The Sandpile Group on a Hexagonal Grid

Eli Johnson

A thesis
submitted in partial fulfillment of the
requirements for the degree of

Master of Science

University of Washington

2019

Reading Committee:

Sara Billey, Chair

Ken Bube

Program Authorized to Offer Degree:
Mathematics

University of Washington

Abstract

The Sandpile Group on a Hexagonal Grid

Eli Johnson

Chair of the Supervisory Committee:
Professor Sara Billey
Mathematics

We introduce the notion of an abelian sandpile, including the basic definitions, its use as a model, and some of the foundational theorems and observations. We describe a constructive method for reducing redundant computations when the underlying graph has one or more symmetries. We motivate the analysis of a particular two-dimensional model: the hexagonal tiling of $\mathbb{R}^2$, and extend a convergence result for the integer lattice $\mathbb{Z}^d$ to this hexagonal grid, using discrete approximation of partial differential equation methods. We end with some asymptotic conjectures on this grid.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

The author wishes to thank Sara Billey for her invaluable guidance and encouragement at all steps of the process. Her leadership on this project began at its inception and continued thanks to her infectious enthusiasm and wellspring of expertise. Thanks to Ken Bube for his helpful comments and edits. The author thanks Timea Tihanyi for providing inspiration and an artist's perspective, as well as many interesting questions and avenues of investigation. Much is owed to the collaboration and brainstorming sessions done with Catherine Babecki and James Pedersen. Both helped in programming, as well as in reading and responding to early drafts. Thank you Kelsey Yale for help with editing, lending a critical eye and a supportive ear.

Chapter 1

INTRODUCTION

Imagine a tall pile of sand balanced precariously in a narrow column over a flat surface. The pile will be unstable, sending a wave of grains of sand propagating outward. The base of the pile will enlarge, reducing the slope of its upper surface until the pile stabilizes and sand stops moving. This is not to say the sand is far from instability. A single extra grain placed in an arbitrary location might be enough to cause another chain reaction through part or all of the pile.

Although physical grains of sand move continuously into an infinity of possible states, in certain circumstances it may be warranted to consider a discrete model with a finite number of positions. A 1987 paper by Bak, Tang and Weisenfeld [1] proposed such a model, which has since developed into the *abelian sandpile*. Despite the name, the model captures characteristics present in many systems beyond simply sand. Indeed, the authors motivated their construction with the example of a two-dimensional square array of pendulums, connected to neighbors by torsion springs. A single rotation of any pendulum would affect the resting state of each of its neighbors, possibly causing a chain reaction of rotations backward and forward in the surrounding pendulums. The end behavior was qualitatively distinct from that of a similar one-dimensional array, since the greater connectivity of the two dimensional array requires a greater degree of stability to avoid effects that propagate infinitely. The authors noted that the array will globally stabilize when such "minimally stable" states have been broken into local structures. That is, there is a self-organized process by which certain areas may be on the cusp of instability, yet there is enough stability in the system as a whole so that noise can't travel infinitely far.

One characteristic we might expect from such a barely-stable state is scale-invariance.

If the "slack" in the system, the leeway to allow a certain degree of noise without setting off unending cascades, has occurred through self-organized processes, then it may be unpredictable how large the effect will be from a small perturbation at a single location. Indeed, the authors in [1] experimentally measured the size of such cascades in their model, and found them to follow a power law distribution. If $D(s)$ is the number of cascades eventually effecting s sites, then $D(s) \sim s^{-\tau}$, with $\tau \approx 0.98$. The authors found a similar power law with an exponent of $\tau \approx 1.35$ for the three-dimensional model.

In both of these cases, the measured exponent of the power law falls within the range of a phenomenon known as $1/f$ noise, or pink noise. Pink noise describes noise with an observed power spectrum intensity that is roughly inversely proportional to frequency f. In the example above, this translates to the property that $D(s) \sim s^{-\tau}$, where $0.5 \leq \tau \leq 1.5$. In broad strokes, low frequency noise is much more common than high frequency noise. In the sandpile model which we will describe, we are considering the area of effect of the cascading waves of sand to be the noise that accompanies the "signal" of the addition of chips at a single location. Small cascades are much more common than large cascades.

Pink noise has been observed in the fields of astrophysics, technology, biology, geophysics, economics, psychology, language and even music [15]. The purpose of the original model was to put this widespread phenomenon into an explanatory framework. The model does not require fine-tuning to produce configurations that show the power law relationship in the distribution of areas affected by small disturbances. Bak in particular went on to propose the underlying concept of self-organized criticality as an explanatory framework for understanding phenomena like forest fires, stock markets, and neural activity in the brain [9]. Later contributions by Dhar [4] expanded and generalized the theory, and many others have since further illuminated various aspects.

Beyond the physical applications, sandpiles have a natural aesthetic appeal. The scale-invariance that characterizes the power law dynamics also seems to produce beautiful and intricate fractal-like images associated to many configurations. Remarkably, after scaling to account for number of chips or the size of the grid, many images appear to converge to

a limiting object. Certain properties of these images also persist across different choices for the underlying grid and consequential dependencies between positions. It might seem paradoxical to have a system that is sensitive to variation of initial data, yet produces images that are robust with respect to variation of parameters.

Figure 1.1: A sandpile. Height of the pile is represented by color, with lighter blue indicating a higher chip count. Here the allowable positions of sand are at the centers of octagons and squares that tessellate a large square region. This particular sandpile is the identity element on the region, a concept we will make precise in Chapter 3.

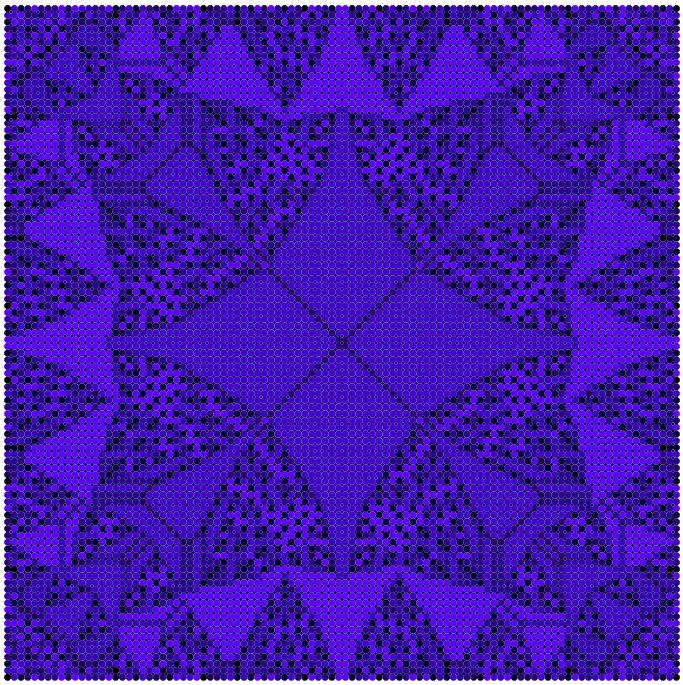

In this paper we will describe the foundations of the model and some fundamental results about it. In the chapter on symmetry we will show how the computations involved in stabilizing a pile can be significantly streamlined by paying attention to the symmetries involved. We will then specialize to a particular grid based on hexagons which has received substantially less attention, motivating this choice and extending a convergence result to it.

Chapter 2

NOTATION

We will typically use G to refer to a graph or directed graph (digraph for short). We will refer to V and E (rather than $V(G)$ and $E(G)$) as the vertex set and edge set of G respectively, even if the correspondence $G = (V, E)$ hasn't been made explicit. If G is a digraph then the ordered pair $(x, y) \in E$ refers to the directed edge going from x to y. We may implicitly turn a graph into a digraph by considering each edge in E as the pair of directed edges between its incident vertices. For this reason any result here that applies to digraphs applies to graphs as long as the other conditions are satisfied.

When G is a graph we will use $N(x)$ to mean the (immediate) neighbors of x.

$$N(x) := \{y \in V \,:\, xy \in E\}.$$

We will define vertex degree to be $d(x) := |N(x)|$. The length of the shortest path in G from x to y will be denoted by $d(x, y)$, which is a metric on V.

If G is a digraph we let

$$N^+(x) := \{y \in V \,:\, (x, y) \in E\},$$

and we define $\deg^+(x) := |N^+(x)|$. The operator $d(x, y)$ in a digraph is the length of the shortest path from x to y. This is not necessarily a metric since $d(x, y)$ and $d(y, x)$ may be different.

If $W \subseteq V$ is a subset of the vertices of G, then the *induced subdigraph* $G[W]$ is the digraph with vertex set W, and edge set $\{(x, y) \in E : x, y \in W\}$.

A convenient invariant we will use when given a finite digraph or multi-digraph G (without loops) is the *Laplacian matrix of* G, which is the $|V| \times |V|$ matrix $\Delta(G)$, with $\Delta(G)_{ii} = -\deg^+(v_i)$ and $\Delta(G)_{ij} = \#\{$ edges from v_i to $v_j\}$ when $i \neq j$.

A vertex x of a digraph is a *sink* if $\deg^+(x) = 0$. In our case we will only look at digraphs with at most one sink, in which case we will understand the *reduced Laplacian matrix* $\Delta'(G)$ to be the matrix obtained by removing the row and column of the sink vertex from $\Delta(G)$. We will define other Laplacian operators later, so to avoid confusion we will never omit the parentheses and argument G when referring to either matrix.

Finally, $\mathbb{N}$ will always mean the nonnegative integers.

Chapter 3

THE BASICS OF SANDPILES

3.1 Definitions, structure and properties

An abelian sandpile has as its starting point a directed graph $G = (V, E)$. Along with the graph structure, we think of each vertex x of the graph as holding a nonnegative number of chips $\sigma(x)$, and we call this $\sigma : V \to \mathbb{N}$ a *chip configuration*. We will sometimes abbreviate this as simply a configuration. A vertex x is *stable* if $\sigma(x) < \deg^+(x)$, and *unstable* otherwise. If a vertex is unstable it can *fire* or *topple*, sending exactly one chip along every edge originating at x. The new chip configuration σ' resulting from firing x will be identical to σ except that $\sigma'(x) = \sigma(x) - \deg^+(x)$, and $\sigma'(y) = \sigma(y) + 1$ for every $y \in N^+(x)$. Such a chip configuration is called *successive* to σ. A chip configuration is considered *stable* if every vertex is stable. To stabilize a chip configuration σ is to list a sequence of successive chip configurations beginning with σ and ending with a stable configuration. We denote such a configuration σ° if it exists, and call it the *stabilization* of σ. As the notation suggests, when it exists the stabilization of a chip configuration is always unique. This is perhaps surprising, since there are possibly many distinct sequences of successive chip configurations beginning with a given one. In fact, the following stronger claim can be made, due to Dhar from 1990 [4]

Lemma 3.1.1. *Let G be any digraph, and let $\sigma_0, \ldots, \sigma_n$ be a sequence of successive chip configurations. Let $\sigma'_0, \sigma'_1, \ldots, \sigma'_m$ be another sequence of successive chip configurations, with $\sigma_0 = \sigma'_0$. Then*

1. *If σ_n is stable, then $m \leq n$, and no vertex fires more times in $\sigma'_0, \ldots, \sigma'_m$ than in $\sigma_0, \ldots, \sigma_m$.*

2. *If σ_n and σ'_m are both stable, then $m = n$, $\sigma'_m = \sigma_n$, and every vertex fires the same number of times in both histories.*

A consequence of Lemma 3.1.1 is that if σ eventually stabilizes, then the number of firings in going from σ to σ° at a vertex x is a well-defined function on V. We call this function $v : V \to \mathbb{N}$ the *odometer function* associated with σ. The reason we write v instead of v_σ is that for our purposes, context or indexing is sufficient to make the initial configuration understood. This function will become important in the final chapter on convergence.

An interesting structure of sandpiles arises by singling out a certain vertex $s \in V$ with $\deg^+(s) = 0$ which we call the *sink*. We say a sink s is *global* if for any $x \in V$ there is a directed path of finite length from x to s. It is easy to see that any global sink must be unique. The stability and number of chips on s are ignored. For this reason we take the domain of any chip configuration to be $V \setminus s$. It may be illustrative to think of any chips on s as having "fallen off the graph". The following definitions will assist in proving a useful result about digraphs with a global sink.

Definition 3.1.2. If G is a finite or infinite digraph with a global sink s, and if σ is any chip configuration on G, we define the *sink distance vector* of σ to be the vector $S_\sigma = (n_1(\sigma), n_2(\sigma), \ldots)$, where

$$n_i(\sigma) = \sum_{x : d(x,s) = i} \sigma(x).$$

The first entry of the sink distance vector is the number of chips at distance 1 from the sink. The second entry is the number of chips at distance 2 from the sink, and so on. We will write $\mathbb{N}^\mathbb{N}$ to mean the set of all infinite vectors of nonnegative integers, which we consider endowed with the *lexicographical order*. That is, $(n_1(\sigma), n_2(\sigma), \ldots) > (n_1(\tau), n_2(\tau), \ldots)$ if and only if $n_i(\sigma) > n_i(\tau)$ at the first index i where terms are not equal. Let

$$N_\sigma := \sum_{x \in V \setminus s} \sigma(x) = \sum_{i=1}^{\infty} n_i(\sigma)$$

be the total number of chips in the configuration σ.

Lemma 3.1.3. *Let $G = (V, E)$ be a digraph with a global sink s. Suppose σ, σ' are chip configurations on G, such that σ' is successive to σ by firing x, and $(x, s) \notin E$. Then*

P1. $N_{\sigma'} = N_\sigma$.

P2. $(n_1(\sigma'), n_2(\sigma'), \ldots) > (n_1(\sigma), n_2(\sigma), \ldots)$.

Proof. Statement **P1** is immediate from the fact that $(x, s) \notin E$, so firing x does not change the number of chips on $V \setminus s$. Statement **P2** follows from the fact that s is a global sink, so there is a finite path in G from x to s. This implies there is some $y \in N^+(x)$ with $1 \leq d(y, s) = d(x, s) - 1$. Firing x will increase $n_{d(x,s)-1}(\sigma)$ by at least 1, while only the $d(x, s)$-th entry will decrease. □

Lemma 3.1.4. *If G is a finite or infinite digraph with a global sink, then any chip configuration on G with a finite number of chips will stabilize in a finite number of steps.*

Proof. Fix $G = (V, E)$ and a chip configuration $\sigma : V \to \mathbb{N}$. When $N_\sigma = 0$ the statement is trivial. Suppose by induction the statement is proved for all chip configurations with fewer than N_σ chips. If σ stabilizes or fires into the sink in a finite number of steps, we are done by induction. Otherwise assume for a contradiction there exists an infinite sequence of successive configurations $\sigma = \sigma_0, \sigma_1, \sigma_2, \ldots$ obtained by firing vertices $x_1, x_2, \ldots$ such that for all i, $(x_i, s) \notin E$. By **P1**, we know $N_{\sigma_i} = N_\sigma$ for all i. By **P2**, the sequence $(S_{\sigma_i})_{i \geq 1} := S_{\sigma_1}, S_{\sigma_2}, \ldots$ is strictly increasing lexicographically.

Note $(S_{\sigma_1})_{i \geq 1}$ can never increase beyond $(N_\sigma, 0, \ldots)$, since this is the lexicographic maximum chip vector whose coordinates sum to N_σ. Any bounded subset of $\mathbb{N}^\mathbb{N}$ has a lexicographic supremum [8], so let $\overline{S} = (s_1, s_2, \ldots)$ be the supremum of $(S_{\sigma_1})_{i \geq 1}$.

Suppose for a contradiction $\overline{S}$ is such a supremum. If $(S_{\sigma_i})_{i \geq 1}$ never reaches s_1 in its first coordinate, then it must reach $s_1 - 1$ in the first coordinate and increase without bound in the second coordinate, or there is a lower lexicographic supremum. This is impossible since each coordinate is bounded above by N_σ, so eventually $n_1(\sigma_i) = s_1$. By the same reasoning

$(S_{\sigma_1})_{i\geq 1}$ must increase at least to $(s_1, s_2, 0, \ldots)$. Proceeding in this way, since N_σ is finite we reach $\overline{S}$ in a finite number of steps. At this point the chip configuration must increase lexicographically beyond $\overline{S}$, contradicting the fact that $\overline{S}$ is the supremum of $(S_{\sigma_1})_{i\geq 1}$.

□

Remark 3.1.5. If the global sink condition is removed from Lemma 3.1.4, then it becomes false. For example, no nonzero chip configuration will stabilize on K_2, the complete graph on 2 vertices. For an infinite digraph counterexample, let $G_2 = (V, E)$ with $V = \mathbb{N}$ and $E = \{(i, i+1) : i \in \mathbb{N}\}$. Again, no nonzero chip configurations will stabilize.

3.2 Monoids and Ideals

The set of all finite chip configurations on a particular digraph forms an abelian monoid according to vertex-wise addition of chips. If we write $\sigma + \tau$, we refer to this vertex-wise addition, which so far is no different from addition of $\mathbb{N}$-valued functions on a graph. If G has a sink then by Lemma 3.1.1 and Lemma 3.1.4, the operator $(\,\cdot\,)^\circ$ is always defined. We may define a new monoid addition on the set of all *stable* chip configurations, given by $(\sigma, \tau) \mapsto (\sigma + \tau)^\circ$. The operation remains commutative since the underlying chipwise addition commutes. If G is finite, this new operation on stable configurations has the added benefit of being a finite monoid since the number of stable chip configurations is exactly

$$\prod_{x \in V} \deg^+(x).$$

Definition 3.2.1. Let M be an abelian monoid, with operation $(a, b) \mapsto a \cdot b$. An *ideal* of M is a subset I such that $x \cdot y \in I$ whenever $x \in I$ and $y \in M$.

We will use the following algebraic fact to endow an important subset of the monoid of stable chip configurations with a group structure.

Lemma 3.2.2. *The unique minimal ideal of a finite abelian monoid is a group.*

Proof. Let $(M, \cdot)$ be a finite abelian monoid. Let

$$J := \bigcap_{\substack{I \subseteq M \text{ is an ideal} \\ I \neq \varnothing}} I$$

The index of the intersection defining J is nonempty, since M is itself a nonempty ideal of M. We claim J is the unique nonempty minimal ideal of M. To see $J \neq \varnothing$, label the elements of M as $m_1, \ldots, m_n$. Let

$$m^* := \prod_{i=1}^{n} m_i.$$

If $I \subseteq M$ is a nonempty ideal with $m_k \in I$, then $m_k \cdot \prod_{i \neq k} m_i = m^*$. By the fact that I is an ideal, $m^* \in I$. As I was arbitrary, $m \in J$. In particular J is nonempty. Clearly J is an ideal since if $j \in J$ and $m \in M$, and if $I \subseteq M$ is any other ideal, then $j \in I$ and so $j \cdot m \in I$, which since I was arbitrary implies $j \cdot m \in J$. As J is the intersection of all ideals, it must be the smallest.

For $a \in J$ arbitrary the sequence $a, a^2, a^3, \ldots$ must eventually become periodic since J is finite. This means there are $j, k \in \mathbb{N}$ such that for all $n \in \mathbb{N}$, $a^j = a^{j+nk}$. Set $e = a^{jk}$. Then

$$e^2 = a^{2jk} = a^{j+jk} a^{j(k-1)} = a^j a^{j(k-1)} = a^{jk} = e$$

so e is idempotent.

If we take another $h \in J$ then the ideal hM is an ideal contained in J, so must be J itself. This is to say that any element of J is expressible as hy for some $y \in M$. Therefore $hy = e$ for some $y \in M$. By the same reasoning, $eM = J$, so $ez = h$ for some $z \in M$. Putting these together,

$$eh = e(ez) = e^2 z = ez = h$$

By the fact that M is abelian, e also acts as a right inverse for h. Since $h \in J$ was arbitrary, e is the identity on J. Further,

$$h(ye) = (hy)e = e^2 = e$$

and as $ye \in J$ by the property of an ideal, therefore $ye = h^{-1}$ and J is a group. $\square$

3.3 The sandpile group

In what follows we assume G is a finite digraph with a global sink. As we saw earlier, the stable chip configurations on G form a finite commutative monoid under the operation $(\sigma, \tau) \mapsto (\sigma + \tau)^\circ$. If we can characterize this monoid's unique minimal ideal, then by Lemma 3.2.2, we will have described a group with the same operation on some subset of stable configurations.

Definition 3.3.1. A chip configuration is *accessible* if, given any other chip configuration, it can be reached by a combination of adding chips and firing vertices. A chip configuration is *recurrent* if it is both stable and accessible. We let $\mathcal{S}(G)$ denote the set of all recurrent chip configurations on G.

A bit of reflection will show that the recurrent property exactly characterizes membership in the minimal ideal. To see this, for the moment let J be the minimal ideal on stable configurations of G (we will see shortly that $J = \mathcal{S}(G)$). Let $\sigma \in J$ and let τ be any stable chip configuration. Then $(\tau + \sigma)^\circ \in J$, so $((\tau + \sigma)^\circ + \sigma)^\circ - (\tau + \sigma)^\circ = \sigma$. This shows σ to be recurrent and that $\mathcal{S}(G) \supseteq J$. On the other hand if $\sigma \notin J$ then it can not be reached from anything inside J or it would contradict the fact that J is an ideal, so $\mathcal{S}(G) \subseteq J$. Combining these facts with Lemma 3.2.2, $\mathcal{S}(G)$ is the *group* of recurrent chip configurations.

In terms of an explicit algebraic construction, another equivalent method for arriving at the same group structure of $\mathcal{S}(G)$ is as a quotient from a larger group. Whereas chip addition defines a monoid on chip configurations, we can start with the *group* that is integer-valued functions on non-sink vertices under vertex-wise addition. The group is isomorphic to $\mathbb{Z}^{V \setminus s}$. Starting with a given chip configuration—thought of as an integer-valued $|V| - 1$-tuple—firing a particular vertex x then becomes addition of a vector which is $-\deg^+(x)$ at the position corresponding to x, while 1 at all of the outgoing neighbors of x, and 0 everywhere else. This vector will be a row vector of $\Delta'(G)$. Since stabilization is uniquely defined and our attention is on a subset of stable configurations, it may be natural to take the quotient

of $\mathbb{Z}^{n-1}$ with the integer row span of $\Delta'(G)$

$$\mathcal{S}(G) := \mathbb{Z}^{n-1}/\mathbb{Z}^{n-1}\Delta'(G),$$

which will define an equivalence relation between successive configurations (and others).

Lemma 3.3.2. *Every equivalence class of $\mathbb{Z}^{n-1}$ modulo $\Delta'(G)$ contains exactly one recurrent chip configuration. Therefore*

$$\mathcal{S}(G) \cong \mathbb{Z}^{n-1}/\mathbb{Z}^{n-1}\Delta'(G).$$

Proof. See [7] (Lemmas 2.13 and 2.15) for a proof of the first statement. The second statement an easy consequence of the first. □

The isomorphism in Lemma 3.3.2 justifies alternating between either characterization of $\mathcal{S}(G)$ as it suits our immediate needs, which we will do by abuse of notation. It is natural to choose recurrent chip configurations as equivalence class representatives of $\mathcal{S}(G)$ since these are the configurations common to both characterizations. The quotient construction has the benefit of generalizing to quotients of $\mathbb{Z}^n$ with the row span of other matrices (see [2]).

3.4 Finding recurrent configurations

Up until now we have not said anything concrete about what these recurrent configurations actually look like. The following lemma gives a process for computing the identity element of $\mathcal{S}(G)$.

Definition 3.4.1. Let δ be the chip configuration on G given by $\delta(x) = \deg^+(x)$, and let

$$\epsilon := ((2\delta - 2) - (2\delta - 2)^\circ)^\circ.$$

Lemma 3.4.2. *The configuration ϵ is the unique recurrent configuration in the equivalence class of the identity on $\mathcal{S}(G)$. Furthermore, for any chip configuration σ, $(\sigma + \epsilon)^\circ = \sigma$ if and only if σ is recurrent.*

Figure 3.1: The recurrent identity configuration of the sandpile group on the Petersen graph. An arbitrary vertex has been made the sink. This is computable by hand using Lemma 3.4.2. Any constant configuration $\sigma(x) \equiv c$ with $c \geq 2$ will stabilize to the identity.

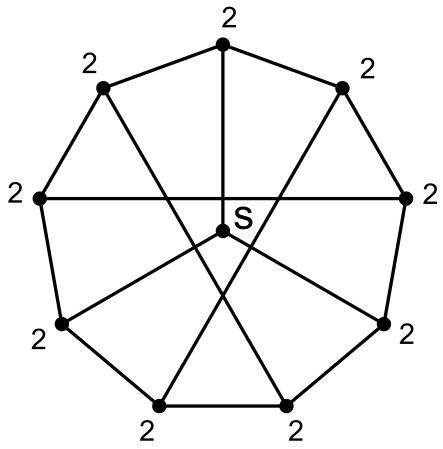

Proof. As $(2\delta - 2)$ and $(2\delta - 2)^\circ$ are in the same equivalence class modulo $\Delta'(G)$, their difference is in the linear span of $\Delta'(G)$, and stabilizing does not change that, so ϵ is in the equivalence class of the zero configuration, with no chips on any vertex. The zero configuration is itself the identity on $\mathbb{Z}^{V \setminus s}$, which means ϵ is in the equivalence class of the identity. Clearly ϵ is stable, so we need only show ϵ is accessible. Since $(2\delta - 2)^\circ$ is stable, we know $0 \leq (2\delta - 2)^\circ(x) < \deg^+(x)$ for every $x \in V \setminus s$. Thus, $((2\delta - 2) - (2\delta - 2)^\circ)(x) \geq \deg^+(x) - 1$ for all $x \in V \setminus s$. If σ is any other chip configuration, then $\sigma^\circ(x) \leq \deg^+(x) - 1$, so $(2\delta - 2) - (2\delta - 2)^\circ$ can be reached from σ by adding chips. Thus its stabilization can be reached from σ as well, so ϵ is accessible.

The "if" part of the second statement is a consequence of the first, using the characterization of $\mathcal{S}(G)$ and its group operation as addition and stabilization of recurrent configurations. Since ϵ is accessible, $(\sigma + \epsilon)^\circ$ is recurrent, and so $(\sigma + \epsilon)^\circ = \sigma$ only if σ is also recurrent. $\square$

3.5 Examples of sandpile groups

The sandpile groups of various graphs and digraphs have been made the objects of study. One of the most investigated examples is the d-dimensional integer lattice, with vertex set $\mathbb{Z}^d = \{(x_1, \ldots, x_d) : x_i \in \mathbb{Z}\}$, with edges between vertices that differ in one coordinate by exactly one, and are equal in every other coordinate.

The graph may be made finite by requiring coordinates be bounded by a uniform bound. Take $V = \{(x_1, \ldots, x_d) \in \mathbb{Z}^d : |x_i| \leq M \text{ for all } i \leq d\}$. Typically the parameter M is taken to be large. In this case we may connect all boundary vertices to the sink. Corner vertices, edge vertices, etc., are multiply connected to the sink so that each vertex has degree $2d$. Alternatively, every boundary vertex might be connected once to the sink so that the graph has no multiple edges, producing a different sandpile group structure.

We have found that many of the same properties of the sandpile group on this square grid carry over to other tessellations of the plane, notably the equilateral, hexagonal and octagon/square tessellations.

Figure 3.2: A portion of the hexagonal grid, with colors representing various chip sizes.

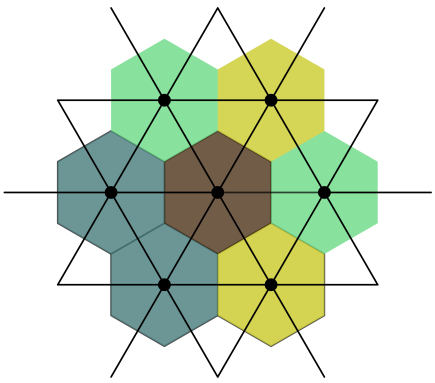

We have paid particular attention to the 6-regular grid whose vertices are the centers of tessellating hexagons (Fig. 3.2). We will call this the *hexagonal grid*, though some readers might insist that this is actually a triangular grid since every face of the graph is a triangle.

Although triangular would be perhaps more accurate from a strictly graph-theoretic perspective, since we will use colored regions to represent chip counts on vertices, the edges are essentially invisible (we are instead looking at the dual graph), and the associated symmetry group (more later) is that of a hexagon, so this sandpile group tends to have a hexagonal feel. The reasons for investigating this particular grid are

1. Sandpiles can be thought of as a discrete model of continuous physical systems exhibiting criticality. This grid is a tessellation whose graph distance better approximates Euclidean distance than the 2-dimensional integer lattice, and which has a larger degree of freedom of motion in moving to an adjacent vertex.

2. Higher vertex degree means stable configurations have more possible states at each vertex, so better resolution of large-scale structures is possible.

3. The hexagonal grid admits interesting generalizations to higher dimension, in the form of sphere packing lattices.

Figure 3.3: The recurrent identity element on the hexagonal grid with side length 400.

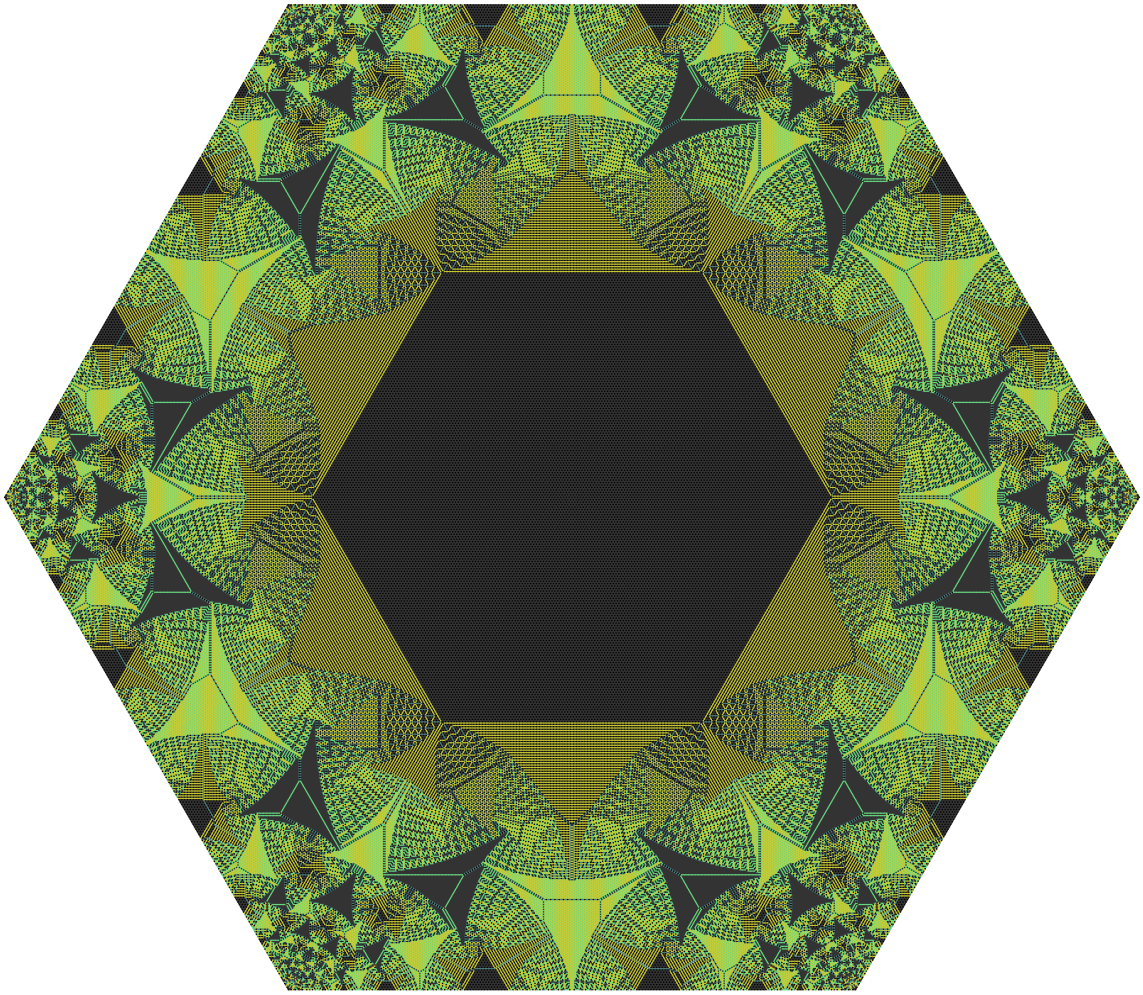

Chapter 4

EXPLOITING SYMMETRIES

4.1 Symmetries as automorphisms

It can be computationally intensive to stabilize a chip configuration on a large graph. In a square $N \times N$ subset of the grid $\mathbb{Z}^2$, for example, when a single chip is added to a random stable configuration, the expected number of firings before stabilization occurs goes to infinity as $N \to \infty$ [13, p. 1]. The purpose of this section is to clarify how the underlying symmetries in the graph may be used to speed up computation.

Lemma 4.1.1. *Let $G = (V, E)$ be a digraph, and assume $\varphi : V \to V$ induces a graph automorphism. Let σ be a chip configuration on G that is invariant with respect to φ, i.e., $\sigma = \sigma \circ \varphi$. Then σ° is invariant with respect to φ.*

Proof. Let $\sigma = \sigma_0, \sigma_1, \ldots, \sigma_n = \sigma^\circ$ be a sequence of successive chip configurations on G. Let $\tau_i := \sigma_i \circ \varphi$, for every $0 \leq i \leq n$. Then $\tau_0 = \sigma_0 \circ \varphi = \sigma \circ \varphi = \sigma$ by assumption. If σ_{i+1} results from starting with σ_i and firing $x_i \in V$, then τ_{i+1} results from starting with τ_i and firing $\varphi(x_i) \in V$, so τ_{i+1} is successive to τ_i. Stability of a chip configuration is unchanged by graph automorphism, since for any $x \in V$, $d(x) = d(\varphi(x))$, so $\sigma(x) \geq d(x)$ if and only if $(\sigma \circ \varphi)(x) \geq d(\varphi(x))$. Thus $\tau_n := \sigma^\circ \circ \varphi$ is stable. Therefore by Lemma 3.1.1 we have $\tau_n = (\sigma \circ \varphi)^\circ = \sigma^\circ$. We have shown $\sigma^\circ = \sigma^\circ \circ \varphi$, so by definition σ° is invariant with respect to φ.

□

4.2 The reflection group of a geometric graph

For the remainder we let $G = (V, E)$ be a digraph together with an embedding $|G| \to \mathbb{R}^n$. By abuse of notation we associate G with the image of this embedding, to make G a geometric

object. If $\alpha : \mathbb{R}^n \to \mathbb{R}^n$ is any reflection or rotation that sends vertices to vertices and edges to edges (preserving the direction on each edge), then we consider G symmetric with respect to α. If in addition, σ is a chip configuration on G and $\sigma(\alpha(x)) = \sigma(x)$ for all $x \in V$ (here we are also associating a vertex with its image in $\mathbb{R}^n$), then we may also consider σ symmetric with respect to α. If G has a global sink s, then we will want to consider symmetries without consideration given to s, so we may as well consider s as living outside of $\mathbb{R}^n$ and fixed by every symmetry. For example, s could be along the x_{n+1} axis in $\mathbb{R}^{n+1}$.

Corollary 4.2.1. *Let G be a geometric digraph embedded in $\mathbb{R}^n$, and let σ be a chip configuration on G. Let $\alpha : \mathbb{R}^n \to \mathbb{R}^n$ be a reflection or rotation such that both G and σ are symmetric with respect to α. Then σ° is symmetric with respect to α.*

Proof. This follows immediately from Lemma 4.1.1 and the fact that any reflection or rotation induces an automorphism on G. In this context invariance with respect to the automorphism is equivalent to symmetry with respect to the rotation or reflection. □

In fact the set of all reflections and compositions of reflections $\mathbb{R}^n \to \mathbb{R}^n$ with respect to which G is symmetric form a group under composition, called a *reflection group*, which we denote R. The group R acts on the vertices of G by the rule $\alpha \cdot x = \alpha(x)$, that is, the group action takes each vertex to its image under the given composition of reflections. Now define an equivalence relation $\sim$ on V, where $x \sim w$ if and only if there exists some $\alpha \in R$ such that $x = \alpha \cdot w$. Each equivalence class is exactly one orbit under the group action. This clearly satisfies reflexivity, symmetry and transitivity. Let $R \cdot x$ denote the equivalence class of x, i.e. the orbit under the group action, and let $|R \cdot x|$ be the size of this equivalence class. Up to translation, the group R is determined by a central arrangement of hyperplanes of reflection H_i. We may select one connected component of $\mathbb{R}^n \setminus \cup H_i$, and call the closure of this component the *fundamental chamber*, which is a closed cone with (cone) vertex at the origin containing exactly one (graph) vertex in every orbit.

We would like to exploit the symmetries of a chip configuration in order to reduce the calculations needed to arrive at its stabilization. If $W \subset V$ are the vertices in the fundamental

chamber, we might hope that a graph H could be derived from the induced graph $G[W]$ in some way so that $\sigma^\circ|_H = (\sigma|_H)^\circ$, since the right hand side will be considerably easier to calculate. If some version of this holds, then by the above corollary we can recover all of σ° from $(\sigma|_H)^\circ$.

Definition 4.2.2. The stabilizer R_x of x is the set of all group elements in R that fix x:

$$R_x := \{\alpha \in R : \alpha \cdot x = x\}.$$

If $\alpha \in R$ is a reflection, the *hyperplane of symmetry* H_α associated with α is the set

$$H_\alpha := \{x \in \mathbb{R}^n : \alpha(x) = x\}.$$

If $n = 2$ we will say *line of symmetry.*

An alternate characterization of R_x is the number of chambers whose closure contains x. This is because the only elements of a reflection group that can fix a vertex x are either the identity, or a composition of reflections, each of which fix x. If a reflection α fixes x, then $x \in H_\alpha$.

4.3 Restricting to the fundamental chamber

The orbit-stabilizer theorem (see e.g. [11, p. 28]) says that for any $x \in V$ we have

$$|R \cdot x| \cdot |R_x| = |R|. \tag{4.1}$$

The following notation will simplify the constructions and statements to follow. For any $(x, \overline{y}) \in V \times V/R$, we let $d_{x,\overline{y}}$ be the number of edges in G from x to a vertex in $R \cdot y$. In addition, let

$$m_{x,\overline{y}} := d_{x,\overline{y}} \cdot \min\left\{1, \frac{|R_y|}{|R_x|}\right\}.$$

Finally, let

$$k_{x,\overline{y}} := \max\left\{1, \frac{|R_y|}{|R_x|}\right\}.$$

Lemma 4.3.1. *For any $x, y \in V$, the quantities $m_{x,\overline{y}}$ and $k_{x,\overline{y}}$ are nonnegative integers.*

Proof. When $\frac{|R_y|}{|R_x|} \leq 1$, then $m_{x,\overline{y}} = d_{x,\overline{y}}$ which is clearly a nonnegative integer. Otherwise $m_{x,\overline{y}} = d_{x,\overline{y}} \cdot \frac{|R_y|}{|R_x|}$. Note that $\frac{|R_y|}{|R_x|} = \frac{|R \cdot x|}{|R \cdot y|}$ by Eq. (4.1). The total outdegree from all of $R \cdot x$ into $R \cdot y$ is $|R \cdot x| \cdot d_{x,\overline{y}}$, which must equal the total indegree into $R \cdot y$ from $R \cdot x$, which is an nonnegative integer multiple of $|R \cdot y|$. Thus $m_{x,\overline{y}}$ is a nonnegative integer.

For clarity we will show $k_{x,\overline{y}}$ is a nonnegative integer when G is a planar geometric digraph (i.e. G or $G \setminus s$ is a subset of $\mathbb{R}^2$), though the result holds more generally in $\mathbb{R}^n$. Assume $|R_y| \geq |R_x|$, since the claim is trivially true otherwise. If $x \in V$ there are then three possible values for $|R_x|$. The first possibility is that x is in the interior of some chamber, in which case $|R_x| = 1$. Secondly, we could have x on the boundary of exactly two chambers, in which case $|R_x| = 2$. Finally, x could be at the origin, in which case $|R_x| = 2\#\{H_i\}$ (where the H_i are lines of symmetry through the origin). The same possibilities hold for $|R_y|$. Given any two of these choices, the smaller always divides the larger, so $k_{x,\overline{y}}$ is a nonnegative integer. □

We are now ready to construct an appropriate graph on the vertices of the fundamental chamber. Recall we are assuming $G = (V, E)$ is a geometric digraph (with or without a global sink s). Let $W \subseteq V$ be the vertices in the fundamental chamber.

Let $H = (W, F)$, where $(x, y) \in F$ has multiplicity $m_{x,\overline{y}}$. Add s to W regardless of whether G has a global sink s, and give $(x, s) \in F$ multiplicity $\deg^+_G(x) - m_{x,\overline{y}}$.

Give H the following amended sandpile rule: whenever a chip travels along the edge $(x, y) \in F$ to y, we add not one chip to y but $k_{x,\overline{y}}$ chips.

Lemma 4.3.2. *Let G be any digraph, and construct H on the fundamental chamber of G as above. If σ is a chip configuration symmetric with respect to the reflection group R for G, then $\sigma^\circ|_H = (\sigma|_H)^\circ$.*

Proof. Note that the added sandpile rule defaults to our usual notion of $k_{x,\overline{y}} = 1$ except when y is on the boundary of the fundamental chamber. By Lemma 4.3.1, H is well-defined.

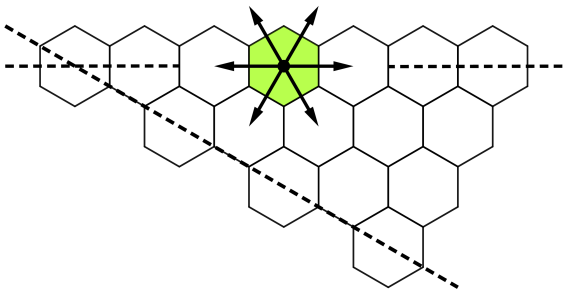

Figure 4.1: The fundamental chamber of a large hexagonal grid. The stabilizer of the firing vertex is larger than or equal to the stabilizers of its neighbors, so only one edge to each vertex equivalence class will be used in H. The remaining two chips will go to the sink.

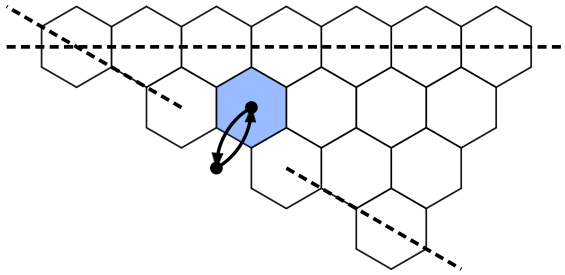

Figure 4.2: Here the stabilizers are the same size, so the edge is used, and H contains a loop.

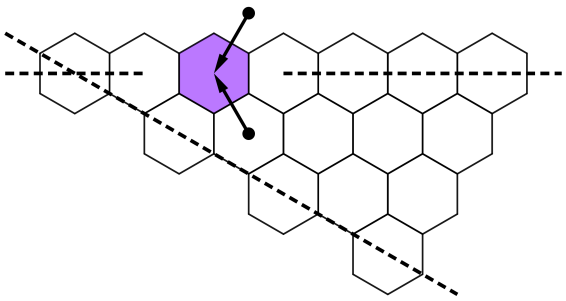

Figure 4.3: The amended sandpile rule: two chips enter the purple hex when the indicated vertex fires.

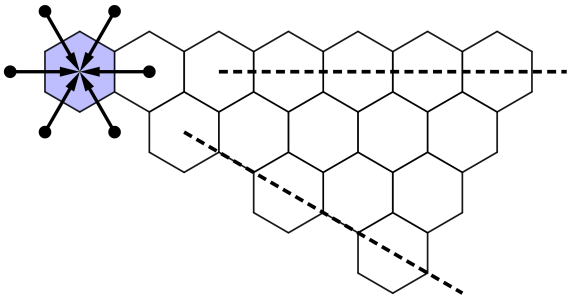

Figure 4.4: The stabilizer of the central vertex has cardinality 12. The stabilizer of its unique neighbor in H has cardinality 2. Thus $12/2 = 6$ chips will enter the central vertex when its neighbor in H donates 1.

By Lemma 3.1.1, we are free to choose the order in which we fire the vertices of G in stabilizing σ. Define a sequence of chip configurations $\sigma = \sigma_0, \sigma_1, \ldots, \sigma_m = \sigma^\circ$ in which σ_{i+1} is obtained by selecting an unstable vertex and firing *every* vertex in its orbit. Firing all vertices in an orbit preserves symmetry of the chip configuration. By the symmetry of σ_i, $x \sim w$ implies x is unstable if and only if w is unstable (and firing x cannot stabilize w), so σ_i is legal for all $0 \leq i \leq m$.

Now define $\tau := \sigma|_H$, and $\tau_i := \sigma_i|_H$. We wish to show that

(**1**) The sequence $\tau = \tau_0, \tau_1, \ldots, \tau_m$ is a successive sequence of configurations on H (according to the modified rule).

(**2**) It ends in a stable configuration $\tau_m = \tau^\circ$.

If we can show (**1**) and (**2**) we will be done. In fact τ_m is certainly stable, since it is the restriction of a stable configuration, with vertex outdegrees the same in both graphs by construction. So (**2**) follows from (**1**), and this is all we need prove.

Assume $\tau_0, \tau_1, \ldots, \tau_k$ is a successive sequence. We will show τ_{k+1} is successive to τ_k. Suppose σ_{k+1} is obtained from σ_k by firing $R \cdot x_k$, where x_k is chosen to be the equivalence

class representative in C. From the construction of H, the outdegree of x_k in H is the same as its outdegree in G, so when x_k fires in H, the same number of chips leave it as leave x_k in G. We can turn our attention to incoming chips.

Suppose (x_k, z) is an edge in H, with $z \neq s$. The number of chips arriving at vertices of $R \cdot z$ in G when $R \cdot x_k$ is fired is $|R \cdot x_k| \cdot d_{x_k,\overline{z}}$. Thus by symmetry the number of chips arriving at the equivalence class representative z, when $R \cdot x_k$ is fired in G (i.e. between σ_k and σ_{k+1}) is

$$\frac{|R \cdot x_k| \cdot d_{x_k,\overline{z}}}{|R \cdot z|} = d_{x_k,\overline{z}} \cdot \frac{|R_z|}{|R_{x_k}|}. \tag{4.2}$$

When we compare Eq. (4.2) to the number of chips arriving at z when x_k is fired in H, there are two possibilities: either $|R_{x_k}| \leq |R_z|$, or $|R_{x_k}| > |R_z|$. In the first case, by construction of H, the multiplicity of the edge (x_k, z) is exactly the right hand side of Eq. (4.2). In the second case, the multiplicity of (x_k, z) in H is simply $d_{x_k,\overline{z}}$, however the modified sandpile rule now applies, so we multiply the incoming chips by $\frac{|R_z|}{|R_{x_k}|}$, which again makes the number of incoming chips equal to the right hand side of Eq. (4.2).

We have shown that beginning with τ_k and firing x_k in H results in exactly the number of chips leaving and arriving at every vertex of W between σ_k and σ_{k+1}. Thus τ_{k+1} is successive to τ_k as promised. □

Chapter 5

CONVERGENCE

5.1 The hexagonal grid

Once you have seen the visual similarities between the identity elements on grids of various scales, it seems very likely that they converge to some fractal-like object. Likewise, in adding chips to a central pile and stabilizing, generally after the graph reaches a certain threshold a picture starts to appear, becoming more stable first around the outside, and gradually beginning to resolve in the interior. Various attributes of the pictures, such as the presence of a large square on the interior of the square grid identity, or the presence of a large central region of 0's surrounded by 4's in the case of the hex grid identity, seem certain. However no one has so far been able to prove anything about any particular trait such as these. Nonetheless, in 2013 Pegden and Smart [14] showed that—at least in the case of the square grid central pile—the stabilized and properly rescaled pile of n vertices on the origin converges to an object that may be studied. In this section we use the methods of [14] to extend this result to the central pile on the hex grid.

To formalize this infinite grid, let

$$e_1 := (1, 0), \qquad e_2 := \left(\frac{1}{2}, \frac{\sqrt{3}}{2}\right) \tag{5.1}$$

Then the hex grid is simply $\mathcal{H} = \mathrm{Span}_{\mathbb{Z}}(e_1, e_2)$. This set is sometimes called the Eisenstein integers. They are made a graph in the natural way by connecting vertices x, y with an edge if $|x - y| = 1$. Let $s_n : \mathcal{H} \to \mathbb{N}$ be the chip configuration obtained from starting with n chips at the origin and stabilizing. Since $\mathcal{H}$ has no sink, we need the following proposition to ensure this sequence is well-defined.

Proposition 5.1.1. *If G is a connected graph on an infinite set of vertices, then any finite chip configuration will stabilize.*

Remark 5.1.2. It is essential that G is a *graph*, i.e. if we are using the digraph definition for sandpiles we introduced earlier, then $(x, w) \in E$ implies $(w, x) \in E$.

We also assume G has no vertices of infinite degree, though such an assumption is not necessary. Such vertices would effectively be sinks, and by identifying all infinite-degree vertices in G and noting that the sandpile structure of such a graph remains unchanged, this modified graph is covered under Lemma 3.1.4.

Proof. Assume for a contradiction there is $\sigma : V \to \mathbb{N}$ that does not stabilize; in that case there is an infinite sequence $\sigma = \sigma_0, \sigma_1, \sigma_2, \ldots$ of successive chip configurations, with corresponding sequence $x_1, x_2, \ldots$ of vertices, where x_i fires between σ_{i-1} and σ_i.

Case 1: There is $x^* \in V$ not appearing in the list $x_1, x_2, \ldots$. Proceeding by induction on $N_\sigma = \sum_{x \in V} \sigma(x)$, the base case $N_\sigma = 1$ is clear since there are no adjacent degree 1 vertices on an infinite connected graph. Inductively suppose that all chip configurations with less than N_σ chips stabilize after a finite number of steps. If $\sigma(x^*) > 0$, the configuration σ' that has a chip removed from x^* and is elsewhere identical to σ has $N_\sigma - 1$ chips and therefore must stabilize. This implies σ must also stabilize after a finite number of steps. Likewise none of the neighbors of x^* can fire without donating a chip to x^* that will never leave it. Since the neighbors of x^* can never fire they too must hold no chips. Continuing in this way outward, no vertex can fire or hold any chips or else the configuration will stabilize. Thus since $N_\sigma > 0$, σ must stabilize.

Case 2: Every vertex in V fires. Define $A_\sigma \subset V$, the *extended support* of σ, to be the set of all vertices x for which $\sigma(x) \neq 0$ (the support of σ) together with any vertices *all* of whose neighbors are in the support of σ. Firing a vertex can't remove a vertex from the extended support, so $A_{\sigma_0}, A_{\sigma_1}, A_{\sigma_2}, \ldots$ is a nondecreasing sequence of subsets.

Choose $y, w \in V$ such that $d(y, w) > 2N_\sigma + 1$. Such a pair exists by the earlier assumption that G had no vertices of infinite degree. Now take $M \geq 1$ large enough so that every vertex

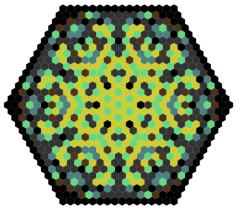

Figure 5.1: s_n on $\mathcal{H}$ when $n = 2 \times 10^3$

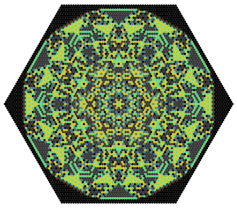

Figure 5.2: s_n when $n = 2 \times 10^4$

in a yw path has fired at least once before σ_M.At least half of the vertices on that yw path must contain at least one chip, giving the contradiction

$$N_\sigma = \sum_{x \in V} \sigma_M(x) \geq \sum_{x \in P} \sigma_M(x) \geq \frac{1}{2}|P| \geq N_\sigma + 1.$$

The leftmost equality above comes from the fact that N_σ is invariant with respect to chip firing on sinkless digraphs.

□

Independent of n, the visual representation of s_n appears to be the same image, but scaling with n and with greater resolution as the larger image allows (see Figs. 5.1 to 5.4). Given that the number of hexes within a given radius r is $O(r^2)$, if we assume there is indeed convergence of a certain type, we will need to consider rescaled configurations $\overline{s}_n : h\mathcal{H} \to \mathbb{N}$ given by $\overline{s}_n(x) := s_n(h^{-1}x)$, where $h := n^{-1/2}$. We will want to sometimes consider $\overline{s}_n(x)$ as a function $\mathbb{R}^2 \to \mathbb{R}$, which we can do by extending the domain via nearest neighbor interpolation. That is, consider $\overline{s}_n(x) : \mathbb{R}^2 \to \mathbb{R}$ as given by

$$\overline{s}_n(x) = s_n(\lfloor h^{-1}x \rceil)$$

We use $\lfloor \cdot \rceil$ here to mean first expressing $x = c_1e_1 + c_2e_2$ (as defined in Eq. (5.1)), and then rounding c_1 and c_2 to the nearest integer, down in the case of ties.

The type of convergence we will use is weak-* convergence. Recall $L^\infty(\mathbb{R}^2)$ is the space of bounded measurable functions on $\mathbb{R}^2$. Certainly the functions $\overline{s}_n : \mathbb{R}^2 \to \mathbb{R}^2$ fall into

Figure 5.3: s_n with $n = 2 \times 10^5$

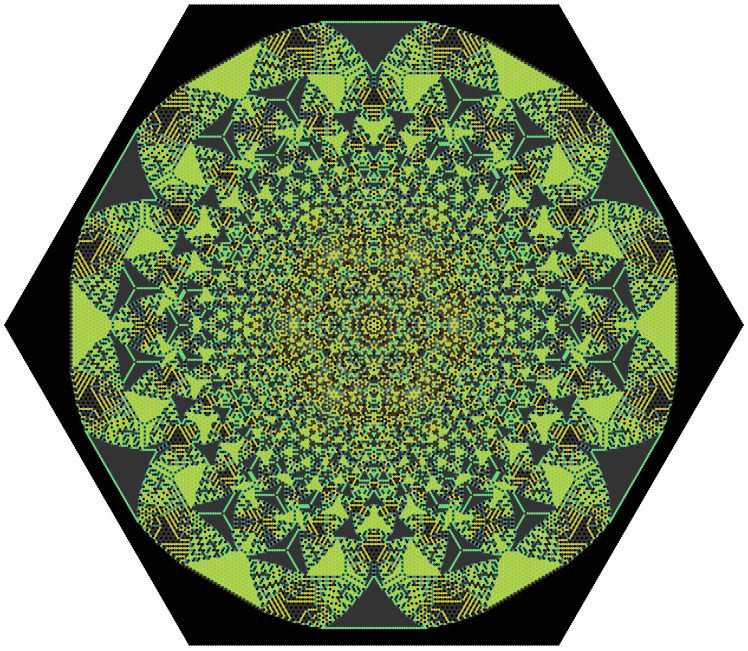

this category. Additionally, $C_0(\mathbb{R}^2)$ denotes the space of all continuous functions on $\mathbb{R}^2$ with compact support, i.e., $f \in C_0(\mathbb{R}^2)$ implies that $\overline{E_f}$ is compact, where $E_f := \{x \in \mathbb{R}^2 : f(x) \neq 0\}$.

Definition 5.1.3. A sequence of functions $f_n \in L^\infty(\mathbb{R}^2)$ converges weakly-* to a function $f \in L^\infty(\mathbb{R}^2)$ if

$$\int_{\mathbb{R}^2} f_n \varphi \, dA \longrightarrow \int_{\mathbb{R}^2} f \varphi \, dA \qquad \text{as } n \to \infty$$

for all functions $\varphi \in C_0(\mathbb{R}^2)$.

With this we are ready to state the main result of this section.

Conjecture 5.1.4. *The rescaled chip configurations $\overline{s}_n(x)$ converge weakly-* to a function $s \in L^\infty(\mathbb{R}^2)$ as $n \to \infty$.*

Our partial proof of Conjecture 5.1.4 follows closely the methods laid out in [14]. We extend to $\mathcal{H}$ their result establishing convergence along subsequences. The only part that remains, which is straightforward from their paper, is to show that the limiting objects are unique, independent of choice of subsequence.

5.2 Differential operators, notation, and analytic estimates

We define the open ball $B_r(x) := \{y \in \mathbb{R}^2 : |y - x| < r\}$, with $B_r := B_r(0)$. We will use $\partial\Omega$ to mean the boundary of an open set Ω, and if $E \subseteq h\mathcal{H}$, we write

$$\partial^h E := \{y \in h\mathcal{H} \setminus E \, : \, d(y, E) = h\},$$

that is, the neighbors of vertices in E that are not themselves in E. We will occasionally write δ_x to mean the Kronecker delta: the indicator function of x, where $x \in V$ and V is considered as a subset of $\mathbb{R}^2$.

The two-dimensional Laplacian

$$\Delta u := \frac{\partial^2}{\partial x_1^2} u + \frac{\partial^2}{\partial x_2^2} u$$

and its equivalent on $h\mathcal{H}$, the 7-point discrete mesh Laplacian

$$\Delta^h u(x) := \frac{2}{3h^2} \sum_{y \in N(x)} u(y) - u(x).$$

will both be at the center of the proof. The coefficient of $2/3$, which makes this operator distinct from the graph Laplacian, was included so that $\Delta^h \varphi \to \Delta \varphi$ locally uniformly as $h \to \infty$. A derivation of Δ^h can be found in [3]. Indeed the discrete operator Δ^h shares various analogous properties with the continuous Laplacian.

Lemma 5.2.1. *If $f, g : h\mathcal{H} \to \mathbb{R}$, with $w := \min\{f, g\}$, then whenever $w(x) = f(x)$ we have*

$$\Delta^h w(x) \leq \Delta^h f(x)$$

Proof. This is a simple consequence of monotonicity: increasing $f(x)$ decreases $\Delta^h f(x)$ whereas increasing $f(y)$ for $y \in N(x)$ increases $f(x)$.

$$\Delta^h w(x) = \frac{2}{3h^2} \sum_{y \in N(x)} w(y) - w(x) \leq \frac{2}{3h^2} \sum_{y \in N(x)} f(y) - f(x) = \Delta^h f(x).$$

□

Proposition 5.2.2 (Discrete Maximum Principle)**.** *Suppose $f, g : h\mathcal{H} \to \mathbb{R}$, and $E \subset h\mathcal{H}$ is a finite set of vertices. If $\Delta^h f \geq \Delta^h g$ in E, then $\max_{\partial^h E}(f - g) \geq \max_E(f - g)$.*

We will want to consider the fundamental solution $\Phi_n : h\mathcal{H} \to \mathbb{R}$ to Δ^h satisfying

$$\Delta^h \Phi_n(x) = \begin{cases} -n & : x = (0,0) \\ 0 & : x \neq (0,0) \end{cases} \tag{5.2}$$

for all $x \in h\mathcal{H}$. To that end we will need the following theorem.

Theorem 5.2.3. *The function $\Phi_n : h\mathcal{H} \to \mathbb{R}$ satisfying Eq.* (5.2) *exists. Up to addition of a constant it is unique.*

Proof. Uniqueness, which we will not prove, is a consequence of Proposition 5.2.2 once an asymptotic estimate has been established, since the difference of two such functions on a large circle must be small if they share the same value at 0.

It will be enough to show existence when $n = 1$, since it is easily verified that $\Phi_n(x) := \Phi_1(h^{-1}x)$ also satisfies Eq. (5.2).

For existence, we begin by examining how the discrete Laplacian on $\mathcal{H}$ interacts with the Fourier transform. If $F : \mathcal{H} \to \mathbb{R}$, we may define

$$\widehat{F}(\theta, \varphi) := \sum_{(x,y) \in \mathcal{H}} F(x,y) e^{-i\theta x} e^{-i\varphi y}.$$

We can break $\Delta^1 F$ into a sum of four terms: the value at the central vertex, and the three pairs of vertices on opposite sides of this vertex. A variable substitution then allows us to write the Fourier transform in terms of cosine multiples of the original function:

$$\begin{aligned}
\widehat{\Delta^1 F} &= \frac{2}{3}\sum_{(x,y)\in\mathcal{H}}\left(\sum_{(w,z)\in N(x,y)} F(w,z) - F(x,y)\right) e^{-i\theta x}e^{-i\varphi y} \\
&= -4\sum_{(x,y)\in\mathcal{H}} F(x,y)e^{-i\theta x}e^{-i\varphi y} \\
&\quad + \frac{2}{3}\sum_{(w,z)\in\mathcal{H}} F(w,z)e^{-i\theta w}e^{-i\varphi z}\left(e^{i\theta}+e^{-i\theta}\right) \\
&\quad + \frac{2}{3}\sum_{(w,z)\in\mathcal{H}} F(w,z)e^{-i\theta w}e^{-i\varphi z}\left(e^{i(\frac{\theta}{2}-\frac{\sqrt{3}\varphi}{2})}+e^{-i(\frac{\theta}{2}-\frac{\sqrt{3}\varphi}{2})}\right) \\
&\quad + \frac{2}{3}\sum_{(w,z)\in\mathcal{H}} F(w,z)e^{-i\theta w}e^{-i\varphi z}\left(e^{i(\frac{\theta}{2}+\frac{\sqrt{3}\varphi}{2})}+e^{-i(\frac{\theta}{2}+\frac{\sqrt{3}\varphi}{2})}\right) \\
&= \widehat{F}\left(\frac{8}{3}\cos\left(\frac{\theta}{2}\right)\cos\left(\frac{\sqrt{3}}{2}\varphi\right)+\frac{4}{3}\cos(\theta)-4\right)
\end{aligned}$$

Combining this with Eq. (5.2) gives us the requirement that $\widehat{\Delta^1\Phi_1} = -1$ and therefore

$$\widehat{\Phi_1} = \left(4-\frac{4}{3}\cos(\theta)-\frac{8}{3}\cos\left(\frac{\theta}{2}\right)\cos\left(\frac{\sqrt{3}}{2}\varphi\right)\right)^{-1}.$$

Now apply the inverse Fourier transform to get the formal expression

$$\Phi_1(x,y) = \frac{3}{16\pi^2}\int_{-\pi}^{\pi}\int_{-\pi}^{\pi}\frac{e^{i(\theta x+\varphi y)}d\theta\,d\varphi}{3-\cos(\theta)-2\cos\left(\frac{\theta}{2}\right)\cos\left(\frac{\sqrt{3}}{2}\varphi\right)}.$$

Unfortunately, this integral does not converge for any (x,y). When $\theta\sim 0$ the Taylor expansion of the denominator will be $O(\theta^2)$. But since modification of a function by a constant does not affect Δ^1, we may try and allow $\Phi_1(0,0)=0$ by setting

$$\Phi_1(x,y) := \frac{3}{16\pi^2}\int_{-\pi}^{\pi}\int_{-\pi}^{\pi}\frac{e^{i(\theta x+\varphi y)}-1}{3-\cos(\theta)-2\cos\left(\frac{\theta}{2}\right)\cos\left(\frac{\sqrt{3}}{2}\varphi\right)}\,d\theta\,d\varphi.$$

The function Φ_1 converges absolutely for all $(x,y)\in\mathbb{R}^2$, and the Laplacian Δ^1 depends on (x,y) and so commutes with the integrals.

$$\begin{aligned}\Delta^1\Phi_1(x,y) &= \frac{3}{16\pi^2}\int_{-\pi}^{\pi}\int_{-\pi}^{\pi}\Delta^1\left(\frac{e^{i(\theta x+\varphi y)}-1}{3-\cos(\theta)-2\cos\left(\frac{\theta}{2}\right)\cos\left(\frac{\sqrt{3}}{2}\varphi\right)}\right)d\theta\,d\varphi \\ &= -\frac{1}{4\pi^2}\int_{-\pi}^{\pi}\int_{-\pi}^{\pi}e^{i(\theta x+\varphi y)}\,d\theta\,d\varphi\end{aligned}$$

The latter integral is checked to satisfy Eq. (5.2), so we have shown existence of Φ_1 and therefore Φ_n.

□

If we desired a more explicit formula, we could continue integrating the function that appeared in Theorem 5.2.3. For our purposes it is only important that $\Phi_n \to \Phi$ locally uniformly as $n \to \infty$, where

$$\Phi = -\frac{1}{2\pi}\log|x|.$$

From [10], the following is a consequence of the Arzela-Ascoli theorem and estimates for the discrete Laplacian.

Proposition 5.2.4. *Suppose $u_n : h\mathcal{H} \to \mathbb{R}$ is a sequence of functions each satisfying*

$$\max_{B_1}|u_n| \le C \qquad \textit{and} \qquad \max_{B_1}|\Delta^h u_n| \le C$$

for some C independent of n. If $u_{n_1}, u_{n_2}, \ldots$ is any subsequence, then there is a further subsequence $u_{n_{k_1}}, u_{n_{k_2}}, \ldots$ together with a function $u \in C(B_1)$ with $u_{n_{k_j}} \to u$ locally uniformly as $j \to \infty$.

Finally we will need a few standard facts about the continuous Laplacian from analysis. When $\Omega \subseteq \mathbb{R}^2$ is an open set and $s \in L^\infty(\Omega)$, then $u \in C(\Omega)$ is a *weak solution* of the Poisson equation

$$\Delta u = s$$

if

$$\int_\Omega u\Delta\varphi\, dA = \int_\Omega s\varphi\, dA$$

for all test functions $\varphi \in C_0^\infty(\Omega)$. The two formulations are equivalent via the integration by parts formula when $u \in C^2(\Omega)$, and when u is not smooth the equivalence is taken as a definition.

Proposition 5.2.5. *If $s \in L^\infty(B_1)$ and $u \in C(B_1)$ is a weak solution of $\Delta u = s$ in B_1, then u is twice differentiable and $\Delta u = s$ almost everywhere in B_1.*

Proposition 5.2.6. *If $s \in L^\infty(B_1)$ and $u \in C(B_1)$ is a weak solution of $\Delta u = s$ in B_1, and if in addition $\sup_{B_1} u > \sup_{\partial B_1} u$, then the set*

$$\{x \in B_1 \,:\, u(x) > \sup_{\partial B_1} u, \quad u \text{ is twice differentiable at } x, \quad D^2u(x) < 0\}$$

has positive measure.

5.3 Convergence along subsequences

Recall for a given initial configuration on $\mathcal{H}$, the odometer function is denoted $v : \mathcal{H} \to \mathbb{N}$. The following property of v, formulated in [6] and adapted here to $\mathcal{H}$, will be crucial in the proof of Conjecture 5.1.4.

Proposition 5.3.1 (Least Action Principle)**.** *If v is the odometer function associated with initial configuration σ on $\mathcal{H}$, then v is the pointwise minimum among all functions $u : \mathcal{H} \to \mathbb{N}$ satisfying*

$$\sigma + \frac{3}{2}\Delta^1 u \leq 5 \tag{5.3}$$

Proof. Let $\mathcal{F}$ be the family of all functions $\mathcal{H} \to \mathbb{N}$ satisfying Eq. (5.3). Observe $\sigma + \frac{3}{2}\Delta^1 u$ corresponds to the (not necessarily legal) configuration achieved by starting with σ and firing each vertex x exactly $u(x)$ many times. The fact is immediately checked when $u(x)$ is the characteristic function of a single vertex, and the general fact holds by linearity of Δ^1. If

$u \in \mathcal{F}$, then an odometer of u results in a stable, though not necessarily legal, configuration. In fact $\mathcal{F}$ exactly characterizes stable (possibly illegal) ending configurations on $\mathcal{H}$.

The condition Eq. (5.3) may be rewritten as

$$\Delta^1 u \leq f,$$

where $f := \frac{2}{3}(5-\sigma)$. For $u, w \in \mathcal{F}$ and arbitrary $x \in \mathcal{H}$, by symmetry suppose $u(x) \leq w(x)$. Letting $z := \min\{u, w\} \in \mathcal{F}$, by monotonicity of Δ^1,

$$\Delta^1 z(x) = \frac{2}{3} \sum_{y \in N(x)} z(y) - z(x) \leq \frac{2}{3} \sum_{y \in N(x)} u(y) - u(x) = \Delta^1 u(x) \leq f.$$

Thus $z \in \mathcal{F}$. Taking $v_{\mathcal{F}} := \min_{\mathcal{F}}\{u \in \mathcal{F}\}$, which is well-defined since every function in $\mathcal{F}$ is nonnegative, we have shown $v_{\mathcal{F}} \in \mathcal{F}$. By our earlier remark $v_{\mathcal{F}}$ corresponds to a stable configuration, and for any $u \in \mathcal{F}$ with $u \neq v_{\mathcal{F}}$, $v < u$ and so Lemma 3.1.1 implies u results in an illegal configuration. Since the odometer function $v \in \mathcal{F}$, yet v results in a legal configuration, $v \neq u$ for any $u \neq v_{\mathcal{F}}$. Thus $v = v_{\mathcal{F}}$. □

If our initial configuration is the pile with n chips at the origin, denoted $n\delta_0$, and the stabilization is denoted s_n as we wrote earlier, then we denote the odometer function v_n. Using the same reasoning as in the proof of Proposition 5.3.1, the relationship of the three in terms of the discrete Laplacian Δ^1 is easily checked to be

$$n\delta_0 + \frac{3}{2}\Delta^1 v_n = s_n. \tag{5.4}$$

Though the form of the last two equations might suggest using the graph Laplacian rather than Δ^1 as defined, we will see that the continuity in the limit with the continuous Laplacian Δ is what is more important in the proof.

Define the rescaled odometer function $\overline{v}_n : h\mathcal{H} \to \mathbb{R}_{\geq 0}$ by

$$\overline{v}_n(x) := h^2 v_n(h^{-1}x),$$

and let

$$\overline{w}_n := \overline{v}_n - \Phi_n.$$

From here by linearity of Δ^h, Eq. (5.4), and definition of Φ_n,

$$\begin{aligned} \Delta^h \overline{w}_n &= \Delta^h \overline{v}_n - \Delta^h \Phi_n \\ &= \Delta^1 v_n(h^{-1}x) + n\delta_0 \\ &= s_n(h^{-1}x) - n\delta_0 + n\delta_0 \\ &= \overline{s}_n \end{aligned} \tag{5.5}$$

over all of $\mathbb{R}^d$. This suggests Poisson's equation and is the reason for the excursion into elliptic PDE theory. We want to show that this relationship continues in the limit as $n \to \infty$, which is the subject of the next lemma. We will first show that an appropriate limit exists along subsequences, and then later show uniqueness of the limit.

Lemma 5.3.2. *For every sequence $n_1, n_2, \ldots$, there is a subsequence $n_{k_1}, n_{k_2}, \ldots$ together with functions $w \in C(\mathbb{R}^2)$ and $s \in L^\infty(\mathbb{R}^2)$ such that*

1. *$\overline{w}_{n_{k_j}} \to w$ locally uniformly.*

2. *$\overline{s}_{n_{k_j}} \to s$ weakly-$*$.*

3. *$\Delta w = s$*

The proof relies partially on an *a priori* bound $\{s_n > 0\} \subseteq B_{h^{-1}R}$ for some $R > 0$ not depending on n. The best such bound comes from [12], though we only need the existence of a finite R. This should not be surprising since $B_{h^{-1}R}$ has $O(n^2)$ cells, and so taking R sufficiently large would guarantee inclusion unless the support were somehow growing in a "one-dimensional" way.

Proof. Since $\{s_n > 0\} \subseteq B_{h^{-1}R}$, which implies $\{\overline{s}_n > 0\} \subseteq B_R$, therefore $\{\overline{v}_n > 0\} \subseteq B_R$. Together with nonnegativity of $\overline{v}_n$, this implies

$$\overline{w}_n \geq -\Phi_n \quad \text{in } h\mathcal{H} \qquad \text{and} \qquad \overline{w}_n = -\Phi_n \quad \text{on } h\mathcal{H} \setminus B_R.$$

Set $E := h\mathcal{H} \cap B_R$ and define the test functions

$$\varphi(x) := |x|^2 - (R+h)^2 + \inf_{\partial^h E} -\Phi_n \qquad \text{and} \qquad \psi(x) := \sup_{\partial^h E} -\Phi_n.$$

These satisfy

$$\varphi \leq \overline{w}_n \leq \psi \qquad \text{on } \partial^h E$$

and

$$4 = \Delta^h \varphi \geq \Delta^h \overline{w}_n \geq \Delta^h \psi = 0.$$

Thus by the discrete maximum principle (Proposition 5.2.2), it follows that

$$|x|^2 - (R+h)^2 + \inf_{\partial^h E} -\Phi_n \leq \overline{w}_n \leq \sup_{\partial^h E} -\Phi_n$$

over all of B_R. Choosing $R' > R$ and noting that convergence of $\Phi_n \to \Phi$ is uniform in $B_{R'} \setminus B_R$, and so the sequence $\|\overline{w}_n\|_{L^\infty(B_{R'})}$ is uniformly bounded, i.e.

$$0 < \|\overline{w}_n\|_{L^\infty(B_{R'})} \leq C.$$

Here $C = C(R')$ is independent of n. By Eq. (5.5) and stability of s_n we also have the uniform bound $|\Delta^h \overline{w}_n| = \overline{s}_n \leq 5$. Proposition 5.2.4 implies there exists a sequence n_k and a function $w \in C(\mathbb{R}^2)$ such that $\overline{w}_{n_k} \to w$ locally uniformly as $j \to \infty$. A uniformly bounded sequence of functions always converges weakly-$*$ along subsequences [5, page 7], so select a further subsequence and $s \in L^\infty(\mathbb{R}^2)$ such that $\overline{s}_{n_k}$ converges weakly-$*$ to s as $k \to \infty$.

Taking any $\varphi \in C_0^\infty(\mathbb{R}^2)$,

$$\int_{\mathbb{R}^2} \varphi s_{n_k} \, dA = \int_{\mathbb{R}^2} \varphi \Delta^{h_k} \overline{w}_{n_k} \, dA = \int_{\mathbb{R}^2} (\Delta^{h_k} \varphi) \overline{w}_{n_k} \, dA, \tag{5.6}$$

If the second inequality is reminiscent of the integration by parts formula, it is nonetheless a discrete analog, which can be confirmed by writing Δ^{h_k} as a finite sum, commuting the sum with the integral, and making the proper changes of variables to align the translated $\overline{w}_{n_k}$ terms.

On any bounded neighborhood of the support of φ, the locally uniform convergence $\overline{w}_{n_k}$ will be uniform. The convergence $\Delta^h\varphi \to \Delta\varphi$ is also uniform as $n \to \infty$. Therefore Eq. (5.6) gives us

$$\lim_{k\to\infty} \int_{\mathbb{R}^2} s_{n_k}\varphi = \int_{\mathbb{R}^2} w\Delta\varphi.$$

Since $s_n \to s$ weakly-$*$, by definition the left hand side above converges to $\int_{\mathbb{R}^2} s\varphi$. We therefore have by definition that w is a weak solution of $\Delta w = s$. □

We have not so far shown that Conjecture 5.1.4 is true, only that there is convergence along subsequences. It remains to show only that the limiting s and w are unique irrespective of the subsequences chosen. We will refer the reader to the methods of [14] for the final more technical parts of the proof. However, we hope the previous proof suggests to the reader that the methods used are not special to the integer lattice, and may be adapted to other lattices accordingly.

Figure 5.4: s_n with $n = 2 \times 10^6$

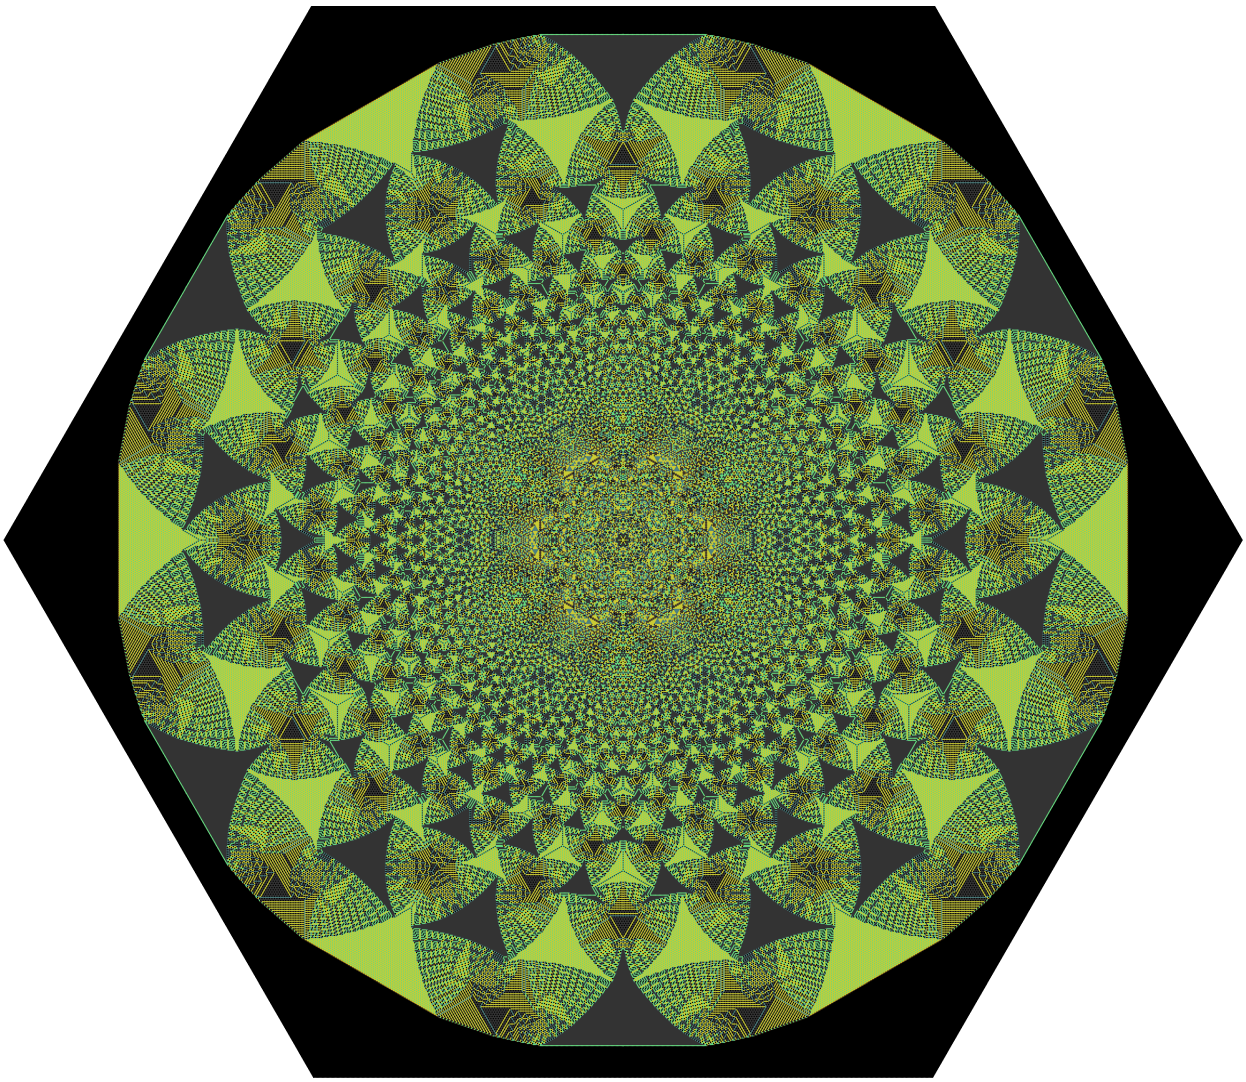

Chapter 6

CONJECTURES AND EXPERIMENTAL DATA

Though we have shown the rescaled central pile converges, the various sources attest that saying anything at all about the object to which it converges has been extremely difficult. Nonetheless, we have gathered experimental data that suggests various conjectures.

Recall we defined A_σ, the extended support of sigma, to be the support of σ together with any vertex x such that $N(x)$ is contained in the support of σ. Let s_n again be the stabilized central pile with n chips. Let $A_n := A_{s_n}$. Let $B_n = A_{\epsilon_n}$, where ϵ_n is the recurrent identity configuration on the hex grid of side length n.

Conjecture 6.0.1. *Let*

$$p_n := \frac{\{x \in A_n : s_n(x) = 1\}}{|A_n|} \qquad \textit{and} \qquad q_n := \frac{\{x \in B_n : \epsilon_n(x) = 1\}}{|B_n|}$$

Then $p_n \to 0$ *and* $q_n \to 0$ *as* $n \to \infty$.

The above conjecture is motivated by the fact that when $n = 2 \times 10^6$, $p_n \approx 0.01017$. When $n = 400$, we have also experimentally verified that $q_n \approx 0.000873$.

Conjecture 6.0.2.

$$\lim_{n\to\infty} \frac{1}{|A_n|} \int_{\mathbb{R}^2} s_n \, dA = \frac{n\sqrt{3}}{2|A_n|} = 3$$

The first equality is a direct computation based on the area of a single hex.

The way the support of s_n increases with $n \to \infty$ is by expanding outward a single ring at a time. Naturally the average value of s_n will then steadily rise before falling sharply as the support expands. We have recorded that the average value is always above 3 before this drop, and below 3 after the drop, and that the difference in these two values immediately before and after the drop generally (though not monotonically) appears to decrease. At $n = 2 \times 10^6$, the average chip value is approximately 3.03731.

Remark 6.0.3. This conjecture was used as a way to dynamically increase the size of matrices used to store central pile chip configurations. Without knowing exactly how big a central pile would ultimately become, the estimate was useful in populating matrices that were roughly the correct size.

BIBLIOGRAPHY

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

[2] G. Benkart, C. Klivans, and V. Reiner, *Chip firing on dynkin diagrams and mckay quivers*, Mathematische Zeitschrift, 290 (2018), pp. 615–648.

[3] K. Crane, F. de Goes, M. Desbrun, and P. Schröder, *Digital geometry processing with discrete exterior calculus*, in ACM SIGGRAPH 2013 courses, SIGGRAPH '13, New York, NY, USA, 2013, ACM.

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

[5] L. Evans and C. B. of the Mathematical Sciences, *Weak Convergence Methods for Nonlinear Partial Differential Equations: [expository Lectures from the CBMS Regional Conference Held at Loyola University of Chicago, June 27-July 1, 1988]*, CBMS Regional Conference Ser. in Mathematics Series, American Mathematical Soc., 1990.

[6] A. Fey, L. Levine, and Y. Peres, *Growth rates and explosions in sandpiles*, J. Stat. Phys., 138 (2010), pp. 143–159.

[7] A. E. Holroyd, L. Levine, K. Mészáros, Y. Peres, J. Propp, and D. B. Wilson, *Chip-firing and rotor-routing on directed graphs*, in In and out of equilibrium. 2, vol. 60 of Progr. Probab., Birkhäuser, Basel, 2008, pp. 331–364.

[8] H. M. (https://math.stackexchange.com/users/14366/henning makholm), *Show $\mathbb{N}^{\times}$ with lexicographic ordering has the least upper bound property (any nonempty bounded subset has a supremum).* Mathematics Stack Exchange. URL:https://math.stackexchange.com/q/3267798 (version: 2019-06-19).

[9] C. J. Klivans, *The Mathematics of Chip-firing*, CRC Press, 11 2018.

[10] H.-J. Kuo and N. S. Trudinger, *Estimates for solutions of fully nonlinear discrete schemes*, in Trends in partial differential equations of mathematical physics, vol. 61 of Progr. Nonlinear Differential Equations Appl., Birkhäuser, Basel, 2005, pp. 275–282.

[11] S. Lang, *Algebra*, Springer, New York, NY, 2002.

[12] L. Levine and Y. Peres, *Strong spherical asymptotics for rotor-router aggregation and the divisible sandpile*, Potential Anal., 30 (2009), pp. 1–27.

[13] L. Levine and J. Propp, *What is ... a sandpile?*, Notices Amer. Math. Soc., 57 (2010), pp. 976–979.

[14] W. Pegden and C. K. Smart, *Convergence of the Abelian sandpile*, Duke Math. J., 162 (2013), pp. 627–642.

[15] L. M. Ward and P. E. Greenwood, *1/f noise*, Scholarpedia, 2 (2007), p. 1537. revision #137265.