D(v)$ refer to the amount of dollars that vertex v has. We refer to the *degree* of a divisor as the sum of these coefficients, i.e., the net wealth. That is,

Definition 1.1.3. The *degree* of a divisor D is given by $\deg(D) = \sum_{v \in V} D(v)$.

Note that for $D, E \in \mathrm{Div}(G)$, we have $\deg(D) + \deg(E) = \deg(D + E)$. We will let $\mathrm{Div}^k(G)$ denote the set of divisors with degree k.

For a vertex v, however, $\deg(v)$ may either refer to its degree as a divisor, in which case $\deg(v) = 1$ or to the number of edges to which v is adjacent. We will use $\deg_G(v)$ to refer to the latter to avoid confusion.

To represent firing formally, we define the following:

Definition 1.1.4. For two divisors $D, D' \in \text{Div}(G)$ and a vertex $v \in V$ we say that D' is obtainable from D by a *lending move at* v, denoted $D \xrightarrow{v} D'$, if

$$D' = D - \deg_G(v)v + \sum_{vw \in E} w.$$

A borrowing move, correspondingly, will have the signs reversed on $\deg_G(v)v$ and the sum. As it turns out, borrowing and lending moves are Abelian: for any ordering of some firings or lendings, we will come to the same final state.

We will call two divisors D, D' *linearly equivalent* if D is obtainable from D' by some sequence of lending and borrowing moves, and vice versa, and say $D \sim D'$. The set of divisors linearly equivalent to some D will be denoted $[D]$. Note that if $D \sim D'$, then we have $\deg(D) = \deg(D')$, since lending and borrowing moves do not affect the degree.

Definition 1.1.5. The (discrete) *Laplacian operator* on G is the linear mapping $L \colon \mathbb{Z}V \to \mathbb{Z}V$ determined by

$$L(v) := \sum_{vw \in E} (v - w).$$

The *Laplacian matrix* $\mathcal{L}$ is the matrix representation of the Laplacian operator, and encodes set firings. That is, for some divisor $D = (D(v_1), D(v_2), \cdots, D(v_n))$ and a *firing script* $\sigma = (a_1, a_2, \cdots, a_n)$, where a_i denotes firing v_i a total of a_i times, the divisor $D' = D - \mathcal{L}\sigma$ gives the resulting divisor after each $v_i \in V$ has been fired a_i times. For example, for the diamond graph in figure 1 in the introduction, $\sigma = (1, 0, 1, 0)$ would denote firing v_1 and v_3 once each, so $\mathcal{L}\sigma$ would have the net effect $(-1, 2, -2, 1)$. Note that if $D \sim D'$ then there exists a firing script σ such that $D' = D - \mathcal{L}\sigma$.

Let us return to the notion of *stability*.

Definition 1.1.6. A vertex v is *stable* for $D \in \text{Div}(G)$ if $D(v) < \deg_G(v)$. A divisor D is stable if all of its vertices are stable. A divisor D is *stabilizable* if it is linearly equivalent to some divisor D which is stable. A divisor which is not stabilizable is *alive*.

Naturally, if a divisor D is stable, then each divisor in $[D]$ is stabilizable.

Next we consider some properties of linear equivalence which will lead us to introduce the notion of the sink. We define the addition of divisors vertex-wise, that is,

$$D + F = \sum_{v \in V} (D(v) + F(v))v.$$

Definition 1.1.7. The *Picard group*, $\mathrm{Pic}(G)$, is defined as the set of linear equivalence classes of divisors,

$$\mathrm{Pic}(G) = \mathrm{Div}(G)/\sim,$$

with addition

$$[D] + [F] = [D + F].$$

Definition 1.1.8. The subgroup of $\mathrm{Pic}(G)$ consisting of divisors with degree zero is called the *Jacobian group*, denoted $\mathrm{Jac}(G)$:

$$\mathrm{Jac}(G) = \mathrm{Div}^0(D)/\sim .$$

Since the two divisors with degree zero will also have degree zero, the group is well-defined. As such, we can establish an isomorphism between the $\mathrm{Pic}(G)$ and $\mathbb{Z} \times \mathrm{Jac}(G)$. Picking some $q \in V$, an isomorphism is given by,

$$\begin{aligned} \mathrm{Pic}(G) &\to \mathbb{Z} \times \mathrm{Jac}(G) \\ [D] &\to (\deg(D), [D - \deg(D)q]). \end{aligned}$$

This means any divisor class is equivalent to some element in the Jacobian, with one vertex adjusted to make the total degree of the divisor zero.

1.1.2 With Sink

Rather than thinking of a divisor on a graph as representing an amount of wealth which can be borrowed and lent, we may think of a graph as having an amount of "grains of sand" on it instead. Suppose we pick some arbitrary vertex and denote it the *sink*, with the property that any sand that enters it will disappear. With regard to this, we will redefine several terms: rather than talking about divisors, we will instead speak of *configurations* and *sandpiles*.

Definition 1.1.9. On a graph $G = (V, E, s)$ with sink $s \in V$, a *configuration* $c \in \mathrm{Config}(G)$ is defined as an element of the free Abelian group on $\widetilde{V} = V \setminus s$:

$$\mathrm{Config}(G) = \mathbb{Z}\widetilde{V} = \sum_{v \in \widetilde{V}} D(v)v.$$

A configuration is essentially just a divisor D with $D(s) = 0$, so $\mathrm{Config}(G)$ is naturally a subset of $\mathrm{Div}(G)$. The degree is defined accordingly:

Definition 1.1.10. The *degree* of a configuration c is given by $\deg(c) = \sum_{v \in \widetilde{V}} c(v)$.

Firing vertices functions exactly as it does in the sinkless model with lending moves, except that the sink cannot fire.

Definition 1.1.11. For two configurations, $c, c' \in \mathrm{Config}(G)$ and a vertex $v \in \tilde{V}$ we can say that c' is obtainable from c by a *firing* at v, denoted $c \xrightarrow{v} c'$ if

$$c' = c - \deg_G(v)v + \sum_{vw \in E : w \neq s} w.$$

A *reverse firing* at v, denoted $c \xrightarrow{-v} c'$ is then given by

$$c' = c + \deg_G(v)v - \sum_{vw \in E : w \neq s} w.$$

So a reverse firing corresponds to a borrowing move in the sinkless model.

Stability and instability are defined the same way: a configuration c is stable if all of its vertices (except the sink) are stable, that is, if $c(v) < \deg_G(v)$ for all $v \in \tilde{V}$.

The main difference between configurations and divisors is the notion of stabilization and stabilizability.

Definition 1.1.12. Let $c \in \mathrm{Config}(G)$. It is *legal* to fire v if $c(v) \geq deg_G(v)$, that is, if v is unstable.

Since sand can disappear down the sink, the degree of a configuration can decrease after successive firings. In fact, however large the degree of an unstable configuration, after a finite number of firings enough sand will have gone down the sink for the configuration to reach stability.

Definition 1.1.13. For any configuration $c \in \mathrm{Config}(G)$, the *stabilization* of c is denoted c°.

Proposition 1.1.1. *For any configuration $c \in \mathrm{Config}(G)$, its stabilization c° is unique, and its firing script σ is also unique.*

Next, we define an important subset of $\mathrm{Config}(G)$.

Definition 1.1.14. If $c \in \mathrm{Config}(G)$ and $c(v) \geq 0$ for all $v \in \tilde{V}$ we write $c \geq 0$ and c is called a *sandpile*. The set of all stable sandpiles is denoted $\mathrm{Stab}(G)$. We define the operation *stable addition*, on two stable sandpiles $a, b \in \mathrm{Stab}(G)$ denoted by $a \circledast b$, as vertex-wise addition followed by stabilization, that is,

$$a \circledast b = (a + b)^\circ.$$

This means $\mathrm{Stab}(G)$ is very nearly a group under $\circledast$—it only lacks inverses. As we will soon see, there is a subset of $\mathrm{Stab}(G)$ which does form a group, and we will call this group $S(G)$ and its elements *recurrents*. It is worth noting that dropping many grains of sand on G and allowing it to stabilize, noting the number of times each configuration in Stab(G) appears, we find that certain sandpiles only appear once or never, and others appear many times. The latter will turn out to be those recurrents, which we will define later.

In the previous section we examined the groups formed by modding out divisors via equivalence. We will do something similar for Config(G), modding out via firings and reverse-firings.

Definition 1.1.15. The *critical group* $K(G)$ of a graph $G = (V, E, s)$ with sink s is given by

$$K(G) := \text{Config}(G)/\sim .$$

Here, $a \sim b$ means a is obtainable from b via a sequence of (not necessarily legal) firings and reverse firings.

1.2 Introducing Markov Chains

We will now define Markov chains, which will allow us to examine the long-term behavior of a group of states with a set of probabilities for going from one state to another. As a general reference for Markov chains, the reader may consult Levin et al. [2009].

Definition 1.2.1. A finite Markov chain consists of:

1. A finite set of states, Ω.
2. A function, $P\colon \Omega \times \Omega \to [0, 1]$ with the property that for all $x \in \Omega$, $P(x, \cdot)$ is a probability distribution on Ω, i.e.: $\sum_{y\in\Omega} P(x, y) = 1$.
3. A sequence of random variables $(X_0, X_1, X_2, \dots)$ which satisfy the law of the chain,
$$\mathbb{P}(X_{t+1} = y\colon X_t = x) = P(x, y).$$

We now define several terms which will give us a definition for the recurrent states we referred to earlier in the model with sink.

Definition 1.2.2. In a Markov chain $(\Omega, P, (X_t))$, for two states $x, y \in \Omega$, the state y is said to be *accessible* from x if there exists an $n > 0$ such that $P^n(x, y) > 0$. If x is also accessible from y, then x and y are said to *communicate*. If all states communicate then the Markov chain is said to be *irreducible*.

Definition 1.2.3. A state $x \in \Omega$ of a Markov chain is said to be *essential* if its communicating class (i.e., the set of states with which it communicates) consists of all $y \in \Omega$ that are accessible from x.

This means that if a state x is essential, then any state which can be accessed from x can also access x. Now, we define a *recurrent* state on a Markov chain.

Definition 1.2.4. For a Markov chain, (Ω, P, X_t), a state $x \in \Omega$ is *recurrent* if, starting from x, the probability that the chain returns to x at some point is 1.

Proposition 1.2.1. *On a Markov chain (Ω, P, X_t), a state is recurrent if and only if it is essential. Every Markov chain has at least one essential state, and for any inessential state there is some essential state which is accessible from it.*

There are several statements we can make about the long-term behavior of a Markov chain in terms of *stationary distributions*. Given some initial probability distribution π_0 on our set of states, we can write π_t to be the probability distribution at some time (i.e., number of steps) t, where $\pi_t = \pi_{t-1}P$. If there exists a limiting distribution π, that is, if $\pi = \lim_{t\to\infty} \pi_t$ for some initial π_0 then it must be the case that $\pi = \pi P$.

Definition 1.2.5. A probability distribution π on Ω is a *stationary distribution* if $\pi = \pi P$.

Proposition 1.2.2. *Suppose π is a stationary distribution for (Ω, P, X_t). Then if $x \in \Omega$ is inessential, $\pi(x) = 0$.*

We will need one more definition to guarantee the existence of a stationary distribution.

Definition 1.2.6. For a finite Markov chain, (Ω, P, X_t) we define the *period* of a state x to be the greatest common divisor of the set $T(x) := \{n \geq 1 : P^n(x, x) > 0\}$, that is, the greatest common divisor of the set of times for which it is possible for x to return to itself. If no such set exists then the period is not defined. If each state has period 1, we say the Markov chain is *aperiodic*.

Theorem 1. *Suppose (Ω, P, X_t) is an irreducible finite Markov chain. Then there exists a unique stationary distribution π. Furthermore, if the chain is also aperiodic, then for any initial distribution π_0, the chain will converge to π.*

1.3 The Abelian Sandpile Model

1.3.1 Without Sink

We will apply our theory of Markov chains first to the sinkless sandpile model. At each step in time we will randomly drop a grain of sand on some vertex of a graph and attempt to stabilize until the system becomes alive. We write,

Definition 1.3.1. The *Sinkless Abelian Sandpile Model* on a graph $G = (V, E)$ is a finite Markov chain (Ω, P, D_t) where the state space Ω consists of the set of divisors, $\mathrm{Div}(G)$. The probability matrix, given $D, D' \in \mathrm{Div}(G)$ and some probability distribution on the vertices $\alpha : V \to [0, 1]$ where $\alpha(v) > 0$ for all v, is given by

$$P(D, D') := \begin{cases} \alpha(v), & \text{if } D' = D + v \text{ and } D + v \text{ is alive} \\ \sum_{v \in V : (D+v)^\circ = D'} \alpha(v), & \text{otherwise,} \end{cases}$$

where $(D+v)^\circ$ denotes the stabilization of $D+v$.

Proposition 1.3.1. *Let $D \in \mathrm{Div}$ be a stabilizable divisor. Then the firing script σ for the stabilization of D is unique up to the all 1s vector, $1_V = \sum_{v \in V} v$.*

Additionally, we define an operator a_v to represent dropping a grain of sand on $v \in V$ and attempting to stabilize. So,

$$a_v D := \begin{cases} D+v, & \text{if } D+v \text{ is alive} \\ (D+v)^\circ, & \text{otherwise.} \end{cases}$$

The random variable (D_t), beginning at $D_0 \in \mathrm{Div}(G)$, is given by $D_{t+1} = a_{v_{t+1}} D_t$.

We will want to consider the states at which the sandpile first becomes alive. We call these *threshold states* and define a time τ for the first divisor D_τ which is alive.

Definition 1.3.2. The *threshold time from starting state* D_0 is the random variable on the sinkless Abelian Sandpile defined as,

$$\tau := \tau(D_0) = \min\{t \geq 0 \colon D_t \text{ is alive}\}.$$

Accordingly, D_τ is a *threshold state*, and v_τ is the *vertex of first unstabilizability* or *epicenter*.

These are all random variables on our Markov chain.

Definition 1.3.3. The *threshold density* of graph G in the sinkless Abelian Sandpile model is the expected value of the average amount of sand on any given vertex at threshold, as given by

$$\zeta_\tau(D_0) := \mathbb{E}_{D_0} \frac{\deg(D_\tau)}{|V|}.$$

1.3.2 With Sink

The Abelian Sandpile model with a sink will also be a Markov chain, defined in a similar fashion to the sinkless model, except with the property that the chain will not terminate, since a sandpile cannot become alive. This will allow us to find a limiting distribution and define the recurrent.

Definition 1.3.4. Let $\alpha \colon V \to [0,1]$ be a probability distribution on the vertices of G, with $\alpha(v) > 0$ for all v (so that there is a non-zero chance of sand being dropped on any vertex, including the sink). The *Abelian Sandpile with a sink* is defined as the Markov chain (Ω, P, c_t) on $G(V, E, s)$, with the state space $\Omega = \mathrm{Stab}(G)$, and probability matrix given by

$$P(c, c') = \sum_{v \in V \colon (c+v)^\circ = c'} \alpha(v).$$

Here, our sequence of random variables c_t starts at $c_0 \in \mathrm{Stab}(G)$ and proceeds as $c_{t+1} = a_{v_{t+1}} c_t$, where we define $a_v c = (c+v)^\circ$ if $v \neq s$ and $a_s c = c$.

Definition 1.3.5. A stable sandpile $c \in \mathrm{Stab}(G)$ is a *recurrent* if for every $a \in \mathrm{Config}(G)$, there exists $b \in \mathrm{Config}(G)$ where $b \geq 0$ such that $a \circledast b = c$.

Proposition 1.3.2. *The set of recurrents, $S(G)$, a subset of* $\mathrm{Stab}(G)$ *forms a group with respect to the operation $\circledast$, which denotes addition followed by stabilization. Further, this group is isomorphic to $K(G)$, the critical group of a graph, as well as* $\mathrm{Jac}(G)$*, the Jacobian group.*

Though we cannot always easily find the recurrent sandpiles, the maximal stable sandpile, $c_{\max}$, is always a recurrent, therefore universally accessible.

Proposition 1.3.3. *The* maximal stable configuration*, given by*

$$c_{\max} := \sum_{v \in \tilde{V}} (\deg_G(v) - 1)v$$

is always a recurrent.

Proof. For any $a \in \mathrm{Config}(G)$, let $b = c_{\max} - a^\circ$, and let σ be the firing script that stabilizes a. Then, since $c_{\max}$ is the largest possible stable sandpile, $b \geq 0$. Firing σ on $a + b = a + (c_{\max} - a^0)$ will then also be legal, and yield $a^\circ + c_{\max} - a^\circ = c_{\max}$. Hence, $(a + b)^\circ = c_{\max}$ and it follows that $c_{\max}$ is recurrent. ☐

Example of Recurrents

We will examine the recurrents of the diamond graph we presented in the introduction. There are several methods for finding the set of recurrents on a graph, analytically or experimentally, but we will not go into detail about this. Suppose we again have the diamond graph:

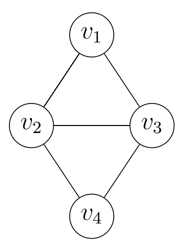

If we choose v_4 as the sink, and write configurations as $(c(v_1), c(v_2), c(v_3))$, the recurrents will be given by,

$$S(G) = \{(1,2,2), (0,2,2), (0,1,2), (1,2,0), (1,0,2), (0,2,1), (1,1,2), (1,2,1)\}.$$

Observing the addition of two elements, say, $(0,2,1)$ and $(1,1,2)$, we see,

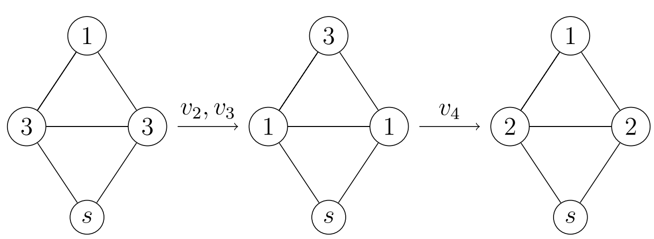

Likewise, if we choose v_2 to be the sink, writing configurations as $(c(v_1), c(v_3), c(v_4))$ we will get the following:

$$S(G) = \{(1,2,1), (0,1,1), (0,2,1), (1,0,1), (1,1,1), (1,2,0), (0,2,0), (1,1,0)\}.$$

Under the operation $\circledast$, these two groups are isomorphic.

Proposition 1.3.4. *For a graph $G = (V, E, s)$ with sink s, the set $S(G)$ of recurrent configurations is identical to the set of recurrent states on the Markov chain $(\text{Stab}(G), P, c_t)$ for the Abelian Sandpile.*

Proof. In the Markov chain, we said that a state was recurrent if and only if it was essential. However, it immediately follows from the definition of a recurrent sandpile that every pair of recurrent sandpiles communicate. □

In order to guarantee a limiting distribution and convergence to it from any starting distribution, we recall, we specified that a Markov chain had to be irreducible and aperiodic. Accordingly, to assure irreducibility, we first modify the Abelian Sandpile model with sink by restricting its state space to $S(G) \subset \text{Stab}(G)$.

Theorem 2. *The stationary distribution of the chain $(S(G), P, c_t)$ is the uniform distribution $\pi(c) = u(c) = \frac{1}{|S(G)|}$ for all $c \in S(G)$.*

Proof. We know the corresponding Markov chain must be aperiodic since at any point a grain of sand may be dropped on the sink, with no effect on the graph. So the period of each state must be 1, and so the chain will converge to the stationary distribution from any initial distribution.

To show that $u(c)$, the uniform distribution, is a limiting distribution, it will suffice to show that $u = uP$. For $c, c' \in S(G)$,

$$uP(c') = u\left(\textstyle\sum_{c \in S(G)} P(c, c')\right) = \frac{1}{|S(G)|} \sum_{c \in S(G)} \sum_{v \in V : (c+v)^\circ = c'} \alpha(v).$$

For $a, b \in S(G)$ and $v \in V$, $(a + v)^\circ = (b + v)^\circ$ if and only if $a = b$. So, for each v there will be a unique c such that $(c + v)^\circ = c'$. Summing α over each $v \in V$ gives 1 since α is a probability distribution on V. Therefore,

$$\frac{1}{|S(G)|} \sum_{c \in S(G)} \sum_{v \in V :\ (c+v)^\circ = c'} \alpha(v) = \frac{1}{|S(G)|} = u(c').$$

□

Thus, the probability of being at any given recurrent c becomes equally likely as time goes on.

Finally we state a theorem due to Dhar [1990], which will motivate something called a *basic alive divisor* in the next chapter.

Theorem 3. *Suppose $c \in \text{Stab}(G)$ is a configuration on some graph $G = (V, E, s)$ with sink s and c' is the sandpile obtained from c by firing the sink. Then c is recurrent if and only if in the stabilization of c' each nonsink vertex fires exactly once, in which case $(c')^\circ = c$.*

Chapter 2

The Threshold Density Theorem

In this chapter we relate the Abelian Sandpile Model with the sink to that without. We then invoke a theorem called the "Markov Renewal Theorem," which will allow us to compute the limit of the threshold density, as the degree of our starting divisor tends towards $-\infty$. Our reference for this chapter is Perkinson and Corry [2017] and Levine [2015].

Definition 2.0.1. A *basic alive divisor* B on a graph $G = (V, E, s)$ consists of a recurrent $c \in S(G)$ plus an amount of sand on the sink equal to $\deg_G(s)$. That is, $B = c + \deg_G(s)s$. The set of basic alive divisors is denoted $B(G, s)$ or just $B(G)$.

A basic alive divisor is clearly alive, since by Dhar's Theorem, firing the sink causes every nonsink vertex to fire once and return to the original divisor. The following number gives the expected amount of sand on any given vertex at the time it becomes alive, under the assumption that reaching any basic alive divisor at the threshold is equally likely.

Definition 2.0.2. The *stationary density* of a graph $G = (V, E, s)$ is

$$\zeta_{st} = \frac{1}{|S(G)|} \sum_{B \in B(G)} \frac{\deg(B)}{|V|} = \frac{\deg_G(s)}{|V|} + \frac{1}{|S(G)|} \sum_{c \in S(G)} \frac{\deg(c)}{|V|}.$$

Two Examples

The Diamond Graph

We previously found all the recurrents of the diamond graph for $s = v_4$ and $s = v_2$. We will calculate the stationary density for both and show that they are the same.

$$S(G, v_4) = \{(1, 2, 2), (0, 2, 2), (0, 1, 2), (1, 2, 0), (1, 0, 2), (0, 2, 1), (1, 1, 2), (1, 2, 1)\}$$

$$S(G, v_2) = \{(1, 2, 1), (0, 1, 1), (0, 2, 1), (1, 0, 1), (1, 1, 1), (1, 2, 0), (0, 2, 0), (1, 1, 0)\}$$

$$\begin{aligned}\zeta_{st,v_1} &= \frac{\deg_G(v_1)}{|V|} + \frac{1}{|S(G)|}\sum_{c\in S(G)}\frac{\deg(c)}{|V|} = \frac{2}{4} + \frac{1}{8}\cdot\frac{5+4+3+3+3+3+4+4}{4} \\ &= \frac{1}{2} + \frac{29}{32} = \frac{45}{32} \\ \zeta_{st,v_2} &= \frac{3}{4} + \frac{1}{8}\cdot\frac{4+2+3+2+3+3+2+2}{4} = \frac{3}{4} + \frac{21}{32} = \frac{45}{32} = \zeta_{st,v_1}.\end{aligned}$$

It turns out that for an undirected graph, the stationary density is always independent of the choice of sink vertex.

Combining Two Graphs

The nature of the stationary density, being dependent only on the recurrents of a graph, makes it easy to examine the concatenation of two graphs, connected at a single point.

Definition 2.0.3. Let $G = (V_G, E_G)$ and $H = (V_H, E_H)$ both be graphs. The *concatenation* of G and H at $(v, w) \in V_G \times V_H$ is the graph $G_v \vee_w H$ formed from the disjoint union of G and H by identifying the vertices v and w.

For example, combining the diamond graph with the complete graph K_4 by combining v_3 and w_2 would look like,

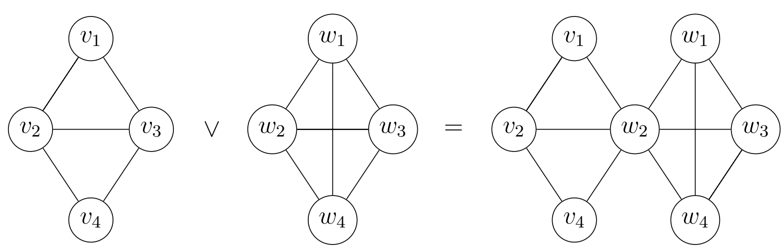

Since the stationary density of a graph is defined as,

$$\zeta_{st} = \frac{\deg_G(s)}{|V|} + \frac{1}{|S(G)|}\sum_{c\in S(G)}\frac{\deg(c)}{|V|},$$

it will suffice to find all the recurrents on $G_s \vee_t H$ to find its stationary density. To do so we will examine configurations on $G_s \vee H_t$.

Definition 2.0.4. For a graph $G = (V_G, E_G, s)$ with sink s and a graph $H = (V_H, E_H, t)$ with sink t, a configuration on $G_s \vee_t H$ will be uniquely given by $a = (a_G, a_H)$, where $a_G \in \text{Config}(G, s)$ and $a_H \in \text{Config}(H, t)$.

Since the two graphs are divided by a sink, there is no way for any sand to go from one side to the other. This means addition on the respective halves will function as it does for $\text{Config}(G, s)$ and $\text{Config}(H, t)$, respectively.

Proposition 2.0.1. *Suppose $G \vee_{s,t} H$ is the concatenation of graphs G and H, and $a = (a_G, a_H), b = (b_G, b_H) \in \mathrm{Config}(G_s \vee_t H)$, where $a_G, b_G \in \mathrm{Config}(G)$ and $a_H, b_H \in \mathrm{Config}(H)$. Then,*

$$a + b = (a_G + b_G, a_H + b_H)$$

and further,

$$(a + b)^\circ = ((a_G + b_G)^\circ, (a_H + b_H)^\circ).$$

Proposition 2.0.2. *The recurrents on $G_s \vee_t H$ will simply be all combinations (c_G, c_H), where c_G is a recurrent on G with respect to sink $s = t$ and c_H a recurrent on H with respect to sink t.*

Proof. Suppose $c = (c_G, c_H)$ is a recurrent on $G_s \vee_t H$. We claim that c_G and c_H are recurrent, that is, $c \in S(G)$ and $c_H \in S(H)$. First note that since c is stable, so are c_G and c_H. Now, take $a_G \in \mathrm{Config}(G)$ and $a_H \in \mathrm{Config}(H)$, and define $a = (a_G, a_H)$. Since c is recurrent, there exists $b = (b_G, b_H) \in \mathrm{Config}(G_s \vee_t H)$ with $b \geq 0$ such that

$$(a + b)^\circ = ((a_G + b_G)^\circ, (a_H + b_H)^\circ) = (c_G, c_H) = c.$$

So, for any configurations a_G and a_H, we have configurations $b_G \geq 0$ and $b_H \geq 0$ such that $(a_G + b_G)^\circ = c_G$ and $(a_H + b_H)^\circ = c_H$. This proves that if $c \in S(G_s \vee_t H)$, then $c_G \in S(G)$ and $c_H \in S(H)$.

Now we examine the converse. Suppose $c_G \in S(G)$ and $c_H \in S(H)$. We would like to show $c = (c_G, c_H) \in S(G_s \vee_t H)$.

First note that since c_G and c_H are both stable sandpiles, (c_G, c_H) is also a stable sandpile.

Next, take $a = (a_G, a_H) \in \mathrm{Config}(G_s \vee_t H)$. Since c_G and c_H are recurrent, there exist $b_G \geq 0$ and $b_H \geq 0$ such that $(a_G + b_G)^\circ = c_G$ and $(a_H + b_H)^\circ = c_H$. Letting $b = (b_G, b_H)$, we have $b \geq 0$ and $(a+b)^\circ = ((a_G+b_G)^\circ, (a_H+b_H)^\circ) = (c_G, c_H)$. Hence, c is recurrent. □

Corollary 3.1. The stationary density of $G_s \vee_t H$, the concatenation of the graphs at s, t is given by

$$\zeta_{st}(G_s \vee_t H) = \frac{\zeta_{st}(G)|V_G| + \zeta_{st}(H)|V_H|}{|V_G| + |V_H| - 1}.$$

Proof. Since the recurrents on $G \vee_{s,t} H$ are all the combinations (c_G, c_H) where $c_G \in S(G)$ and $c_H \in S(H)$, the size of $|S(G_s \vee_t H)|$ will be $|S(G)||S(H)|$, and the size of the vertex set will be $|V_G| + |V_H| - 1$, since one pair of vertices is being combined.

This means

$$\begin{aligned}
\zeta_{st} &= \frac{\deg_G(s)}{|V|} + \frac{1}{|S(G)|}\sum_{c\in S(G)}\frac{\deg(c)}{|V|} \\
&= \frac{\deg_G(s)+\deg_H(t)}{|V_G|+|V_H|-1} + \frac{1}{|S(G)||S(H)|}\sum_{c_G\in S(G)}\sum_{c_H\in S(H)}\frac{\deg(c_G)+\deg(c_H)}{|V_G|+|V_H|-1} \\
&= \frac{\deg_G(s)+\deg_H(t)}{|V_G|+|V_H|-1} + \frac{1}{|S(G)|}\sum_{c_G\in S(G)}\frac{\deg(c_G)}{|V_G|+|V_H|-1} \\
&+ \frac{1}{|S(H)|}\sum_{c_H\in S(H)}\frac{\deg(c_H)}{|V_G|+|V_H|-1} \\
&= \frac{1}{|V_G|+|V_H|-1}\Big(\deg_G(s) + \frac{1}{|S(G)|}\sum_{c_G\in S(G)}\deg(c_G) + \deg_H(t) \\
&+ \frac{1}{|S(H)|}\sum_{c_H\in S(H)}\deg(c_H)\Big) \\
&= \frac{\zeta_{st}(G)|V_G| + \zeta_{st}(H)|V_H|}{|V_G|+|V_H|-1}.
\end{aligned}$$

□

So, for the concatenation of a graph, the stationary density will be closely related to the sum of their respective stationary densities, adjusted by a factor of the ratio of the sizes of the vertex sets. Additionally, since it turns out the stationary density of the concatenation is independent of where we choose to combine the graphs, we can even find the threshold density of the concatenation of any number of graphs.

Corollary 3.2. The stationary density of the concatenation of a series of graphs, $G_1, G_2 \cdots G_n$, with respective stationary densities $\zeta_{st}(G_1), \zeta_{st}(G_2)\cdots\zeta_{st}(G_n)$ is given by

$$\zeta_{st}(\vee_{i=1}^n G_i) = \frac{\sum\limits_{i=1}^n \zeta_{st}(G_i)|V_{G_i}|}{\sum\limits_{i=1}^n |V_{G_i}| - n + 1}.$$

Proof. We will prove this by induction. The first case, regarding one concatenation, we just proved. Now, suppose that the concatenation of some n graphs is given by $\zeta_{st}(G) = \frac{\sum_{i=1}^n \zeta_{st}(G_i)|V_{G_i}|}{\sum_{i=1}^n |V_{G_i}|-n+1}$. Since each concatenation reduces the size of the total vertex set by 1, we have $|V_G| = \sum\limits_{i=1}^n |V_{G_i}| - n + 1$, since for n graphs we have $n-1$ concatenations.

We will combine this graph with another, G_{n+1}, which has stationary density $\zeta_{st}(G_{n+1})$. By the previous corollary, the stationary density is given by,

$$\zeta_{st}(G\vee H) = \frac{\zeta_{st}(G)|V_G| + \zeta_{st}(H)|V_H|}{|V_G|+|V_H|-1}$$

So we have,

$$\begin{aligned}
\zeta_{st}(G \vee G_{n+1}) &= \frac{\zeta_{st}(G)|V_G| + \zeta_{st}(G_{n+1})|V_{G_{n+1}}|}{|V_G| + |V_{G_{n+1}}| - 1} \\
&= \frac{\frac{\sum_{i=1}^{n} \zeta_{st}(G_i)|V_{G_i}|}{\sum_{i=1}^{n} |V_{G_i}| - n + 1}\left(\sum_{i=1}^{n} |V_{G_i}| - n + 1\right) + \zeta_{st}(G_{n+1})|V_{G_{n+1}}|}{\sum_{i=1}^{n} |V_{G_i}| - n + 1 + |V_{G_{n+1}}| - 1} \\
&= \frac{\sum_{i=1}^{n} \zeta_{st}(G_i)|V_{G_i}| + \zeta_{st}(G_{n+1})|V_{G_{n+1}}|}{\sum_{i=1}^{n} |V_{G_i}| - n + 1 + |V_{G_{n+1}}| - 1} \\
&= \frac{\sum_{i=1}^{n+1} \zeta_{st}(G_i)|V_{G_i}|}{\sum_{i=1}^{n+1} |V_{G_i}| - (n+1) + 1} \\
&= \zeta_{st}(\vee_{i=1}^{n+1} G_i).
\end{aligned}$$

□

Basic alive divisors allow us to think of a divisor on the sinkless model in terms of the sink model and the stationary density as a quantity related to the threshold density. We may drive this relation further: any divisor can be represented uniquely in terms of a basic alive divisor, along with two other quantities:

Theorem 4. *Let $G = (V, E)$, and $D \in \mathrm{Div}(G)$ a divisor. For any choice of sink $s \in V$ we can uniquely represent D as a triple, (B, m, σ), where B is a basic alive divisor $c + \deg_s(s)$, $m \in \mathbb{Z}$, and $\sigma : V \to \mathbb{Z}$ a firing script with $\sigma(s) = 0$, such that,*

$$D = B + ms - \mathcal{L}\sigma.$$

Corollary 4.1. A divisor $D \in \mathrm{Div}(G)$, written as $D = B + ms + \mathcal{L}\sigma$, will be stable if and only if $m < 0$.

Now, we have a preliminary way of relating the threshold density to the stationary density. However, in the sinkless model, it is not the case that every basic alive divisor will be equally likely for representing a threshold state, since threshold is reached after a finite number of steps, which will not allow for sufficient mixing.

The other problem with our comparison of models is that in the sinkless model, not every threshold state is reached with equal probability. We will address this in the next chapter, where we will compute the limiting distribution of the sinkless Abelian Sandpile Model as the degree of the initial divisor goes to negative infinity. We will find that, though basic alive divisors are not reached with equal probability in the limit, basic alive divisors of any given degree *are*. This will be enough to prove the convergence of the threshold density to the stationary density.

Non-Basic Threshold Divisors

Let D be a threshold divisor and suppose v is the epicenter vertex. Thus, the addition of a grain on s caused the chain to reach an unstabilizable state D. Then D is linearly equivalent to $c + deg_G(s)s$. However, it may be the case that $D \neq c + \deg_G(s)s$ for any choice of s and $c \in S(G, s)$.

For example, consider the graph G consisting of a line segment with two vertices attached to a triangle:

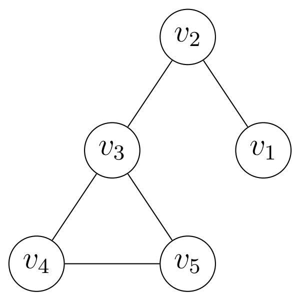

For the sink model, this graph has three recurrents for any of the five choices of sink. However, in the sinkless model, it has 17 threshold states—this means that two of these cannot be represented as $c + \deg_G(s)s$. Namely, the threshold states $B_1 = (0, 0, 2, 2, 1)$ and $B_2 = (0, 0, 2, 1, 2)$ have no such representations: $B_1(v_4) = deg_G(v_4)$ and $B_2(v_5) = deg_G(v_5)$, the configurations with respect to v_4 and v_5, respectively being the sink are not recurrent.

2.1 The Stationary Density Theorem

We will map out a way to prove that the threshold density converges to the stationary density as $D_0 \to \infty$. Since we have a way of uniquely decomposing a divisor in terms of a triple, we will want to find a distribution which tells us the likelihood of reaching that exact triple at threshold. To do so, we will use something called the *Markov Renewal Theorem*, which gives a stationary distribution for a Markov chain once it reaches a certain cumulative "length." In our case, this length will be the cumulative amount of sand which has fallen down the sink, and we will find the distribution of c being the recurrent and v being the vertex at the time when this amount of sand on the sink overflows and causes the divisor to become alive. This limiting distribution will tell us that as the initial amount of sand on the sink (hence the degree of the original divisor) approaches $-\infty$, the limiting distribution of any given basic alive divisor being the threshold state does not become uniform, but the number of basic alive divisors reached with degree n *does* reach a uniform distribution, which means that the threshold density converges to the stationary density.

We next define the burst size, which tracks the amount of sand going into the sink for any firing, and allows us to examine the decomposition of $a_v D$.

Definition 2.1.1. For some basic alive divisor, $B = c + ms \in B(G, s)$, the *burst size*, $\beta_v(B)$ denotes the amount of sand which flows down the sink in the stabilization of

c from the addition of v:

$$\beta_v(B) = \beta_v(c) := \deg(\tilde{c}) - \deg(c) + 1$$

where $\tilde{c}$ is the unique recurrent such that $(\tilde{c} + v)^\circ = c$.

We will also need the *odometer function*, $\mathrm{odo}(c)$, which assigns to each configuration c the firing script σ of its stabilization. This means the odometer function will record the number of times each vertex fires in the stabilization of $c \to c^\circ$. It follows that the basic decomposition of $a_v D$ will be,

$$a_v D = (c+v)^\circ + (m + \deg_G(s) + \beta)s - \mathcal{L}(\sigma - \mathrm{odo}(c+v) + \mu - \mu(s)).$$

Here, β refers to $\beta_v((c+v)^\circ)$, μ is the stabilization script for $D + v$, and $\mu(s)$ is the constant script, $\mu(s) \cdot \vec{1}$.

Finally, we state the Threshold Density Theorem.

Theorem 5. *Threshold Density Theorem Let $G = (V, E)$ be a sinkless Abelian Sandpile. Then, as $\deg(D_0) \to -\infty$,*

$$\zeta_{D_0} \to \zeta_{st}.$$

In order to prove this theorem, we will examine a related Markov chain which keeps track of the vertices chosen at each point in time. This Markov chain will be endowed with a "length" function based on the burst sizes, which will keep track of the amount of sand on the sink, and we will want to figure out when the chain surpasses a certain length. The key technical result upon which we ultimately rely is given by the following, the Markov Renewal Theorem, a general proof of which can be found in Kesten [1974].

Theorem 6. *(Markov Renewal Theorem) Consider an irreducible Markov chain, $(\Omega, P, (X_t))$ and its set of* edges, *$\mathcal{E} = \{(x,y)\colon P(x,y) > 0\}$. Let $l\colon \mathcal{E} \to \mathbb{N}$ be a* length *function, and assume the chain is aperiodic with respect to l: $\gcd\{n \geq 1\colon P^n(x,x) > 0\} = 1$. Let π be the stationary distribution. For each $n \in \mathbb{N}$, the random time is defined as $\tau_n = \min\{t\colon \lambda_t \geq n\}$, for the cumulative length function, $\lambda_t := \sum_{j=1}^{t} l(X_{j-1}, X_j)$. Then, as $n \to \infty$, for any states $x_0, x, y \in \Omega$ and $m \in \mathbb{N}$ we will have:*

$$\mathbb{P}_{x_0}\{(X_{\tau_n - 1}, X_{\tau_n}, \lambda_{\tau_n} - n) = (x, y, m)\} \to \frac{1}{Z}\pi(x)P(x,y)\boldsymbol{1}\{0 \leq m \leq l(x,y)\}.$$

The normalization constant is $Z = \sum_{(x,y)\in\mathcal{E}} \pi(x)P(x,y)l(x,y)$.

This tells us that, starting from some initial position (x_0, y_0) on our Markov chain, as the required length approaches infinity, the likelihood of reaching that required length at some particular (x, y) will converge to the stationary distribution, $\pi(x)P(x, y)$, as long as the excess is not more than the length of the last edge traveled (in which case the chain would have finished in the prior step).

Theorem 7. *For a graph G, a probability distribution $\alpha\colon V \to [0,1]$, with $\alpha(v) > 0$ for all $v \in V$, let D_t be the Markov chain for the sinkless Abelian Sandpile, with threshold time $\tau = \tau(D_0)$ and epicenter v_τ. With respect to some sink s, our divisor may be written as*

$$D_\tau \sim B_\tau + m_\tau s - \mathcal{L}\sigma_\tau.$$

Then, as $\deg(D_0) \to -\infty$*, the joint probability of the triple (v_τ, B_τ, m_τ) converges:*

$$\mathbb{P}_{D_0}\{(v_\tau, B_\tau, m_\tau) = (v, B, m)\} \to \begin{cases} \frac{\alpha(v)}{|S(G)|}, & \text{if } 0 \le m < \beta_v(B) \\ 0, & \text{otherwise.} \end{cases}$$

Proof. We begin by defining a new Markov chain:

$$\begin{aligned} &\Omega = V \times S(G); \\ &P((v', c'), (v, c)) = \alpha(v)\mathbf{1}\{c = (c' + v)^\circ\}; \\ &X_t = (v_t, c_t). \end{aligned}$$

The above chain is nearly identical to our Abelian Sandpile Model with sink, except that we keep track of which vertices we fire. Since we are restricting ourselves to the recurrents, the Markov chain will be irreducible, and so have stationary distribution $\pi(v, c) = \frac{\alpha(v)}{|S(G)|}$. Since we will reach a basic alive divisor once the amount of sand on the sink exceeds the degree of the sink, we will want to know the expected time, after which we have $m > 0$, as well as how much "excess" there will be. Conveniently, this is just what the Markov Renewal Theorem allows us to do.

In our case, the initial state will be some recurrent c, paired with the vertex to be fired, and the length the total amount of sand which will be required to fill up the sink.

We define the length function:

$$l((v', c'), (v, c)) = \beta_v(c) = \deg(c') - \deg(c) + 1.$$

That is, the amount of sand that flows into the sink going from c' to c by firing v. We note that $l((s, c), (s, c))$ will trivially be zero, and will also mean that our length function is aperiodic (since we can always drop sand on the sink with no effect.) The cumulative length function will be:

$$\lambda_t := \sum_{j=1}^{t} l(X_{j-1}, X_j) = \sum_{j=1}^{t} l((c_{j-1}, v_{j-1}), (c_j, v_j)).$$

The amount of excess sand on the sink, will be denoted m, as it was in our decomposition. Recall that the divisor will be stable if and only if $m_t < 0$, so our threshold τ will naturally be $\tau := \min\{t \ge 0\colon m_t \ge 0\}$. The length and m_t will thus be intimately related: $m_t = m_0 + \lambda_t$. To show this, we recall the basic decomposition, $D_t = B_t + m_t s + \mathcal{L}\sigma_t$, and the decomposition of $a_v D$ to find that for m,

$$m_t = m_{t-1} + \beta_{v_t}(c_t) = m_{t-1} + l((c_{t-1}, v_{t-1}), (c, v)).$$

This means we can also write τ as $\tau = \min\{t \geq 0\colon \lambda_t \geq -m_0\}$. We now have everything in place to invoke the Markov renewal theorem: a set of states (c, v) endowed with a length function, a variable target $-m_0$, a random time τ, and a stationary distribution $\pi(x)$. So, we find that,

$$\begin{aligned}&\mathbb{P}_{x_0}\{((v_{\tau-1}, c_{\tau-1}), (v_\tau, c_\tau), m_\tau) = ((v', c'), (v, c), m)\}\\ &\to \frac{1}{Z}\frac{\alpha(v')}{|S(G)|}\alpha(v)\mathbf{1}\{c = (c' + v)^\circ, 0 \leq m < \beta_v(B)\}.\end{aligned}$$

The normalization constant Z turns out, conveniently, to be 1. If we sum over all possible (v', c') which would yield the (v, c) we seek, we find that, as $m_0 \to -\infty$,

$$\mathbb{P}_{D_0}\{(v_\tau, B_\tau, m_\tau) = (v, B, m)\} \to \frac{\alpha(v)}{S(G)}\mathbf{1}\{0 \leq m \leq \beta_v(B)\}.$$

Because of the uniqueness of our decomposition, $m_0 \to -\infty$ will necessarily mean $\deg(D) \to -\infty$. So, we conclude our proof, to find that as the degree of the divisor approaches negative infinity, the likelihood of the epicenter being some v, the final recurrent being c, and the excess being m approach the stationary distribution biased by α. □

This result will give us several convenient corollaries. Summing first over all possible m and v, we get the following:

Corollary 7.1. For all basic alive divisors B, as $\deg(D_0) \to -\infty$,

$$\mathbb{P}_{D_0}\{B_\tau = B\} \to \frac{1}{|S(G)|}\sum_{v \in V}\alpha(v)\beta_v(B).$$

That is, the limiting distribution of a basic alive divisor at the threshold is determined by the distribution α and the burst size β. Again, we had been expecting that in order for the threshold density to approach the stationary density, we would want each basic alive divisor to occur with equal frequency. However, as we will soon see, all we actually need is the probability for one of those basic alive divisors to be of degree n to be equal for the Threshold Density Theorem to hold.

Next we consider the distribution for when the epicenter is located at the sink. Since D is stabilizable if and only if $m < 0$, we know that the epicenter will necessarily be at the sink if $m_\tau = 0$. So,

$$\mathbb{P}_{D_0}\{(v_\tau = s, B_\tau = B\} = \mathbb{P}_{D_0}\{((v_\tau, B_\tau, m) = (s, B, 0)\} \to \frac{\alpha(s)}{|S(G)|}.$$

Conveniently, since the choice of sink is arbitrary, we may pick any given sink and sum over all the recurrents with respect to it, to find the limiting distribution for any given vertex being the epicenter:

Corollary 7.2. For any $v \in V$, as $\deg(D) \to -\infty$,

$$\mathbb{P}_{D_0}\{v_\tau = v\} \to \alpha(v).$$

Finally, in order to prove the actual Threshold Density Theorem, we look at the limiting distribution of the size of a divisor at the moment it becomes critical. We find,

Corollary 7.3. For any $n \geq 0$ and sink s, as $n \to -\infty$,

$$\mathbb{P}_{D_0}\{\deg(D\tau) = n\} \to \frac{|B_s\colon\ \deg(B) = n|}{|S(G)|}.$$

In order to prove the above, we must consider the breakdown of the divisor with respect to any possible sink. So, we will have,

$$D_t = B_t^s + m_t^s s - L(\sigma_t^s).$$

In order for the sink to be the epicenter, m_τ must be zero. So $\deg(D_\tau) = \deg(B_\tau^s)$. Setting $v_\tau = s$ and invoking the distribution for when the epicenter is located at the sink, we find,

$$\mathbb{P}_{D_0}\{\deg(D\tau) = n\} = \sum_{s \in V} \mathbb{P}_{D_0}\{v_\tau = s, \deg(D_\tau) = n\} = \sum_{s \in V} \mathbb{P}_{D_0}\{v_\tau = s, \deg(B_\tau^s) = n\}.$$

As $\deg(D) \to -\infty$, this sum converges to,

$$\sum_{s \in V} \frac{\alpha(s)}{|S(G)|} |\{B_s\colon\ \deg(B) = n\}|.$$

We invoke Merino's theorem from Merino López [1997], and find that the number of elements in B_s with degree n is independent of the choice of sink. This proves the threshold density theorem:

$$\begin{aligned}
\zeta_\tau(D_0) :&= \mathbb{E}_{D_0} \frac{\deg(D_\tau)}{|V|} \\
&= \sum_{n \geq 0} \mathbb{P}_{D_0}\{\deg(D\tau) = n\} \frac{n}{|V|} \\
&\to \sum_{n \geq 0} \mathbb{P}_{D_0} \frac{|\{B_s\colon\ \deg(B) = n\}|}{|S(G)|} \frac{n}{|V|} \\
&= \frac{1}{|S(G)|} \sum_{B \in B(G)} \frac{\deg(B)}{|V|} \\
&= \zeta_{st}. \quad \square
\end{aligned}$$

Chapter 3

Threshold Density of the Banana Graph

The *banana graph*, B_n consists of two vertices, connected by some n number of edges. For example, B_3 is,

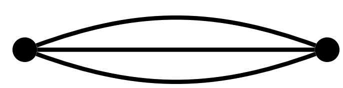

Figure 3.1: The Banana Graph.

To find the threshold density we must find the probability that the divisor becomes alive after some number of grains of sand have been dropped on it. We will suppose that there is an equal probability of dropping the grain at either vertex.

3.1 From the Zero Divisor

We begin with the zero divisor, $(0, 0)$. The first grain will fall on the first or second vertex, leading to the divisor $(1, 0)$ or $(0, 1)$. The next grain of sand will fall on either vertex, which will give us $(2, 0)$, $(1, 1)$, or $(0, 2)$, and then $(3, 0), (1, 2), (2, 1)$, and $(0, 3)$, and so on. Once either vertex reaches n grains, it will fire, sending n grains of sand to the other, which will fire, and so the divisor will be alive. So, in the case of the banana graph, the threshold states will be given by $(n, 0), (n, 1) \cdots (n, n-1)$ and $(0, n), (1, n) \cdots (n-1, n)$. As example, the possible ways B_3 can evolve are given here:

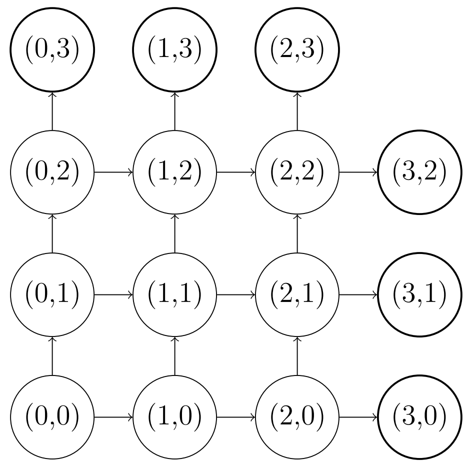

Proposition 3.1.1. *Let B_n be the sinkless Abelian Sandpile on the banana graph B_n. The probability of reaching some non-threshold state (i, j) from $(0, 0)$ after $i+j$ grains of sand have been dropped is given by*

$$P(i,j) = \frac{1}{2^{i+j}}\binom{i+j}{i} = \frac{1}{2^{i+j}}\binom{i+j}{j}.$$

Proof. Suppose $\lambda(i, j)$ represents the number of possible paths from the divisor $(0, 0)$ to the non-threshold divisor (i, j). Since (i, j) can be immediately preceded by $(i-1, j)$ or $(i, j-1)$ or both, we can say $\lambda(i, j) = \lambda(i-1, j) + \lambda(i, j-1)$, where $\lambda(0, 0) = 1$ and $\lambda(i, j) = 0$ for $i < 0$ or $j < 0$.

By induction, we see that $\lambda(i, j) = \binom{i+j}{i} = \binom{i+j}{j}$:

$$\lambda(i,j) = \binom{i+j}{i} = \binom{i+j-1}{i-1} + \binom{i+j-1}{j-1} = \lambda(i-1,j) + \lambda(i,j-1).$$

The base for induction is $\binom{0}{0} = 1$ and $\binom{i+j}{i}$ is 0 if i or $j < 0$.

Since there are two vertices to choose from at each moment in time, there will be a total of 2^{i+j} paths to the divisors of weight $i + j$.

So, the probability that the sandpile reaches the divisor (i, j) after $i + j$ grains of sand have been dropped is given by the number of paths to (i, j) divided by the total number of paths: $\frac{1}{2^{i+j}}\binom{i+j}{i}$. □

From this, we can deduce the threshold density.

Corollary 7.4. The threshold density of a banana graph B_n, starting at the zero divisor $(0, 0)$ is given by

$$\zeta_{\vec{0}} = \frac{1}{2^n}\sum_{k=0}^{n-1}\frac{n+k}{2^k}\binom{n+k-1}{k} = n\left(1 - \frac{1}{4^n}\binom{2n}{n}\right).$$

Proof. For $0 \le k \le n-1$, the states which immediately precede the threshold states are $(n-1, k)$ and $(k, n-1)$. That is, $(n-1, 0), (n-1, 1), \cdots, (n-1, n-1)$ and $(0, n-1), (1, n-1), \cdots, (n-1, n-1)$.

The probability of reaching any of these states is $P(n-1,k) = P(k,n-1) = \frac{1}{2^{n+k-1}}\binom{n+k-1}{k}$.

Since, once one of these states is reached, the probability of reaching the threshold states (n,k) or (k,n) is $\frac{1}{2}$, we may say $P(n,k) = P(k,n) = \frac{1}{2}P(n-1,k) = \frac{1}{2}P(k,n-1)$.

So,

$$P(n,k) = P(k,n) = \frac{1}{2} \cdot \frac{1}{2^{n+k-1}}\binom{n+k-1}{k} = \frac{1}{2^{n+k}}\binom{n+k-1}{k}.$$

Recall that the threshold density is given by $\zeta_\tau(D_0) := \mathbb{E}_{D_0}\frac{\deg(D_\tau)}{|V|}$. In our case, $|V| = 2$ and $\deg(D_\tau) = n+k$, so,

$$\zeta_\tau(D_0) = \mathbb{E}_{D_0}\frac{\deg(D_\tau)}{|V|} = \sum_{k=0}^{n-1}(P(k,n)+P(n,k))\frac{n+k}{2} = \frac{1}{2^n}\sum_{k=0}^{n-1}\frac{n+k}{2^k}\binom{n+k-1}{k}.$$

We have proved the first half of the identity.

Next, we prove the second equality. An established identity (5.20) in Graham et al. [1994] states,

$$\sum_{k=0}^{n}\frac{1}{2^k}\binom{n+k}{k} = 2^n.$$

We split the final nth term off the sum on the left and move it to the right side to get,

$$\sum_{k=0}^{n-1}\frac{1}{2^k}\binom{n+k}{k} = 2^n - \frac{1}{2^n}\binom{2n}{n}.$$

Now we use the combinatorial identity $\binom{n+k}{k} = \frac{n+k}{k}\binom{n+k-1}{k}$:

$$\sum_{k=0}^{n-1}\frac{1}{2^k}\binom{n+k}{k} = \sum_{k=0}^{n-1}\frac{1}{2^k}\binom{n+k-1}{k}\frac{n+k}{n} = 2^n - \frac{1}{2^n}\binom{2n}{n}.$$

Multiplying both sides by $\frac{n}{2^n}$, we get our result:

$$\frac{1}{2^n}\sum_{k=0}^{n-1}\frac{n+k}{2^k}\binom{n+k-1}{k} = n\left(1-\frac{1}{4^n}\binom{2n}{n}\right) = \zeta_{D_0}.$$

□

3.2 From an Arbitrary Divisor

We will now examine the Abelian Sandpile model, starting from some arbitrary (α,β). We want to prove that as $\deg(D) = \alpha+\beta$ approaches $-\infty$, the threshold density of the banana graph will approach the stationary density, as per the Threshold Density

Theorem. In this proof we will find some good intuition for the Threshold Density Theorem itself and its proof.

Recall that the stationary density is given by,

$$\zeta_{st} = \frac{1}{|S(G)|} \sum_{B \in B(G)} \frac{\deg(B)}{|V|}.$$

On the banana graph, picking either vertex as the sink, we will have exactly n elements in $S(G)$, namely, $(s,0),(s,1),\cdots,(s,n-1)$. Likewise, the elements of $B(G)$ will just be $(n,0),(n,1),\cdots,(n,n-1)$. In our case $|V|=2$. So, we find,

$$\zeta_{st} = \frac{1}{n}\sum_{i=0}^{n-1}\frac{n+i}{2} = \frac{1}{2n}\left(\textstyle\sum_{i=0}^{n-1} n + \sum_{i=0}^{n-1} i\right) = \frac{1}{2n}\left(n^2 + \frac{n(n-1)}{2}\right) = \frac{3n-1}{4}.$$

In order to examine the behavior of the banana graph at increasingly negative (α,β), we will first make several remarks. Recall that on our Sandpile model, we allowed negative amounts of sand to represent "holes" which could be filled in by sand, and then function normally. We can imagine, then, starting at some (α,β), going either to $(\alpha+1,\beta)$ or $(\alpha,\beta+1)$, and then dropping another grain to get either $(\alpha+2,\beta),(\alpha+1,\beta+1)$, or $(\alpha,\beta+2)$, and so on.

We can say that, at least until one of the vertices accumulates n grains of sand, the model will evolve as it did starting from zero—that is, in a binomial fashion. Once the divisor reaches degree n, we will begin getting to threshold states.

Proposition 3.2.1. *Let B_n be the sinkless Abelian Sandpile on the banana graph of weight n. The probability $p_{i,j}$ of reaching some divisor (i,j), with $i,j \geq 0$ and $i+j \leq n$, from a starting divisor (α,β), after $(i+j)-(\alpha+\beta)$ grains of sand have been dropped is given by*

$$p_{i,j} = \sum_{r \in \mathbb{Z}} \frac{2^{\alpha+\beta}}{2^{i+j}} \binom{(i+j)-(\alpha+\beta)}{i+rn-\alpha}.$$

Proof. Let us momentarily disregard stabilization. Then, we can say that the probability, starting from (α,β) of reaching any (i,j), where $i+j \leq n$ is given by,

$$\frac{2^{\alpha+\beta}}{2^{i+j}} \binom{(i+j)-(\alpha+\beta)}{i-\alpha}.$$

We are operating on the same principle as before, starting from $(0,0)$, except offset by (α,β).

Consider, the case of B_3, starting at $(-5,-2)$. If we continually drop sand on the second vertex, we will eventually reach $(-5,3)$. Its stabilization is $(-2,0)$, so $(-5,3) \sim (-2,0)$. Continuing to add sand to the second vertex, we reach $(-5,6) \sim (-2,3) \sim (1,0)$. Since $(-5,6)$ is similar to a stable divisor, it will behave identically with regards to reaching threshold states and we may apply our method from the previous section.

We re-introduce the notion of divisor classes. Since a firing will move n grains of sand from one vertex to the other, we may represent a firing as $r(n, -n)$, where $r \in \mathbb{Z}$. A positive r corresponds to firing v_2 a total of r times and a negative r corresponds to firing v_1 a total of r times. The set of divisor classes, then, will be given by,

$$[(i, j)] = \{(i, j) + r(n, -n) \colon r \in \mathbb{Z}\}.$$

However, not all the divisors in a divisor class will be reachable. Consider the diagram below:

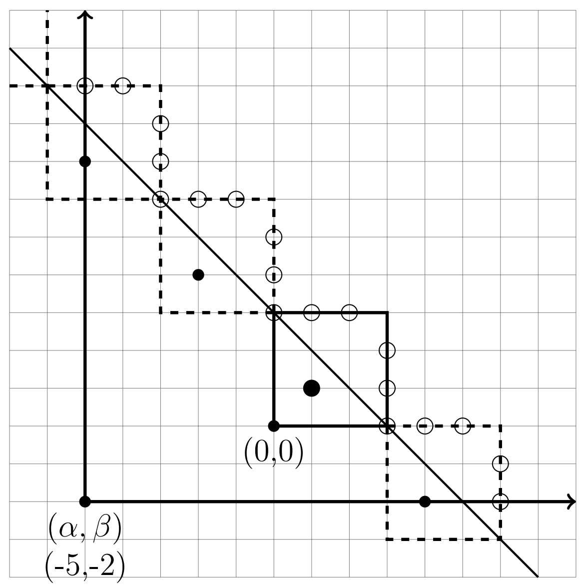

It depicts the possible evolution of B_3, starting at $D_0 = (-5, -2)$, with rightward movement corresponding to dropping a grain on the first vertex and upward movement corresponding to dropping a grain on vertex two. The dashed boxes represent sets of equivalence classes, corresponding to the set of stable divisors, in the un-dashed box. Being located at some point in one of these boxes corresponds to linear equivalence with any correspondingly located point. Thus the four black dots correspond to linearly equivalent divisors, with the larger one being the stable one. The circles on the boundary of the un-dashed box represent the threshold divisors, while the circles on the boundary of the dashed boxes represent linearly equivalent divisors to the corresponding ones on the main box.

The set of possible divisors reachable from (α, β) through the addition of sand are those enclosed by the vertical and horizontal lines emanating from (α, β). The diagonal line represents the set of divisors with degree n.

The probability of reaching some (i, j) with $i, j \geq 0$ and $i + j \leq n$ through the evolution of the chain will then be the sum of the probabilities of reaching all the possible linearly equivalent divisors.

So, for all permissible r, the probabilities of reaching $(i, j) + r(n, -n) = (i + rn, j - rn)$, will be given by

$$\sum_{r\in\mathbb{Z}} \frac{2^{\alpha+\beta}}{2^{i+j}} \binom{(i+j)-(\alpha+\beta)}{i+rn-\alpha}.$$

Note that when r is too large or small, $\binom{(i+j)-(\alpha+\beta)}{i+rn-\alpha}$ will be zero. So summing over all possible r, we find the probability of reaching any divisor (i, j) in terms of reaching any possible linearly equivalent divisor. □

We now use series multisection, as established in Riordan [1968], which will allow us to compute the limit in the Threshold Density Theorem.

Proposition 3.2.2. *Fix some $q, d, n \in \mathbb{N}$. Then,*

$$\binom{q}{d} + \binom{q}{d+n} + \binom{q}{d+2n} + \ldots = \frac{1}{n}\sum_{i=0}^{n-1}\Big(2\cos(\frac{\pi i}{n})\Big)^q \cos(\frac{\pi(q-2d)i}{n}).$$

Proof. For any $n \in \mathbb{N}$, and $\omega = e^{\frac{2i\pi}{n}}$ a primitive nth root of unity, we have,

$$\prod_{l=0}^{n-1}(x-\omega^l) = x^n - 1 = (x-1)(1+x+x^2+\cdots+x^{n-1}).$$

So, for any $k \in \mathbb{Z}$,

$$\sum_{l=0}^{n-1}\omega^{kl} = \begin{cases} n, & \text{if } k = 0 \mod n \\ 0, & \text{otherwise.}\end{cases}$$

For a formal power series, $F(x) = \sum_{k\geq 0} a^k x^k$, and fixing an offset constant $d \in \{0, 1\cdots n-1\}$, let

$$G(x) := \frac{1}{n}\sum_{l=0}^{n-1}\omega^{-dl}F(\omega^l x).$$

We will want to prove that,

$$G(x) = \sum_{m\geq 0} a_{d+mn}x^{d+mn} = a_d x^d + a_{d+n}x^{d+n} + a_{d+2n}x^{d+2n} + \cdots$$

The coefficient of x^k in $G(x)$ is,

$$\begin{aligned}
[x^k]G(x) &= [x]^k \frac{1}{n}\sum_{l=0}^{n-1}\omega^{-dl}F(\omega^l x) \\
&= \frac{1}{n}\sum_{l=0}^{n-1}\omega^{-dl}[x^k]F(\omega^l x) \\
&= \frac{1}{n}\sum_{l=0}^{n-1}\omega^{-dl}a_k\omega^{kl} \\
&= a_k\frac{1}{n}\sum_{l=0}^{n-1}\omega^{(k-d)l} \\
&= \begin{cases} a_k, & \text{if } k = d \mod n \\ 0, & \text{otherwise.}\end{cases}
\end{aligned}$$

So the coefficients of $G(x)$ are the coefficients of $F(x)$, starting at d and spaced apart by factors of n.

Now, let

$$F(x) = (1+x)^q = \binom{q}{0} + \binom{q}{1}x + \cdots = \sum_{k\geq 0}\binom{q}{k}x^k$$

and $G(x)$ be as it was above. Then, evaluating $G(x)$ at 1, we get,

$$\begin{aligned}
G(1) &= \frac{1}{n}\sum_{l=0}^{n-1}\omega^{-dl}F(\omega^l)\\
&= \frac{1}{n}\sum_{l=0}^{n-1}\omega^{-dl}(1+\omega^l)^q\\
&= \frac{1}{n}\sum_{l=0}^{n-1}e^{(\frac{2i\pi}{n})(-dl)}(1+e^{\frac{2i\pi l}{n}})^q\\
&= \frac{1}{n}\sum_{l=0}^{n-1}e^{(\frac{2i\pi}{n})(-dl)}((e^{\frac{\pi il}{n}}+e^{-\frac{\pi il}{n}})(e^{\frac{\pi il}{n}}))^q\\
&= \frac{1}{n}\sum_{l=0}^{n-1}e^{\frac{l\pi(q-2d)i}{n}}\left(2\cos(\frac{\pi l}{n})\right)^q
\end{aligned}$$

We know the sum is real, so we can replace the exponential with the real part:

$$\begin{aligned}
&= \frac{1}{n}\sum_{l=0}^{n-1}\cos(\frac{\pi(q-2d)l}{n})\left(2\cos(\frac{\pi l}{n})\right)^q\\
&= a_d + a_{d+n} + a_{d+2n} + \cdots\\
&= \binom{q}{d} + \binom{q}{d+n} + \binom{q}{d+2n} + \dots
\end{aligned}$$

□

As a result of this, we find a limit for the probability of reaching $(n-k,k)$.

Corollary 7.5. For a banana graph B_n , the probability $p_{n-k,k}$ of reaching some divisor, $(n-k,k)$ from initial divisor $D_0 = (\alpha,\beta)$ approaches $\frac{1}{n}$ as $\deg(D_0) \to -\infty$.

Proof. The probability of reaching a divisor $(n-k,k)$ is the sum of the probabilities of reaching a divisor linearly equivalent to it. We define $d \in \{0,1,\dots,n-1\}$ such that $d \equiv k - \alpha \mod n$ to be our offset coefficient. So, by our prior proposition,

$$\begin{aligned} p_{i,j} &= \sum_{r\in\mathbb{Z}} \frac{2^{\alpha+\beta}}{2^{i+j}} \binom{(i+j)-(\alpha+\beta)}{i+rn-\alpha} \implies \\ p_{n-k,k} &= \sum_{r\in\mathbb{Z}} \frac{2^{\alpha+\beta}}{2^{n}} \binom{n-(\alpha+\beta)}{d+rn} \\ &= \frac{2^{\alpha+\beta}}{2^{n}} \left(\binom{n-(\alpha+\beta)}{d} + \binom{n-(\alpha+\beta)}{d+n} + \binom{n-(\alpha+\beta)}{d+2n} + \dots \right). \end{aligned}$$

Given our identity from Riordan, we can rewrite our sum as

$$\begin{aligned} p_{n-k,k} &= \frac{1}{n2^{n-(\alpha+\beta)}} \sum_{i=0}^{n-1} \left(2\cos(\frac{\pi i}{n})\right)^{n-(\alpha+\beta)} \cos(\frac{\pi(n-(\alpha+\beta)-2d)i}{n})\Big) \\ &= \frac{1}{n} + \frac{1}{n}\sum_{i=1}^{n-1} \left(\cos(\frac{\pi i}{n})\right)^{n-(\alpha+\beta)} \cos(\frac{\pi(n-(\alpha+\beta)-2d)i}{n})\Big). \end{aligned}$$

Since $\left|\cos(\frac{\pi i}{n})\right| < 1$, as $n-(\alpha+\beta)$ approaches ∞, the second sum will go to zero and the total sum will approach,

$$p_{n-k,k} \to \frac{1}{n}.$$

□

Now, we will examine the behavior of the sandpile starting at $(n-k, k)$.

Proposition 3.2.3. *Let $p_{n-k,k}$ be the probability of reaching a divisor $(n-k, k)$ starting from (α, β). Then the probability $P(l)$ that the threshold weight will be $n+l$, where $1 \le l \le n-1$, is given by*

$$P(l) = \frac{1}{2^l}\sum_{k=0}^{l-1}\binom{l-1}{k}(p_{n-k-1,k+1} + p_{k+1,n-k-1}).$$

The probability $P(0)$ is given by $p_{n,0}$. Thus the expected threshold density, in terms of $P(l)$, is given by

$$\zeta_\tau = \sum_{l=0}^{n-1} P(l)\frac{n+l}{2}.$$

Proof. As we saw in the previous section, a threshold divisor, (n, l) or (l, n), where $0 \le l \le n-1$ and $l \in \mathbb{N}$ will be reached with a $\frac{1}{2}$ probability, only from a divisor $(n-1, l)$ or $(n, l-1)$. Any one of these pre-critical divisors will be reachable by exactly $\binom{l-1}{k}$ paths from some divisor $(n-k, k)$. So, the probability that (n, l) or (l, n) will be reached will be

$$P(l) = \frac{1}{2^l}\sum_{k=0}^{l-1}\binom{l-1}{k}p_{n-k-1,k+1} + \frac{1}{2^l}\sum_{k=0}^{l-1}\binom{l-1}{k}p_{k+1,n-k-1}$$
$$= \frac{1}{2^l}\sum_{k=0}^{l-1}\binom{l-1}{k}(p_{n-k-1,k+1} + p_{k+1,n-k-1}).$$

$P(0)$ is given to be $p_{n,0}$ since $(0, n)$ and $(n, 0)$ are threshold states. □

Since we know that the probability $p_{n-k,k}$ of reaching each $(n-k, k)$ converges to $\frac{1}{n}$ as $\deg(D_0) \to -\infty$,

$$P(l) = \frac{1}{2^l}\sum_{k=0}^{l-1}\binom{l-1}{k}(p_{n-k-1,k+1} + p_{k+1,n-k-1})$$
$$\to \frac{1}{2^l}\sum_{k=0}^{l-1}\binom{l-1}{k}\frac{2}{n} = \frac{1}{2^l}\cdot\frac{2^{l-1}\cdot 2}{n} = \frac{1}{n}.$$

Therefore, if it is equally likely that the divisor starts at any $(n-k, k)$, then the likeliness of attaining weight $n+l$ will be $\frac{1}{n}$.

This means our threshold density will be, as $\deg(D) \to -\infty$,

$$\zeta_{D_0} = \sum_{l=0}^{n-1} P(l)\frac{n+l}{2} \to \sum_{l=0}^{n-1}\frac{n+l}{2n} = \frac{3n-1}{4} = \zeta_{st}.$$

Conclusion

In this paper we have given an exposition of the Threshold Density theorem. The equality between the stationary density and the threshold density in the limit as the degree of the starting divisor goes to $-\infty$ came as a result of the uniform stationary distribution of threshold divisors of a certain degree. We gave an explicit calculation of the threshold density for the banana graph.

Attempts were also independently made to explicitly find the threshold density, starting from zero, for the complete graph, K_n, but the complex nature of finding the probability of reaching a threshold divisor made finding a closed-form solution difficult. That work does not appear here. For most general graphs, other than trees, the banana graphs, and cycle graphs, the threshold density starting at zero is still unknown.

References

Deepak Dhar. Self-organized critical state of sandpile automaton models. *Phys. Rev. Lett.*, 64(14):1613–1616, 1990. ISSN 0031-9007. doi: 10.1103/PhysRevLett.64.1613. URL `http://dx.doi.org/10.1103/PhysRevLett.64.1613`.

Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete mathematics*. Addison-Wesley Publishing Company, Reading, MA, second edition, 1994. ISBN 0-201-55802-5. A foundation for computer science.

Harry Kesten. Renewal theory for functionals of a Markov chain with general state space. *Ann. Probability*, 2:355–386, 1974.

David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009. ISBN 978-0-8218-4739-8. With a chapter by James G. Propp and David B. Wilson.

Lionel Levine. Threshold state and a conjecture of Poghosyan, Poghosyan, Priezzhev and Ruelle. *Comm. Math. Phys.*, 335(2):1003–1017, 2015. ISSN 0010-3616. doi: 10.1007/s00220-014-2216-5. URL `http://dx.doi.org/10.1007/s00220-014-2216-5`.

Criel Merino López. Chip firing and the Tutte polynomial. *Ann. Comb.*, 1(3):253–259, 1997. ISSN 0218-0006. doi: 10.1007/BF02558479. URL `http://dx.doi.org/10.1007/BF02558479`.

David Perkinson and Scott Corry. Divisors and sandpiles. 2017. URL `http://people.reed.edu/~davidp/divisors_and_sandpiles/`.

John Riordan. *Combinatorial identities*. John Wiley & Sons, Inc., New York-London-Sydney, 1968.