The Structure of the Sandpile Group

Akhil Jalan

May 20, 2019

Abstract

The abelian sandpile model is a discrete dynamical system defined on a graph, in which grains of "sand" are placed on the vertices and move along edges. Despite its simple combinatorial description, the sandpile model has surprising connections to a variety of areas, including spectral graph theory, finite group theory, computational complexity, and more. We will survey some of these connections through the lens of the *sandpile group*, a finite abelian group associated to a graph, whose elements can be identified with recurrent states of the sandpile model. After presenting various equivalent formulations of the sandpile group, we will focus on a particular description of the sandpile group, as the cokernel of the graph Laplacian matrix. Under this description, it is (in principle) easy to compute the invariant factor decompositon of the sandpile group, by computing the Smith Normal Form of the reduced Laplacian matrix. Using this technique, we study the invariant factors of the sandpile groups for various families of graphs, especially expanders. We prove lower bounds on the number of trivial invariant factors for families of graphs, such as the hypercube and grid graphs. For more interesting graph families, such as two explicit constructions of expanders, we formulate conjectures about their invariant factors on the base of computer experiments. [1] [2]

Contents

[1] Email: akhiljalan@berkeley.edu

[2] **Acknowledgments**: Many thanks to Prof. Nikhil Srivastava for his patience and advice. Thanks also to Antares Chen for his extensive collaboration and help with checking early computer experiments, as well as Aaron Schild, Brett Kolesnik, Varun Khurana, and Tahsin Saffat for valuable discussions.

1 Introduction

The abelian sandpile model, also called the chip-firing game, is a system in which integers (called grains of sand, or chips) are assigned to the vertices of a graph. Starting from some intial configuration, the model evolves locally according to a firing rule, which a vertex is allowed to *fire* if it has at least as many chips as its degree. When a vertex fires, it sends one chip along each of its incident edges. Thus chips move locally along the graph, with vertices accumulating chips through firings of their neighbors, and then losing chips when they have too many.

In this thesis, we will focus on understanding the sandpile groups of undirected graphs. This chapter will present basic definitions and properties of sandpiles. In chapter 2, we will develop some background on spectral and algebraic graph theory. Chapter 3 presents four characterizations of the sandpile group, from spectral graph theory, algebraic graph theory, combinatorics, and algorithms respectively. Finally, in chapter 4 we will prove that the Smith Normal Form of the graph Laplacian matrix completely determines the sandpile group, through its invariant factor decomposition. As it is both analytically tractable and computable through a simple algorithm, we can employ the Smith Normal Form to prove remarkable facts about sandpiles, such as the fact that every finite abelian group is the sandpile group of a planar graph. Finally, we will prove lower bounds on the number of trivial invariant factors for sandpile groups of certain graphs, such as the Boolean hypercube. We will conclude by presenting computer experiments on the values of Smith invariant factors for explicit families of expander graphs, and use these to make conjectures on the asymptotic behavior of their invariant factors.

1.1 Background

The abelian sandpile model was discovered independently by researchers in several different communities, including statistical physics, probability, and arithmetic geometry [Hol08]. It is just one out of many kinds of discrete processes that one can define on a graph - others include the rotor-routing model, bootstrap percolation, diffusion-limited aggregation, and so on [BL16]. Many such models, including the ones just named, are examples of *abelian networks*, which are discrete processes characterized by the fact that the order of their operations will not affect the final outcome [BL16]. In the case of the sandpile model, one can show that the order of vertices chosen to fire will not affect the resulting stable configuration, if one exists (see section 1.2 for details).

Many distinct versions of the sandpile model are known and studied, depending on the context and questions at hand (see the books [Kli18], [CP18]). Some versions change the firing rules, such as the *dollar game* ([CP18], [Big97]). Others study a Markov process in which an initial chip configuration is randomly chosen, and then additional chips are added to vertices at random ([Lev15], [JLP15]). In computer science, a large body of work concerns the efficient simulation of sandpile processes, or the impossibility thereof (see references in [RS17]). Remarkably, even though efficient simulation algorithms are known for certain families of graphs, the general problem of computing the final configuration of a sandpile on a directed multigraph is NP-complete [FL16].

Given any finite undirected graph G, one can define the *sandpile group* of G in terms of certain states of the sandpile model. Like the sandpile model, the sandpile group is studied by researchers in various communities, due to the fact that it appears in so many distinct contexts. Lorenzini ([Lor08]) identifies at least four unique names for the sandpile group in different research areas:

1. The *group of components* in arithmetic geometry.
2. The *sandpile group* in physics.
3. The *Picard group* in the theory of algebraic curves. A celebrated result in this area gives a graph-theoretic analogue of the Riemann-Roch theorem [BN07]. For a detailed exposition of this

result, see [CP18].

4. The *critical group* in algebraic graph theory (see [Big97], [Big99], and section 3.4).

Despite being of interest to so many communities, the sandpile group is difficult to compute in general. Many basic questions remain unanswered. For example, it is unknown whether the limiting shape of a certain chip configuration on $\mathbb{Z}^2$ even exists (see [Hol08], Question 6.2). Worse, for a version of the finite, square grid in $\mathbb{Z}^2$, even the exact shape of the identity element for the sandpile group is unknown ([Hol08])!

Of the positive results that do exist on sandpile groups, many utilize the Smith Normal Form of the graph Laplacian ([Lor08], [RMW93], [Bai03], [JNR03]). The Smith Normal Form is an invariant of integer matrices, which can be used to compute the invariant factors of the sandpile group (see section 4). A related work uses methods from random matrix theory to study the sandpile groups of Erdos-Renyi graphs [Woo17]. Using a version of the moments method for random variables, they are able to prove exact asymptotic probabilities for the Sylow p-subgroups of the sandpile group on $G(n, q)$ (the Erdos-Renyi random graph), as $n \to \infty$.

1.2 Sandpiles

In this section, we introduce the the abelian sandpile model (also referred to as chip-firing game [Hol08], or the dollar game [CP18]). We discuss the notion of stabilization with/without a sink, and prove the *abelian property*, which justifies the name *abelian* sandpile.

Throughout this section, $G = (V, E)$ will denote a finite, undirected, connected graph. If $v \in V$, then $d_v = deg(v) = |\{e \in E : v \in e\}|$ is the degree of v. Note that if G is a multigraph, then the degree of v counts every incident edge to v, including multi-edges and self-loops.

Definition 1.2.1. Let $G = (V, E)$ be a graph. A **sandpile** (or chip configuration, or dollar configuration) is a function $\sigma : V \to \mathbb{Z}$.

Note: We will usually write $\sigma \in \mathbb{Z}^V$, since we only ever deal with finite graphs, and this notation emphasizes the view that σ is a vector of integers indexed by V.

Definition 1.2.2. Given a graph G and sandpile σ, a vertex $v \in V$ is **stable** if and only if $\sigma_v < d_v$ (that is, if v has fewer chips than its degree). If v is not stable, we say it is **unstable**.

An important variation of the sandpile model involves the notion of a *global sink vertex*. This vertex is reachable from every other vertex, always considered stable, and can never be fired. Including a global sink vertex results in several nice properties - in particular, that every configuration can be stabilized.

Definition 1.2.3. Given a graph G, a **global sink** is a fixed vertex $z \in V$ such that:

i. For all $v \in V$, there is a path from v to z

ii. The vertex z is always considered stable.

Remark 1.2.4. Note that even if there exists a vertex $z \in V$ which is globally reachable, we can define a sandpile model on the graph in which z is treated as every other vertex and allowed to fire. Thus the presence of a global sink is a *choice of model*, in the sense that we can designate a global sink vertex for our sandpile model if we so choose, but we can just as well choose to let every vertex fire.

Remark 1.2.5. The presence of two different models of the sandpile - one with a global sink vertex, and the other without - raise several questions about the relationships between the two choices of model. In 3.3, we will show that there is a surprising connection between the sandpile models with

and without a global sink, called the *z-recurrent decomposition.* Moreover, in 4.3 we will show that the sandpile group of a graph (something we define in 3.1) is independent of the choice of sink, provided that the underlying graph is finite, connected, and undirected.

Having defined stability of individual vertices, we can now define stability for sandpiles.

Definition 1.2.6. *Stability of Sandpiles*: Given a graph G and sandpile σ, we say that σ is **stable** if for all $v \in V$, v is stable. If $z \in V$ is a global sink, we say σ is **stable with respect to** z if for all $v \neq z$, v is stable.

Given an unstable sandpile, the sandpile can evolve by redistributing grains of sand (also called *chips*) from over-burdened vertices to their neighbors. This redistribution occurs through a sequence of firing moves, or chip-firings.

Definition 1.2.7. Let $G = (V, E)$ be a graph, and $\sigma \in \mathbb{Z}^V$. Let $v \in V$. Let $A \in \mathbb{Z}^{V \times V}$ denote the adjacency matrix of V, and $a_{i,j}$ the entry of A at row i, entry j.

A **firing move at v** is a map $f_v : \mathbb{Z}^V \to \mathbb{Z}^V$ such that for all $w \in V$,

$$f_v(\sigma)_w = \begin{cases} \sigma_w + a_{v,w} & w \neq v \\ \sigma_v + a_{v,v} - d_v & w = v \end{cases}$$

Remark 1.2.8. From the definition of a firing move, it is clear that if $\Delta = D - A$ denotes the discrete graph Laplacian[3], and $v \in V$, then a firing move at v corresponds to subtracting the v^{th} row (equivalently, the v^{th} column) of the Laplacian from the original configuration. This characterization of chip-firing will be important in 3.1.

Not every firing move is legal. The abelian sandpile only allows a vertex firing if that vertex is unstable.

Definition 1.2.9. Let $G = (V, E)$ be a graph, $\sigma \in \mathbb{Z}^V$, and $v \in V$. We say that a firing move at v is a **legal firing move at** v if $\sigma_v \geq d_v$.

The definitions are analogous for the global sink model, except the sink is never allowed to fire.

Definition 1.2.10. Let $G = (V, E)$ be a graph, $\sigma \in \mathbb{Z}^V$, and $v \in V$. If $z \in V$ is the fixed global sink vertex, we say that a firing move at v is a **legal firing move at** v **with respect to** z if $\sigma_v \geq d_v$, and $v \neq z$.

As mentioned, one motivation for the global sink model of the abelian sandpile is that every configuration can be stabilized. We are are ready to prove this.

Proposition 1.2.11. Let $G = (V, E)$ be a graph, $\sigma \in \mathbb{Z}^V$, and $z \in V$ be the global sink vertex. There exists a finite sequence of legal firing moves from σ resulting in a configuration that is stable with respect to z.

Proof. If σ cannot be stabilized, then every legal sequence of firing moves will be infinite. Thus, some vertex will be fired an infinite number of times - say $v \in V$.

Since z is globally reachable, there is a path from v to z - say $v_0, ..., v_k$, where $v_0 = v$ and $v_k = z$. Every time v fires, it sends a chip to v_1. Therefore, after v_0 fires at most $d_{v_1} - \sigma_{v_1}$ times, vertex v_1

[3]Here, $D \in \mathbb{Z}^{V \times V}$ denotes the diagonal matrix where $D_{v,v} = deg(v)$. We will formally define the Laplacian in section 2.

can be fired once. Thus after v_0 fires at most $(d_{v_1} - \sigma_{v_1})(d_{v_2} - \sigma_{v_2})$ times, v_1 can be fired $(d_{v_2} - \sigma_{v_2})$ times, and thus v_2 can be fired once.

Proceeding by induction, it follows that after v_0 is fired $\prod_{i=1}^{k-1}(\sigma_{v_i} - d_{v_i})$ times, that v_{i-1} can be fired at least once. Thus the sink vertex $z = v_k$ will receive one chip.

Notice that since the sink vertex never fires, that the total number of chips available to fire decreases by 1. Since v fires an infinite number of times, it follows that an infinite number of chips are lost to the sink. However, σ begins with a finite number of chips. This is a contradiction. Thus σ can be stabilized. □

Both the sink and non-sink versions of the model have a crucial abelian property, which is a consequence of the more general swapping principle for firing sequences.

Proposition 1.2.12. Swapping Principle: Suppose that $G = (V, E)$ is a finite, undirected, connected graph (possibly with a global sink) and that $\sigma \in \mathbb{Z}^V$ is an unstable chip configuration. Suppose that $\sigma_0, ..., \sigma_n$ and $\eta_0, ..., \eta_m$ are two configurations obtained through legal firing moves, such that $\sigma = \sigma_0 = \eta_0$. Further, suppose that σ_n is stable (again, possibly with respect to the global sink). Let $k = min(n, m)$. We claim that we can re-arrange the firing sequence of the first k firings of $\sigma_0 \to \sigma_k$, so as to obtain η_k.

Proof. Let $v_1, ..., v_n$ be the sequence in which chips are fired from $\sigma_0 \to \sigma_n$, and similarly let $w_1, ..., w_m$ be the sequence for $\eta_0 \to \eta_m$. Suppose that $i \leq k$ is the least index such that $w_i \neq v_i$. Then $\sigma_\ell = \eta_\ell$ for $\ell < i$. Schematically,

$$\sigma = \sigma_0 \xrightarrow{v_1} \sigma_1 \xrightarrow{v_2} \sigma_2 \xrightarrow{v_3} \cdots \xrightarrow{v_n} \sigma_n$$

$$\sigma = \eta_0 \xrightarrow{w_1} \eta_1 \xrightarrow{w_2} \eta_2 \xrightarrow{w_3} \cdots \xrightarrow{w_m} \eta_m$$

Observe that both v_i, w_i are unstable in the configuration $\sigma_{i-1} = \eta_{i-1}$. Further, since σ_n is stable, we know that w_i must fire at least once in $v_{i+1}, ..., v_n$. Let $j > i$ be the least index such that $v_j = w_i$.

Now, consider the swapped sequence of firings

$$v_1, ..., v_{i-1}, v_j, v_i, v_{i+1}, ..., v_{j-1}, v_{j+1}, ..., v_n$$

Let $\tau_0, \tau_1, ..., \tau_n$ be the chip configurations obtained in this manner, where $\tau_0 = \sigma_0$, and τ_q is obtained by firing the q^{th} vertex in this sequence, in the configuration τ_{q-1}. Schematically,

$$\tau_0 \xrightarrow{v_1} \tau_1 \xrightarrow{v_2} \tau_2 \xrightarrow{v_3} \cdots \xrightarrow{v_{i-1}} \tau_{i-1} \xrightarrow{v_j} \tau_i \xrightarrow{v_i} \tau_{i+1} \xrightarrow{v_{i+1}} \tau_{i+2} \xrightarrow{v_{i+2}} \cdots \xrightarrow{v_{j-1}} \tau_j \xrightarrow{v_{j+1}} \tau_{j+1} \xrightarrow{v_{j+2}} \cdots \xrightarrow{v_{n-1}} \tau_{n-1} \xrightarrow{v_n} \tau_n$$

We must show that this is a valid firing sequence, and that $\tau_n = \sigma_n$. It suffices to show that $\tau_j = \sigma_j$.

First, it is evident that all firings up to v_j are valid, since $\tau_{i-1} = \sigma_{i-1}$ and both v_i, v_j are unstable in σ_{i-1}. Next, consider τ_{i+1}. This is the same as σ_i, except we have fired v_j early. Firing v_j early does not prevent any vertex from being unstable, except for v_j itself. However, $v_j \notin \{v_{i+1}, ..., v_{j-1}\}$ by construction. Thus, $\tau_j = \sigma_j$, and thus $\tau_n = \sigma_n$.

Thus, we have re-arranged the firing order for $\sigma_0 \to \sigma_n$ such that it agrees with $\eta_0 \to \eta_m$ up to the first i firings. Proceeding inductively, we can re-arrange so that the sequences agree up to the first $min(n, m)$ firings. □

Remark 1.2.13. The modification method in the proof is an analogue of the insertion sort algorithm for lists of integers. At each step we swap the order of two vertices in the firing sequence, so that the order is more closely aligned with the sequence $w_1, ..., w_m$.

Using the swapping principle, we can derive some useful properties, including the abelian property.

Corollary 1.2.14. Using the same notation as before,

i. $m \leq n$.

ii. If η_m is stable, then for all $v \in V$, v is fired the same number of times in the sequence $\eta_0 \to \eta_m$ as it is fired in the sequence $\sigma_0 \to \sigma_m$.

iii. The stabilization of any chip configuration $\sigma \in \mathbb{N}^V$ is unique, if it exists.

iv. If the stabilization of $\sigma \in \mathbb{N}^V$ exists, then it will be reached by any valid sequence of chip firings, in a unique number of steps. Moreover, each chip will be fired exactly the same number of times in each such sequence.

Proof. i. We know that we can rearrange firings so that $\sigma_k = \eta_k$, where $k = min(n, m)$. If $m > n$, then this would imply that η_n is stable. However, η_n can fire a chip, which is a contradiction. Thus $m \leq n$.

ii. Swapping the roles of η_m, σ_n, we find that $m = n$ by (i). Then by the rearrangement argument, we can rearrange firings so that $\eta_i = \sigma_i$ for all $0 \leq i \leq n$. So then the histogram of firings will be equal.

iii. By (ii), if σ_n, η_m are stable configurations both reached from σ, then the histogram of firings must be the same for both of them, which implies $\sigma_n = \eta_m$.

iv. By (ii) and (iii), we know that any sequence of firings which reaches the unique stabilization of σ will have length exactly n, and a unique histogram of firings. By definition, only these firing sequences are valid, and we can take any valid rearrangement to reach the stabilization. □

As a consequence of part (iii) of the corollary, we can unambiguously define the stabilization of any configuration.

Definition 1.2.15. Let $G = (V, E)$ be a graph, $\sigma \in \mathbb{Z}^V$. If there is no global sink, let $S(\sigma)$ be the **stabilization** of σ, if it exists. If there is some global sink vertex $z \in V$, let $S_z(\sigma)$ be the stabilization of σ with respect to z.

Notice that for the sink-vertex model, we have shown that for every chip configuration, there exists a unique stabilization, and this stabilization is reached by the same number of firings at each vertex. In the chapter 3, we will show that under a certain equivalence relation, these stable configurations form a finite abelian group called the *sandpile group*. Remarkably, the sandpile group has several equivalent characterizations, some of which can be defined with no reference to chip-firing whatsoever. But first, we need to develop tools from spectral graph theory and electrical network theory, which we do in the next chapter.

2 Spectral/Algebraic Graph Theory and Expanders

Spectral and algebraic graph theory give powerful tools for analyzing graph properties such as connectedness, chromatic number, cut sparsity, and more. In this chapter we will present some basic results from spectral (2.1) and algebraic (2.2) graph theory, anticipating and motivating later sections. Further, we will introduce expanders (2.3), which are special families of graphs that are simultaneously sparse and well-connected. Because of their remarkable asymptotic properties, expanders are a topic of intense interest for mathematicians and computer scientists. Therefore, we will later seek to understand the structure of the sandpile groups of expanders.

2.1 The Graph Laplacian

Throughout this section, let $G = (V, E)$ be an undirected graph, with self-loops and multi-edges allowed.

Definition 2.1.1. The **graph Laplacian** of G is an integer matrix $\Delta \in \mathbb{Z}^{V \times V}$, whose rows and columns are indexed by the vertices of G. For $u, v \in V$, let $e(u, v)$ denote the number of edges between u and v. Then entries of Δ are given by

$$\Delta_{u,v} = \begin{cases} deg(u) - e(u,u) & u = v \\ -e(u,v) & u \neq v \end{cases}$$

We can equivalently define Δ to be $D - A$, where D is the diagonal degrees matrix and A is the adjacency matrix.

Because G is undirected, Δ is a real symmetric matrix. Thus by the spectral theorem, all its eigenvalues are real and non-negative. The spectrum of the Δ tells us several useful things about the underlying graph. We present some basic facts below:

Proposition 2.1.2. Let Δ be the graph Laplacian of $G = (V, E)$. Let $0 \leq \lambda_1 \leq \lambda_2 \leq ... \leq \lambda_n$ be the eigenvalues of Δ in increasing order, counting multiplicity.

i. $\lambda_1 = 0$, and $\vec{1} \in Ker(\Delta)$, where $\vec{1} \in \mathbb{Z}^V$ is the all-ones vector.

ii. $\lambda_2 > 0$ if and only if G is connected.

Proof. i. Observe that for any $v \in V$, $(\Delta\vec{1})_v = \sum_{w \in V} \Delta_{v,w} = deg(v) - \sum_{w \in V} e(v, w) = 0$. Therefore $\vec{1} \in Ker(\Delta)$, and so the kernel of Δ has dimension at least one. Thus $\lambda_1 = 0$.

ii. $\Longleftarrow$: We show the contrapositive. Suppose that G is disconnected. Then G has at least 2 connected components. Let $U \subset V$ be a connected component. Let $\chi_U \in \mathbb{Z}^V$ be the indicator vector for U, so for $v \in V$,

$$\chi_U(v) = \begin{cases} 1 & v \in U \\ -1 & v \notin U \end{cases}$$

We claim $\Delta\chi_U = 0$. Fix $v \in V$. Then

$$(\Delta\chi_U)_v = \sum_{w \in V} \Delta_{v,w}\chi_U(w) \tag{1}$$

$$= \sum_{w \in U} \Delta_{v,w} \tag{2}$$

If $v \in U$, then $\sum_{w \in U} \Delta_{v,w} = deg(v) - \sum_{w \in U} e(v,w)$, since all edges incident to v are contained in U. If instead $v \notin U$, then $\Delta(v,w) = e(v,w) = 0$ for all $w \in U$, since U is a connected component. Thus in either case, $(\Delta \chi_U)_v = 0$. Since U is a proper subset of V, χ_U is not in the span of $\vec{1}$. Thus the kernel of Δ has dimension at least 2, and so $\lambda_2 = 0$.

$\Longrightarrow$: Suppose $\lambda_2 > 0$. Then for every nonempty proper subset $U \subset V$, $\chi_U \notin Ker(\Delta)$. Therefore by the above argument, no proper subset $U \subset V$ can be a connected component in G. Thus the only connected component of G contains all of the vertices, and thus G is connected. □

In fact, the second eigenvalue of the Laplacian captures a much more precise notion of graph connectivity, called the *edge expansion*. This relationship is given by Cheeger's inequality, which we develop here.

First, we define edge expansion, which measures how robustly a graph is connected. It does so by finding the sparsest deletion of edges which disconnects the graph.

Definition 2.1.3. Let $S \subset V$. The **edge boundary** of S, denoted ∂S, is the set of edges whose endpoints contain one vertex in S and one vertex in $V \setminus S$.

Informally, the size of the edge boundary ∂S captures how connected S is to the rest of a graph. If ∂S is small, then paths between vertices in S and $V \setminus S$ must pass through a bottleneck region, whose removal would disconnect S from the rest of the graph. If ∂S is large, then there are many paths out of S, and it will remain connected even if some of these boundary edges are lost. We can next define the edge expansion of a graph.

Definition 2.1.4. The **edge expansion** of G, also called the **Cheeger constant** and **isoperimetric number**, is the minimum ratio of $|\partial S|$ to $|S|$ for $S \subseteq V$.

$$h(G) = \min_{S \subset V : |S| \leq \frac{|V|}{2}} \frac{|\partial S|}{|S|}$$

Note that for any $S \subseteq V$, $|\partial S| = |\partial (V \setminus S)|$. Hence we could equivalently define the edge expansion as

$$h(G) = \min_{S \subseteq V} \frac{|\partial S|}{\min\{|S|, |V \setminus S|\}}$$

Evidently, $h(G) > 0$ if and only if G is connected. The best and worst cases of edge expansion are the complete and path graphs, respectively. If n is an even integer, then the path graph on n vertices has Cheeger constant $h(P_n) = \frac{2}{n}$, while the complete graph on n vertices, K_n, has Cheeger constant $h(K_n) = 1 + \frac{n}{2}$.

Having high edge expansion is an important measure of robustness, which is useful for many applications. Finding a subset $S \subset V$ which achieves $\frac{|\partial S|}{|S|} = h(G)$ is called the *Sparsest Cut problem*, and it is a well-studied problem in its own right. In fact, the Sparsest Cut problem is known to be NP-hard, with the best approximation algorithm giving an $O(\sqrt{log(n)})$ approximation for a graph on n vertices [ARV09].

Despite the difficulty of finding even an approximate solution to the Sparsest Cut problem, computing an approximation of the Cheeger constant is as easy as computing the second eigenvalue of the Laplacian matrix. This is due to Cheeger's inequality, which gives both lower and upper bounds on the Cheeger constant in terms of the second eigenvalue of the Laplacian.

To state Cheeger's inequality, we first need to define the *normalized* Laplacian matrix.

Definition 2.1.5. Let G be an unweighted graph in which every vertex has nonzero degree. [4] The **normalized Laplacian** of an unweighted graph G,with Laplacian Δ and degree matrix D is given by

$$L = D^{-1/2}\Delta D^{-1/2}$$

Notice that if G is d-regular, then the normalized Laplacian is just $\frac{1}{d}\Delta$. The normalized Laplacian is useful when G does not have constant degrees. We can now state and prove the "easy direction" of Cheeger's inequality.

Proposition 2.1.6. Let L denote the normalized Laplacian of G, and let its eigenvalues be $\nu_1 \leq \nu_2 \leq ... \leq \nu_n$, counting multiplicities. Then

$$h(G) \geq \frac{\nu_2}{2}$$

Proof. If G is disconnected, by 2.1.2 we know $\lambda_2 = h(G) = 0$.

Suppose G is connected. Then if $\langle\vec{1}\rangle$ denotes the linear span of $\vec{1} \in \mathbb{R}^V$, by 2.1.2 it follows that $Ker(\Delta) = \langle\vec{1}\rangle$. Then by the Courant-Fischer theorem, the minimum Rayleigh quotient of Δ over vectors $Ker(\Delta)^\perp$ gives its second eigenvalue. That is,

$$\lambda_2 = \min_{x \in \langle\vec{1}\rangle^\perp} \frac{x^T\Delta x}{x^T x}$$

We can exploit this characterization to lower-bound $h(G)$. Let $S \subset V$. Let $\chi_S \in \mathbb{Z}^V$ be the indicator vector for S, and $s = \frac{|S|}{|V|}$. Observe that $(\chi_S - s\vec{1})x \in \langle\vec{1}\rangle^\perp$, since $(\chi_S - s\vec{1})^T\vec{1} = |S| - s|V| = 0$. Thus, let $x = (\chi_S - s\vec{1})$.

Next, notice

$$x^T\Delta x = \sum_{(u,v)\in E} (x(u) - x(v))^2 \tag{3}$$

$$= \sum_{(u,v)\in E} (\chi_S(u) - \chi_S(v))^2 \tag{4}$$

$$= |\partial S| \tag{5}$$

Further, $x^T x = (1-2s)|S| + s^2|V| = (1-s)|S|$. Combining results, we obtain

$$\lambda_2 \leq \min_{S\subset V} \frac{|\partial S|}{(1-s)|S|} \leq \min_{S\subset V:|S|\leq|V|/2} \frac{2|\partial S|}{|S|} = 2h(G)$$

Finally, we relate λ_2 to ν_2. Notice that $D^{1/2}\vec{1} \in Ker(L)$. Since we assumed G is connected,

[4] This assumption is necessary for $D^{-1/2}$ to be defined.

$Ker(L)$ is generated by $D^{1/2}\vec{1}$. Thus by the Courant-Fischer Theorem,

$$\nu_2 = \min_{y:y\in\langle D^{1/2}\vec{1}\rangle^\perp} \frac{y^T L y}{y^T y} \tag{6}$$

$$= \min_{y:y\in\langle D^{1/2}\vec{1}\rangle^\perp} \frac{(D^{-1/2}y)^T \Delta (D^{-1/2}y)}{y^T y} \tag{7}$$

$$= \min_{y:y\in\langle D^{1/2}\vec{1}\rangle^\perp} \frac{y^T y}{z^T z}\frac{z^T \Delta z}{z^T z} \quad \text{Where } z \text{ denotes } (D^{-1/2}y) \tag{8}$$

$$\leq \min_{y:y\in\langle D^{1/2}\vec{1}\rangle^\perp} \frac{z^T \Delta z}{z^T z} \quad \text{since } D^{-1} \leq I \text{ entry-wise} \tag{9}$$

$$= \min_{z:D^{1/2}z\in\langle D^{1/2}\vec{1}\rangle^\perp} \frac{z^T \Delta z}{z^T z} \tag{10}$$

$$= \lambda_2 \tag{11}$$

Thus we conclude that $\frac{\nu_2}{2} \leq h(G)$. □

The other direction of Cheeger's inequality gives an upper bound on $h(G)$, as $h(G) \leq \sqrt{2\nu_2}$ (see, for example, [CC96], Theorem 3.1). Putting the two together, we obtain:

Theorem 2.1.7. *Let G be an undirected graph and ν_2 be the secondl-largest eigenvalue of its normalized Laplacian, counting multiplicies. Then*

$$\frac{\nu_2}{2} \leq h(G) \leq \sqrt{2\nu_2}$$

We conclude this section with the celebrated Matrix-Tree Theorem, which relates the number of spanning subtrees of a graph with the spectrum of its Laplacian matrix.

Definition 2.1.8. The **tree number** of a graph G on n labeled vertices, denoted $\kappa(G)$, is the number of distinct labeled spanning trees of G. [5]

Theorem 2.1.9. *([CK78], Theorem 1): Let G be an undirected graph on n labeled vertices. Let $0 \leq \lambda_1 \leq \lambda_2 \leq ... \leq \lambda_n$ denote the eigenvalues of the graph Laplacian, counting multiplicities. Then*

$$\kappa(G) = \frac{1}{n}\lambda_2\lambda_3\cdots\lambda_n$$

2.2 Cuts and Flows

Throughout this section, let $G = (V, E)$ denote a finite, connected, undirected graph. We assume that G has no self-loops, although multi-edges are still allowed.[6]. Further, assign an arbitrary orientation to the edges of G, such that each edge $e \in E$ has a head (denoted $h(e)$) and a tail (denoted $t(e)$).

[5] The distinction between spanning trees on labeled and unlabeled vertices is important. On labeled vertices, every unique subset of edges which forms a spanning tree will be counted separately. However, on unlabeled vertices these spanning trees will only be counted up to isomorphism (i.e. up to re-labeling of the vertices). For example, the triangle graph (the cycle graph on 3 vertices) has 3 labeled spanning trees but only one unlabeled spanning tree.

[6] The presence of self-loops would make the incidence matrix ill-defined, since a self-loop at v would mean $h(\{v, v\}) = t(\{v, v\}) = v$

Definition 2.2.1. The **incidence matrix** $D \in \mathbb{R}^{V \times E}$ on G is defined, for $v \in V, e \in E$, as

$$D_{v,e} = \begin{cases} 1 & v = h(e) \\ -1 & v = t(e) \\ 0 & \text{Otherwise} \end{cases}$$

Having assigned an orientation to edges of G, we can identify vectors in $\mathbb{R}^E$ with cuts and flows of G.

Definition 2.2.2. Let G be an oriented, connected graph and $D \in \mathbb{R}^{V \times E}$ its incidence matrix. The **cut space** of G is

$$B = Ker(D)^{\perp}$$

The **flow space** of G is

$$Z = Ker(D)$$

Thus

$$\mathbb{R}^E = Ker(D)^{\perp} \oplus Ker(D) = \text{Cut Space} \oplus \text{Flow Space}$$

As one would expect, the vectors in the cut space can be identified with cuts in the graph, and the same is true of flow space.

Definition 2.2.3. Given a nonempty, proper vertex subset $U \subset V$, we can define the **characteristic vector of** U $b_U \in \mathbb{R}^E$ by

$$b_U(e) = \begin{cases} 1 & U \cap e = \{h(e)\} \\ -1 & U \cap e = \{t(e)\} \\ 0 & \text{otherwise} \end{cases}$$

Observe that if $b_U(e) \neq 0$, then exactly one vertex of e is contained in U. Thus, e is an edge from U to $V \setminus U$. Thus the edges at which b_U is nonzero give a cut of the graph. Moreover, b_U is indeed an element of the cut space.

Proposition 2.2.4. For any nonempty, proper $U \subset V$, $b_U \in Ker(D)^{\perp}$.

Proof. Notice that $b_U = \sum_{v \in U} b_{\{v\}} = \sum_{v \in U} D^T \delta_v$. Let $z \in Ker(D)$. Then

$$\begin{aligned} \langle z, b_U \rangle &= \sum_{v \in U} \langle z, D^T \delta_v \rangle && (12) \\ &= \sum_{v \in U} z^T D^T \delta_v && (13) \\ &= \sum_{v \in U} (Dz)^T \delta_v && (14) \\ &= 0 && (15) \end{aligned}$$

□

Next, we define the analogous notion for the flow space.

Definition 2.2.5. Let $Q = (v_1, e_1, v_2, e_2, ..., v_{r-1}, e_{r_1}, v_r)$ be a cycle. Then the **characteristic vector of** Q $z_Q \in \mathbb{R}^E$ is given by

$$z_Q(e) = \begin{cases} 1 & t(e), e, h(e) \text{in } Q \\ -1 & h(e), e, th(e) \text{in } Q \\ 0 & \text{otherwise} \end{cases}$$

Proposition 2.2.6. For any cycle Q, $z_Q \in Ker(D)$.

Proof. It is enough to show that for $v \in V$, $\langle z_Q, b_{\{v\}} \rangle = 0$, since $b_{\{v\}} = D^T \delta_v$, so $0 = \langle z_Q, b_{\{v\}} \rangle = z_Q^T D^T \delta_v = (Dz_Q)^T \delta_v$.

Consider $v \in V$. If v never occurs in the cycle, then clearly $\langle z_Q, b_{\{v\}} \rangle = 0$. If v does occur in the cycle, there are four cases.

$$z_Q^T b_{\{v\}} = (z_Q(e_1) b_{\{v\}}(e_1)) + (z_Q(e_2) b_{\{v\}}(e_2)) = \begin{cases} (1 \cdot 1) + (-1 \cdot 1) & v = h(e_1) = t(e_2) \\ (-1 \cdot -1) + (1 \cdot -1) & v = t(e_1) = h(e_2) \\ (-1 \cdot 1) + (-1 \cdot -1) & v = h(e_1) = h(e_2) \\ (1 \cdot -1) + (1 \cdot 1) & v = t(e_1) = t(e_2) \end{cases}$$

In all four cases, the inner product is 0. Thus $\langle z_Q, b_{\{v\}} \rangle = 0$, so we conclude $z_Q \in Ker(D)$. □

Finally, we conclude with a simple identity relating the incidence matrix to the graph Laplacian.

Proposition 2.2.7. The graph Laplacian Δ is given by

$$\Delta = DD^T$$

Proof. Let $v, w \in V$ be distinct. Then

$$\begin{aligned} (DD^T)_{v,w} &= \sum_{e \in E} D_{v,e} D^T_{e,w} && (16) \\ &= \sum_{e \in E} D_{v,e} D_{w,e} && (17) \\ &= -|\{e \in E : e = \{v, w\}\}| && (18) \\ &= \Delta_{v,w} && (19) \end{aligned}$$

Next,

$$\begin{aligned} (DD^T)_{v,v} &= \sum_{e \in E} D_{v,e} D^T_{e,v} && (20) \\ &= \sum_{e \in E} D^2_{v,e} && (21) \\ &= d_v = \Delta_{v,v} && (22) \end{aligned}$$

We conclude that all entries of Δ and DD^T are equal. □

Remark 2.2.8. One can define a Laplacian on a *weighted* graph by replacing edge counts with weights, so that $\Delta_{v,w}$ is the sum of the weights on the edges between $v, w \in V$, and $\Delta_{v,v}$ is the sum of the weights on edges incident to v, minus the sum of weights on self-loops at v. It turns out that if $W \in \mathbb{R}^{E \times E}$ is the diagonal weight matrix, then one can generalize this result to

$$\Delta = DWD^T$$

Our version corresponds to the case where every edge has weight 1.

2.3 Expander Graphs

Expanders are, informally, graphs which are simultaneously sparse and robustly connected. They are a major topic of interest to computer scientists and mathematicians, with applications to error-correcting codes, metric geometry, probabilistically checkable proofs, and so on [HLW06].

A natural way of viewing expanders is as optimal asymptotic solutions to the *isoperimetry problem*. Recall from 2.1 that any graph has an associated *Cheeger constant* which is a measure of how robustly connected it is. Graphs with high Cheeger constant will remain connected even when subject to small edge perturbations or deletions.

Clearly, a complete graph is the most robust graph in this sense. However, it is also important to consider the overhead of storing complete graphs, which have $O(n^2)$ edges on n vertices. This becomes cumbersome for applications involving large n, and we would thus like to know if we can retain the high Cheeger constant of complete graphs, but for a much sparser graph. A priori, it is not obvious that we can. For example, the path graph P_n is very sparse, having only $O(n)$ edges, but it also has very low edge expansion.

Miraculously, there are graphs which have high Cheeger constant but are still quite sparse, having only $O(n)$ edges. These are called expander graphs, or expander families.

Throughout this section, $G = (V, E)$ will denote a finite, undirected, connected graph.

Definition 2.3.1. Let $d \in \mathbb{N}$ and $\epsilon > 0$. A family of graphs $(G_i)_{i \in \mathbb{N}}$ is a (d, ϵ)**-expander family** if:

i. For all i, G_i is d-regular (that is, each vertex has exactly d incident edges).

ii. For all i, $h(G_i) \geq \epsilon$.

iii. The sequence $(|V(G_i)|)_{i \in \mathbb{N}}$ is non-decreasing and goes to infinity.

Since an expander family is d-regular for a constant d, every member graph G_i has $O(n)$ edges and expansion bounded away from zero. Notice that these are the asymptotically sparsest connected graphs, since a connected graph on n vertices has at least $n - 1$ edges.

A number of expander family constructions, both deterministic and probabilitistic, are known. We discuss three such families.

Definition 2.3.2. The **Margulis, Gabber, and Galil (MGG) graphs** are a family of expander graphs $(G_k)_{k \in \mathbb{N}^+}$. For each $k \in \mathbb{N}^+$, $V_k = (\mathbb{Z}/k\mathbb{Z}) \times (\mathbb{Z}/k\mathbb{Z})$. For each $(x, y) \in V_k$, its neighbors are $(x+y, y), (x-y, y), (x, y+x), (x, y-x), (x+y+1, y), (x-y+1, y), (x, x+y+1)$, and $(x, y-x+1)$ (all additions and subtractions are modulo k).

The family $(G_k)_{k \in \mathbb{N}^+}$ is an 8-regular expander family ([HLW06], 2.2(1)).

Definition 2.3.3. Let $P \subset \mathbb{N}$ be all positive prime integers. The **chordal cycle graphs** are a family of expander graphs $(G_p)_{p \in P}$. For each prime p, $V_p = (\mathbb{Z}/p\mathbb{Z})$. For each $x \in V_p$, its neighbors are $(x + 1), (x - 1)$, and x^{-1} (all operations are modulo p, and we define 0^{-1} to be 0).

The family $(G_p)_{p \in P}$ is a 3-regular expander family ([HLW06], 2.2(2)).

The proofs that the MGG graphs and chordal cycle graphs are expander families rely on highly technical results in harmonic analysis and number theory, respectively [HLW06]. However, for random graphs the analysis is much simpler, requiring only elementary ideas from probability. Below, we define a family of d-regular bipartite expanders and prove their expansion is bounded away from zero with constant probability.

Definition 2.3.4. Let $d \in \mathbb{N}^+$. Let S_n be the set of permutations of $[n]$, and let $\pi_1, ..., \pi_d \in S_n$ be chosen independently and uniformly at random. Let $G = (V, E)$ be defined as follows.

$$V = \{\ell_1, ..., \ell_n, r_1, ..., r_n\} \tag{23}$$

$$E_k = \{(\ell_1, r_{\pi_m(1)}), (\ell_2, r_{\pi_k(2)}), ..., (\ell_n, r_{\pi_k(n)})\} \tag{24}$$

$$E = \bigsqcup_{k=1}^{d} E_k \tag{25}$$

Note that E is the *disjoint* union of of $E_1, ..., E_d$, so there are possibly multiple edges between a pair of vertices ℓ_i, r_j.

We can show that with nonzero probability, a graph sampled in such a way is an expander for large enough n, d.

Proposition 2.3.5. Let $d > 2$, and $n > 8d$. Then if G is a d-regular random graph on $2n$ vertices described as above, then there exist constants $p, c > 0$ independent of n such that $\mathbb{P}[h(G) \geq c] > p$. Thus G is an expander with positive probability.

Proof. By definition G is d-regular. We want to show that for some $c > 0$ that $\mathbb{P}[\forall S \subseteq V : |S| \leq n \Rightarrow (\frac{|\partial S|}{|S|} \geq c)] > 0$. Notice that

$$\mathbb{P}[\forall S \subset V : |S| \leq n \Rightarrow \frac{|\partial S|}{|S|} \geq c] = 1 - \mathbb{P}[\exists S \subset V : |S| \leq n \wedge \frac{|\partial S|}{|S|} < c] \tag{26}$$

$$\geq 1 - \sum_{S \subset V : |S| \leq n} \mathbb{P}[\frac{|\partial S|}{|S|} < c] \quad \text{Union Bound} \tag{27}$$

Thus, we want to show that $\sum_{S \subset V : |S| \leq n} \mathbb{P}[\frac{|\partial S|}{|S|} < c] < 1$.

There are two cases regarding S.

Case 1: $S \subseteq V_1$ or $S \subseteq V_2$. Then every edge incident to S leaves S to go to the other side of the partition, and so $|E(S, V \setminus S)| = d\,|S|$. Thus $\frac{|E(S, V \setminus S)|}{|S|} = d$.

Case 2: $S \not\subseteq V_1$ and $S \not\subseteq V_2$. Without loss of generality let $|S \cap V_1| \geq |S \cap V_2|$ (the other case is symmetric).

For simplicity of notation let $S_1 = S \cap V_1, S_2 = S \cap V_2$. Then by assumption $|S_2| \leq \frac{n}{2}$. Thus, let B_2 be any $(n/2)$ elements of $V_2 \setminus S_2$. Notice that every edge in $E(S_1, B_2)$ is in ∂S. Thus $\frac{|E(S_1, B_2)|}{2|S_1|} \leq \frac{|\partial S|}{2|S_1|} \leq \frac{|\partial S|}{|S|}$.

Thus $\frac{|\partial S|}{|S|} < c$ implies that $\frac{|E(S_1, B_2)|}{2|S_1|} < c$, and thus $\mathbb{P}[\frac{|\partial S|}{|S|} < c] \leq \mathbb{P}[\frac{|E(S_1, B_2)|}{|S_1|} < 2c]$.

We use a Chernoff bound.

For $i \in [d]$, let $E_i(S_1, B_2)$ denote the edges between S_1, B_2 that are from the i^{th} permutation. Let $X_1, ..., X_d$ be Bernoulli random variables, where for all $i \in [d]$,

$$X_i = \begin{cases} 1 & \frac{|E_i(S_1,B_2)|}{|S_1|} \geq \frac{1}{2} \\ 0 & \text{otherwise} \end{cases}$$

Notice that $E_i(S_1, B_2) \cup E_i(S_1, V_2 \setminus B_2) = E_i(S_1, V_2)$. Further, by symmetry, since $\mathbb{P}[\frac{|E_i(S_1,B_2)|}{|S_1|} \geq \frac{1}{2}] = \mathbb{P}[\frac{|E_i(S_1,V_2\setminus B_2)|}{|S_1|} \geq \frac{1}{2}]$. Thus since at least one of the ratios must be at least $1/2$, it follows that

$$1 = \mathbb{P}[(\frac{|E_i(S_1,B_2)|}{|S_1|} \geq \frac{1}{2}) \cup (\frac{|E_i(S_1,V_2 \setminus B_2)|}{|S_1|} \geq \frac{1}{2})] \quad (28)$$

$$= \mathbb{P}[\frac{|E_i(S_1,B_2)|}{|S_1|} \geq \frac{1}{2}] + \mathbb{P}[\frac{|E_i(S_1,V_2 \setminus B_2)|}{|S_1|} \geq \frac{1}{2}] - \mathbb{P}[\frac{|E_i(S_1,V_2 \setminus B_2)|}{|S_1|} = \frac{|E_i(S_1,V_2 \setminus B_2)|}{|S_1|} = \frac{1}{2}] \quad (29)$$

$$= 2\mathbb{P}[\frac{|E_i(S_1,B_2)|}{|S_1|} \geq \frac{1}{2}] - \mathbb{P}[\frac{|E_i(S_1,V_2 \setminus B_2)|}{|S_1|} = \frac{|E_i(S_1,V_2 \setminus B_2)|}{|S_1|} = \frac{1}{2}] \quad (30)$$

Thus since $\mathbb{P}[\frac{|E_i(S_1,V_2\setminus B_2)|}{|S_1|} = \frac{|E_i(S_1,V_2\setminus B_2)|}{|S_1|} = \frac{1}{2}] \leq \mathbb{P}[\frac{|E_i(S_1,B_2)|}{|S_1|} \geq \frac{1}{2}]$, it follows that $1/3 \leq \mathbb{P}[\frac{|E_i(S_1,B_2)|}{|S_1|} \geq \frac{1}{2}] \leq 2/3$.

We are ready to apply the Chernoff bound. $X_1, ..., X_d$ are i.i.d. Bernoulli variables, and for all i we have $1/3 \leq \mathbb{P}[X_i = 1] \leq 2/3$. Let $X = \sum_{i=1}^{d} X_i$ and $\mu = \mathbb{E}[X] = d\mathbb{P}[X_1 = 1]$. Notice $X = \frac{|E(S_1,B_2)|}{|S_1|}$, and $d/3 \leq \mu \leq 2d/3$.

Let $\delta \in (0, 1)$ be a constant (to be chosen later). Then setting $2c = (1 - \delta)\mu$, by the Chernoff bound we have that

$$\mathbb{P}[\frac{|E(S_1,B_2)|}{|S_1|} < 2c] = \mathbb{P}[X < (1-\delta)\mu] \leq exp(\frac{-\delta^2\mu}{2})$$

Thus,

$$\sum_{S\subset V:|S|\leq n} \mathbb{P}[\frac{|\partial S|}{|S|} < c] \leq \sum_{S\subset V:|S|\leq n} \mathbb{P}[\frac{|E(S_1,B_2)|}{|S_1|} < 2c] \quad (31)$$

$$\leq \sum_{S\subset V:|S|\leq n} exp(\frac{-\delta^2\mu}{2}) \quad (32)$$

$$\leq 2^{2n} exp(\frac{-\delta^2\mu}{2}) \quad (33)$$

$$\leq exp(2n - \frac{\delta^2\mu}{2}) \quad (34)$$

Thus, to obtain $exp(2n - \frac{\delta^2\mu}{2}) \leq e^{-1}$, we set $\delta \geq \sqrt{\frac{4n+2}{\mu}}$, and hence $c \leq \frac{\mu}{2} - \sqrt{\mu(4n+2)} \leq \frac{\mu}{2} \leq \frac{d}{3}$ suffices. Since in case 1 we found that $c = d$ is achieved with certainty, we conclude that

$$\mathbb{P}[\forall S \subset V : |S| \leq n \Rightarrow \frac{|\partial S|}{|S|} \geq \frac{d}{3}] \quad (35)$$

$$\geq 1 - \sum_{S\subset V:|S|\leq n} \mathbb{P}[\frac{|\partial S|}{|S|} < \frac{d}{3}] \quad (36)$$

$$\geq 1 - \frac{1}{e} \quad (37)$$

Thus setting $c = \frac{d}{3}$ and $p = 1 - \frac{1}{e}$, we obtain constants independent of n. $\square$

Notice that at the cost of a weaker lower bound c, we can increase the probability that the Cheeger constant is achieved. Thus for any $\epsilon > 0$, there exists a $c_\epsilon > 0$ such that $\mathbb{P}[h(G) \geq c_\epsilon] \geq 1-\epsilon$. In particular, we have shown that for every admissible choice of (n, d) there exists a d-regular bipartite graph on $2n$ vertices with expansion at least c_ϵ. This gives an expander family.

3 The Sandpile Group

In this chapter we turn our attention back to sandpiles, and present four equivalent characterizations of the sandpile group.

The sandpile group of a graph is a finite abelian group that captures several remarkable properties, one of which is the collection of recurrent chip configurations on the graph under the global sink model of the abelian sandpile (3.2). In addition to this characterization in terms of sandpiles, it is also isomorphic to the cokernel of the Laplacian matrix (3.1), and to a certain quotient lattice defined in terms of integral cuts and flows (3.4).

As mentioned in 1.1, the sandpile group goes by many names in the literature, as its ubiquity makes it interesting to researchers in subfields of mathematics, computer science, statistical physics, etc. In this chapter we will explore just a few of these connections, motivating the focus of chapter 4.

3.1 The cokernel of the Laplacian Matrix

Perhaps the most straightforward definition of the sandpile group is in terms of the cokernel of the Laplacian matrix. Throughout this section, let $G = (V, E)$ be an undirected, connected graph and Δ its Laplacian.

Any linear map is a special case of a group homomorphism, where the groups in question are just the additive groups of the domain and codomain respectively. Just as the kernel of a group homomorphism is a subgroup of its domain, the *cokernel* is a subgroup of its codomain.

Definition 3.1.1. Let G, H be groups and $\phi : G \to H$ be a group homomorphism. The **cokernel of** ϕ is the quotient group $H/\phi(G)$.

$$coker(\phi) = H/\phi(G)$$

We are interested in the graph Laplacian matrix Δ, which is a group homomorphism on the additive group of $\mathbb{Z}^V$. First, note that the image of $\mathbb{Z}^V$ under Δ is orthogonal to the all-ones vector.

Proposition 3.1.2. Let $\mathbb{Z}_0^V = \{x \in \mathbb{Z}^V : \vec{1}^T z = 0\}$. Then $\Delta(\mathbb{Z}^V) \subseteq \mathbb{Z}_0^V$.

Proof. We must show that for all $x \in \mathbb{Z}^V$, (Δx) is in the orthogonal complement of the all-ones vector. Let $z \in \mathbb{Z}^V$. Observe:

$$\vec{1}^T(\Delta x) = (\Delta x)^T \vec{1} = x^T \Delta^T \vec{1} = x^T \Delta \vec{1} = x^T \vec{0} = 0$$

□

Therefore, if we are interested in integer vectors, we can just as well *define* Δ to be a map from $\mathbb{Z}^V$ to $\mathbb{Z}_0^V$. Nothing is lost in this definition, since the image $\Delta(\mathbb{Z}^V)$ is contained in $\mathbb{Z}_0^V$.

The sandpile group is then just the cokernel of the Laplacian matrix.

Definition 3.1.3. Let Δ denote the graph Laplacian of a connected, undirected graph G. Let $\mathbb{Z}_0^V = \{x \in \mathbb{Z}^V : \vec{1}^T z = 0\}$. Then the **sandpile group of** G is given by

$$S(G) = \mathbb{Z}_0^V / \Delta(\mathbb{Z}^V)$$

In other words, if we view $\Delta : \mathbb{Z}^V \to \mathbb{Z}_0^V$ as a group homomorphism, then

$$S(G) = coker(\Delta)$$

Why is this group interesting? One interpretation of the sandpile group is as the set of all chip configurations on a graph with fixed sink, modulo equivalence under chip-firing. To precisely explain this, we first show how chip-firings are just addition and subtraction with columns of the Laplacian.

Proposition 3.1.4. Let $\sigma \in \mathbb{Z}^V$ be a chip configuration, and $\eta \in \mathbb{Z}^V$ be the configuration obtain by firing some $v \in V$ from σ. Then if $\delta_v \in \mathbb{Z}^V$ is the indicator vector at v,

$$\eta = \sigma - \Delta\delta_v$$

Proof. For $w \in V$, let $a_{v,w} \in \mathbb{N}$ denote the number of edges between v and w. Then if d_v denotes the degree of v, $\eta_v = \sigma_v + a_{v,v} - d_v = (\sigma - \Delta\delta_v)_v$, and for $w \neq v$, $\eta_w = \sigma_w + a_{v,w} = (\sigma - \Delta\delta_v)_w$. □

Corollary 3.1.5. Let $\sigma \in \mathbb{Z}^V$ be a chip configuration, and $\eta \in \mathbb{Z}^V$ be a configuration obtained by some sequence of chip firings. Suppose $u \in \mathbb{N}^V$ is the vector such that u_v counts the number of times that v was fired. Then

$$\sigma = \eta - \Delta u$$

Proof. Notice $u = \sum_{v \in V} u_v \delta_v$. Induction on $|u|$ gives the result. □

It follows that σ, η belong to the same coset of $\Delta(\mathbb{Z}^V)$ exactly when one is reachable from the other via chip-firings. Thus, cosets can be considered "firing equivalence classes" in the sense that two elements in the same coset are related by a sequence of firings.[7]

Further, we can identify $\mathbb{Z}_0^V$ as the set of all "interesting" chip configurations when there is a global sink. Suppose that G has a fixed global sink vertex $z \in V$. For the purposes of understanding chip-firing dynamics, the number of chips at z is irrelevant, as z can never fire and is always considered stable. Therefore, if $\sigma, \eta \in \mathbb{Z}^V$ are chip configurations which are equal on every non-sink vertex, then they can both be identified with the unique $\zeta \in \mathbb{Z}_0^V$ such that $\zeta_z = -\sum_{v \in V \setminus \{v\}} \zeta_v$. Conversely, any $\beta \in \mathbb{Z}_0^V$ is such that $\beta_z = -\sum_{v \in V \setminus \{v\}} \beta_v$.

Therefore, when interested in all sandpiles in $\mathbb{Z}^V$, the presence of a global sink means it suffices to consider sandpiles in $\mathbb{Z}_0^V$. Moreover, cosets of $\Delta(\mathbb{Z}^V)$ in $\mathbb{Z}_0^V$ can be viewed as collections of sandiples up to firing-equivalence. Thus the sandpile group is just the collection of all equivalence classes of chip configurations.

3.2 Combinatorial interpretation

In this section we will develop a purely combinatorial perspective on the sandpile group, by viewing it as a collection of chip configurations with a suitable binary operation. This justifies the name of *sandpile* group, since under this perspective the sandpile group is indeed a collection of sandpiles.

Throughout this section, let $G = (V, E)$ be an undirected, connected graph with global sink vertex $z \in V$. Recall that the requirement of a sink vertex ensures that every chip configuration on G has a unique stabilization (see 1.2.15). For a chip configuration σ, we will denote its stabilization with respect to z by $S_z(\sigma)$.

First, we define recurrent configurations.

[7] And of course, the cosets of $\Delta(\mathbb{Z}^V)$ give a partition of $\mathbb{Z}_0^V$, so membership in the same coset is a bona fide equivalence relation. However, our relation of "firing equivalence" is not an equivalence relation, since it is not symmetric; consider the case where σ is reachable from η via a sequence of firings, but σ is stable and so cannot fire at all.

Definition 3.2.1. Let σ, ζ be chip configurations. We say that σ is **reachable** from ζ if there exists some configuration $\alpha \geq 0$ (component-wise) such that $S_z(\zeta + \alpha) = \sigma$

Definition 3.2.2. A chip configuration $\sigma \in \mathbb{Z}^V$ is **accessible** if it is reachable from every chip configuration $\zeta \in \mathbb{Z}^V$.

Note that any configuration which has a negative number of chips somewhere cannot be accessible, since the all-zero configuration is stable and can never reach a negative configuration via chip additions and firings.

We can now define recurrent configurations.

Definition 3.2.3. A chip configuration $\sigma \in \mathbb{Z}^V$ is **recurrent** if it is stable and accessible.

Our goal in this section is to show that the recurrent configurations are bijective with the elements of the sandpile group, which we defined as the cokernel of the graph Laplacian. In fact, under a suitable binary operation, the recurrent configurations form a group which is isomorphic with the sandpile group.

We begin by developing a few key lemmas concerning recurrent configuratinos.

Lemma 3.2.4. Let $\sigma \in \mathbb{Z}^V$ be a chip configuration such that $\sigma_v \geq 0$ for all $v \neq z$. Then $S_z(\sigma + \eta) = S_z(\sigma + S_z(\eta))$.

Proof. Let $\sigma' = \sigma + \eta$. Since σ is nonnegative on non-sink vertices, every vertex unstable in η is also unstable in σ'. Thus, starting from σ', it is legal to apply the (possibly empty) sequence of firings which stabilizes η. This yields $\sigma + S_z(\eta)$. Stabilizing $\sigma + S_z(\eta)$, via another sequence of firings, gives $S_z(\sigma + S_z(\eta))$.

Composing these two firing sequences, we obtain a legal firing sequence from σ' which results in a stable configuration. By 1.2.14, σ' has a unique stabilization. Thus $S_z(\sigma') = S_z(\sigma + \eta) = S_z(\sigma + S_z(\eta))$. □

Corollary 3.2.5. If for some configurations σ, η (where η is non-negative on non-sink vertices) we have $S_z(\sigma + \eta) = \sigma$, then for all $k \in \mathbb{Z}^+$, $S_z(\sigma + k\eta) = \sigma$.

Proof. Since η is non-negative on non-sink vertices, observe by 3.2.4 that

$$S_z(\sigma + k\eta) = S_z(S_z(\sigma + \eta) + (k-1)\eta) \tag{38}$$
$$= S_z(\sigma + (k-1)\eta) \tag{39}$$

It follows by induction on k that $S_z(\sigma + k\eta) = \sigma$. □

We begin by defining a useful configuration ϵ.

Lemma 3.2.6. Let β be the configuration obtained by "firing the sink," so $\beta = -\Delta\delta_z$.

There exists $k > 0$ such that for $\epsilon = (k\beta) - S_z(k\beta)$,

i. ϵ can be selectively fired to obtain a configuration $\alpha \in \mathbb{Z}^V$, such that $\alpha \geq 1$ component-wise. In other words, α has at least one chip at every vertex. then for some $v \neq z$

ii. For every recurrent configuration σ, $S_z(\sigma + \epsilon) = S_z(\sigma)$.

Proof. i. Since G is connected, the sink vertex $z \in V$ has at least one neighbor $v \in V$. Thus, $\beta_v > 0$. Then for large enough m, $m\beta_v > deg(v)$, so $m\beta$ is unstable. Thus if $\epsilon_m = (m\beta) - S_z(m\beta)$, then $(\epsilon_m)(v) > 0$.

Since G is connected, v has a path to every other vertex. Thus for large enough $\ell \in \mathbb{N}^+$, the configuration $(\ell\delta_v)$ can be selectively fired to obtain some $\alpha' \geq 1$. Then since $(\ell\ell_m) \geq (\ell\delta_v)$ component-wise, we can perform the same sequence of selective firings from $(\ell\ell_m)$, to obtain $\alpha = (\ell\ell_m - \ell\delta_v) + \alpha'$. Since $(\ell\ell_m - \ell\delta_v) \geq 0$, it follows that $\alpha \geq \alpha' \geq 1$, so setting $\epsilon = \ell\epsilon_m$ suffices.

ii. By 3.2.5, it suffices to show that $S_z(\sigma + \epsilon_m) = \sigma$. Since σ is accessible, there exists some configuration $\eta \geq 0$ such that $S_z(\beta + \eta) = \sigma$. Then consider $\gamma = \beta + \eta + \epsilon_m = \beta + \eta + k\beta - S_z(k\beta)$.

Observer that we can selectively fire γ to stabilize $k\beta$, obtaining $\beta + \eta + S_z(k\beta) - S_z(k\beta) = \beta + \eta$. Stabilizing this configuration gives $S_z(\beta + \eta) = \sigma$. Since stabilizations are unique, it follows that $S_z(\gamma) = \sigma$.

Since $k\beta - S_z(k\beta)$ is non-negative, we can instead selectively fire γ to stabilize $\beta + \eta$, obtaining $\sigma + \epsilon_m$. Stabilizing this configuration should give $S_z(\gamma) = \sigma$. Thus $\sigma = S_z(\sigma + \epsilon_m)$. □

These two special properties of ϵ allow us to prove that every coset of $\Delta(\mathbb{Z}^V)$ contains exactly one recurrent chip configuration.

First, we show the existence of a recurrent configuration in each coset.

Proposition 3.2.7. ([Hol08] Lemma 2.13): Each coset in $\mathbb{Z}_0^V / \Delta(\mathbb{Z}^V)$ contains at least one recurrent configuration.

Proof. Let $\sigma \in \mathbb{Z}_0^V$ be an arbitrary configuration. Then by the lifting lemma, there exists $\ell > 0$ such that $(\sigma + \ell\beta)$ can be selectively fired to obtain a configuration ζ, where $\zeta \geq \max(\{d_v : v \in V\})$.

We claim that $S_z(\zeta)$ is recurrent. It is stable by definition. For accessibility, let $\eta \in \mathbb{Z}^V$ be arbitrary. Observe that $S_z(\eta) < \zeta$ component-wise, since each vertex in $S_z(\eta)$ has fewer chips than its degree. Thus $S_z(\zeta) = S_z((\zeta - S_z(\eta)) + S_z(\eta)) = S_z(\eta + (\zeta - S_z(\eta)))$. Since $(\zeta - S_z(\eta))$ is non-negative, it follows that we can add non-negative chips to η and then stabilize to obtain $S_z(\zeta)$.

Next, observe that $S_z(\zeta) = S_z(\sigma + \ell\beta)$, since we can first fire chips from $(\sigma + \ell\beta)$ to obtain ζ, and then stabilize.

Finally, $S_z(\zeta)$ is in the same coset as σ. Since $\sigma + \ell\beta = \sigma - \ell\Delta\delta_z$, it follows that $\sigma + \ell\beta$ is in the same coset as σ. Since $S_z(\zeta) = S_z(\sigma + \ell\beta)$ is obtained by firing chips from $(\sigma + \ell\beta)$, it follows that $S_z(\zeta)$ is in the same coset as $(\sigma + \ell\beta)$, and thus the same coset as σ. Since $S_z(\zeta)$ is recurrent, we are done. □

Second, we show the uniqueness of this recurrent configuration.

Proposition 3.2.8. ([Hol08] Lemma 2.15): Let $\sigma, \zeta \in^V$ be recurrent configurations such that for some $w \in \mathbb{Z}^V$, $\sigma - \Delta w = \zeta$. Then $\sigma = \zeta$.

Proof. Let $w^+ \in \mathbb{Z}^V$ be the positive component of w, so $(w^+)_v = \begin{cases} w_v & w_v > 0 \\ 0 & w_v \leq 0 \end{cases}$. Similarly, let $w^- \in \mathbb{Z}^V$ be the negative component, so $(w^-)_v = \begin{cases} w_v & w_v < 0 \\ 0 & w_v \geq 0 \end{cases}$. Notice that $\Delta w = \Delta(w^+ + w^-)$. Therefore, we have

$$\sigma - \Delta w^- = \zeta + \Delta w^+$$

Let $\tau = \sigma - \Delta w^- = \zeta + \Delta w^+$.

Next, let ϵ be as before. We know that ϵ can be selectively fired to obtain some configuration $\alpha' \geq 1$. Thus, for any $m \in \mathbb{N}$, $(m\epsilon)$ be can selectively fired to obtain some $\alpha = m\alpha' \geq m$. Thus, let α be a configuration reachable from $\ell\beta$ for some sufficiently large $\ell \in \mathbb{N}$, such that for all $v \in V \setminus \{z\}$, $(\tau + \alpha)_v \geq |w_v|\, d_v$.

Then, consider $\tau + \ell\epsilon$. By the lemma, we know for any recurrent configuration ζ, $S_z(\eta + \alpha) = S_z(\eta + \ell\epsilon) = \eta$.

We can first selectively fire chips in $(\tau + \ell\epsilon)$ to obtain $\tau + \alpha$. Then, we can fire each vertex v such that $w_v > 0$ exactly w_v times to obtain $(\tau + \alpha - \Delta(w^+)) = \zeta + \alpha$. Since ζ is stable, we obtain $S_z(\tau + \alpha) = S_z(\zeta + \alpha) = \zeta$.

Next, we can instead start from $\tau + \alpha$ and fire each vertex v such that $w_v < 0$ exactly $(-w_v)$ times to obtain $(\tau + \alpha + \Delta w^-) = \sigma + \alpha$. Then since σ is stable, we obtain $S_z(\tau + \alpha) = S_z(\sigma + \alpha) = \sigma$.

Schematically, these two processes are:

$$\tau + \alpha \xrightarrow{Firings} \tau + \alpha - \Delta w^+ = (\zeta + \Delta w^+) - \Delta w^+ + \alpha = \zeta + \alpha \xrightarrow{S_z} \zeta$$

$$\tau + \alpha \xrightarrow{Firings} \tau + \alpha + \Delta w^- = (\sigma - \Delta w^-) + \Delta w^- + \alpha = \sigma + \alpha \xrightarrow{S_z} \sigma$$

Thus we conclude that $\zeta = S_z(\tau + \alpha) = \sigma$. □

Combining the two results, we obtain a bijection between the sandpile group and the collection of recurrent chip configurations. We are almost ready to prove the isomorphism, but we need to define a group operation on the collection of recurrent configurations. This binary operation is simply addition, then stabilization. We show that the recurrent configurations are closed under this operation below.

Proposition 3.2.9. Let σ, ζ be recurrent chip configurations. Then $S_z(\sigma + \zeta)$ is recurrent.

Proof. Since $S_z(\sigma + \zeta)$ is stable by definition, we must show it is accessible. Let η be an arbitrary configuration. Since σ is accessible, there is some configuration $\xi \geq 0$ (component-wise) such that $S_z(\eta + \xi) = \sigma$. Then since ζ is recurrent, it is non-negative. Thus

$$S_z(\sigma + \zeta) = S_z(S_z(\eta + \xi) + \zeta) \tag{40}$$
$$= S_z(\eta + \xi + \zeta) \tag{41}$$

Thus we have shown $S_z(\sigma + \zeta)$ is reachable from η, by first adding $(\xi + \zeta)$, and then stabilizing. Since η was arbitrary. we are done. □

Finally, we can prove the isomorphism.

Theorem 3.2.10. *Let $R_z(G)$ denote all recurrent configurations on G with global sink vertex z. Let $*$ be a binary operation on $R_z(G)$, defined $(\sigma * \eta) = S_z(\sigma + \eta)$. Then $(R_z(G), *)$ is an abelian group isomorphic to the sandpile group.*

$$(R_z(G), *) \cong S(G)$$

Proof. By 3.2.9, we know that $R_z(G)$ is closed under $*$. Further, 3.2.7 and 3.2.8 together show that there is a bijective map $\phi : S(G) \to R_z(G)$, where ϕ maps each coset to its unique recurrent configuration.

We must show that ϕ is a homomorphism. Let $A, B \in \mathbb{Z}_0^V / \Delta(\mathbb{Z}^V)$. Let $\sigma \in A$ and $\zeta \in B$ denote the recurrent representatives. Then notice that $S_z(\sigma + \zeta)$ is obtained by some sequence of firings from $(\sigma + \zeta)$. Let $u \in \mathbb{N}^V$ be the vector such that u_v is the number of times $v \in V$ is fired from $(\sigma + \zeta)$ to obtain $S_z(\sigma + \zeta)$. Then $(\sigma + \zeta) = S_z(\sigma + \zeta) - \Delta u$. Thus $S_z(\sigma + \zeta) \in (\sigma + \zeta + \Delta(\mathbb{Z}^V))$. Thus,

$$\phi(A+B) = \phi(\sigma + \Delta(\mathbb{Z}^V) + \zeta + \Delta(\mathbb{Z}^V)) \tag{42}$$
$$= \phi((\sigma+\zeta) + \Delta(\mathbb{Z}^V)) \tag{43}$$
$$= S_z(\sigma+\zeta) \tag{44}$$
$$= \sigma * \zeta \tag{45}$$

Since ϕ is a bijective homomorphism, which preserves the group operation of $S(G)$, it immediately follows that $*$ is a group operation on $R_z(G)$. Verifying that $*$ obeys the group axioms is simply a matter of checking the corresponding inverse identities with respect ϕ, which will obey the group axioms since $S(G)$ is a group. Thus $R_z(G)$ is an abelian group, and ϕ is an isomorphism. □

3.3 Dhar's Burning Algorithm and the z-Recurrent Decomposition

So far, we have characterized the sandpile group in two ways: First, it is the set of chip configurations "up to firing equivalence," given by the cokernel of the Laplacian matrix. Second, it is the set of all recurrent chip configurations, under the binary operation of addition and then stabilization.

Neither definition is equipped to deal with granular questions. For example, what does the identity element of the sandpile group look like? So far, all we can say is that such an element exists, and it is in the same coset as $\vec{0}$.

Fortunately, a more concrete characterization of the sandpile group exists. It is given by Dhar's burning test, a simple algorithm which quickly verifies whether a configuration is recurrent.

In addition to giving us an easy way of identifying recurrent configurations, the burning test allows us to identify configurations in the sink-free sandpile model with recurrent configurations in the sink model, via the *z-recurrent decomposition*. Since the sink-free model is less well-behaved than the sink model [8] this decomposition gives a useful bridge between the two. For example, the z-recurrent decomposition forms a key component in Levine's proof of the Threshold Energy Density Theorem [Lev15]. [9]

Throughout this section, let $G = (V, E)$ be an undirected, connected graph with global sink vertex $z \in V$.

Definition 3.3.1. Let G be a graph, and let $z \in V$ be the global sink vertex. The **z-recurrent configurations** on G, denoted $Rec(z)$, are the set of all configurations $\rho \in \mathbb{Z}_0^V$ such that:

i. ρ is stable with respect to z

ii. Every non-sink vertex fires exactly once in the stabilization of $(\rho + \Delta\delta_z)$ with respect to z.

Remark 3.3.2. Our definition is a slight modficiation of the definition of $Rec(z)$ in [Lev15] is slightly different - they require that $\rho(z) = deg(z)$ instead of $\rho^T\vec{1} = 0$. This distinction is just a matter of convention. The purpose of both versions of the definition is just to ensure that two configurations are considered the same if they only differ at the sink vertex. In our case, since we identified recurrent configurations with elements of $\mathbb{Z}_0^V/\Delta(\mathbb{Z}^V)$, it is more convenient to require that $\rho \in \mathbb{Z}_0^V$.

Theorem 3.3.3. *Let $\rho \in \mathbb{Z}^V$. Then $\rho \in Rec(z)$ iff ρ is recurrent (that is, ρ is stable and accessible).*

[8] In particular, not every configuration will stabilize in the sink-free model. Consider the configuration where every vertex v begins with $deg(v)+1$ chips. After every finite sequence of firings, either some vertex never fired and so is still unstable, or some vertex stopped firing first, and thus receives all its original chips back from its neighbors.

[9] For a long, detailed exposition of this result, see the first half of the book [CP18].

Proof. $\implies$: Suppose $\rho \in Rec(z)$. Then ρ is stable by definition. We want to show ρ is accessible.

Let ζ be an arbitrary configuration. Recall that β denotes the configuration obtain by firing the sink once from $\vec{0}$. By definition, every non-sink vertex fires exactly once in the stabilization of $\rho + \beta$. Thus $S_z(\rho + \beta) = \rho + (-\Delta\delta_z) - \Delta(\vec{1} - \delta_z) = \rho + \Delta(\vec{1}) = \rho$. Thus $S_z(\rho + k\beta) = S_z(\rho)$ for all integer $k \geq 1$, by 3.2.5.

Next, since z has positive degree, it has at least one neighbor. Thus β has a positive number of chips at some non-sink vertex which is adjacent to v. Since G is connected, this vertex has a path to every other vertex in the graph. Thus for large enough k, $k\beta$ can be selectively fired to obtain some configuration $\alpha \geq 1$ entry-wise.

Let $m \geq 1$ be a large enough integer such that $\zeta \leq (\rho + m\alpha)$. Then since $k\beta$ can be selectively fired to obtain α, $(km)\beta$ can be selectively fired to obtain $m\alpha$.

Therefore, consider $(\rho + mk\beta)$. We can selectively fire to obtain $\rho + m\alpha$, and then stabilize to obtain ρ (since $S_z(\rho + mk\beta) = \rho$). Since $\zeta \leq \rho + m\alpha$, it follows that ρ is reachable from ζ. Simply add $(\rho + m\alpha - \zeta)$, and then stabilize.

$\iff$: Suppose ρ is stable and accessible.

First, observe that $S_z(\sigma + \beta)$ is accessible. To see this, let ζ be an arbitrary configuration. Then since ρ is accessible, there exists α such that $S_z(\zeta + \alpha) = \rho$. Then $S_z(\zeta + \alpha + \beta) = S_z(S_z(\zeta + \alpha) + \beta) = S_z(\rho + \beta)$. Since $S_z(\rho + \beta)$ is stable, it follows that $S_z(\rho + \beta)$ is reachable from arbitrary ζ, and thus $S_z(\rho + \beta)$ is recurrent.

Next, observe that since $S_z(\rho + \beta)$ is reached from $(\rho + \beta)$ by some sequence of firings, there exists $u \in \mathbb{N}^V$ such that $u_z = 0$ and $S_z(\rho + \beta) = \rho + \beta - \Delta u$. Since $\beta = -\Delta\delta_z$, it follows $S_z(\rho + \beta) = \rho - \Delta(u + \delta_z)$. Thus ρ and $S_z(\rho + \beta)$ are in the same coset with respect to $\Delta\mathbb{Z}^V$. Since both are recurrent, and each coset has a unique recurrent configuration, they must be equal. So $S_z(\rho + \beta) = \rho$.

Finally, notice that if we set $u = \sum_{v \neq z} \delta_v$, then $\rho + \beta - \Delta u = \rho - \Delta\vec{1} = \rho$. Since the firing histogram is unique, we conclude that $u = \sum_{v \neq z} \delta_v$, and so each non-sink vertex fires exactly once in the stabilization of $\sigma + \beta$. Thus $\rho \in Rec(z)$. □

Example 3.3.4. Consider the cycle graph on 3 vertices. Suppose the vertices are numbered (v_1, v_2, v_3) and v_1 is the global sink. Let $(-2, 1, 1)$ be a chip configuration, where the i^{th} entry gives the number of chips at the v_i. Dhar's burning test shows $(-2, 1, 1)$ is recurrent, as

$$(-2, 1, 1) \xrightarrow{\text{fire the sink}} (-4, 2, 2) \xrightarrow{\text{fire } v_2} (-3, 0, 3) \xrightarrow{\text{fire } v_3} (-2, 1, 1)$$

With the characeterization of recurrent states via Dhar's algorithm, we can prove a surprising connection between the sink-free sandpile model and the model with sink, called the z-recurrent decomposition. Roughly speaking, this extends the notion of "firing equivalence" developed in 3.1 in order to give a correspondence between chip configurations in the sink-free model.

Definition 3.3.5. Let $\sigma, \eta \in \mathbb{Z}^V$ be chip configurations. We say σ, η are **z-equivalent** if there exist $u \in \mathbb{Z}^V$ and $m \in \mathbb{Z}$ such that

$$\sigma = \eta + m\delta_z + \Delta u$$

Evidently, z-equivalence is an equivalence relation on chip configurations. Note that z-equivalence is very similar to the notion of equivalence of configurations with respect to cosets of $\Delta\mathbb{Z}^V$. It is slightly more powerful, though, because in addition to coset equivalence, the number of chips at the sink does not matter.

Next, we develop a few lemmas.

Lemma 3.3.6. ([Lev15], Lemma 8): Let $\sigma \in \mathbb{Z}^V$ be a chip configuration, and $z \in V$ the global sink vertex.

a. σ is z-equivalent to a unique $\rho \in Rec(z)$

b. σ is stabilizable if and only if there exists $u \in \mathbb{N}^V$ such that for all $v \neq z$,

$$(\sigma + \Delta u)_v < d_v$$

c. If σ is stabilizable, then for all $v \in \mathbb{Z}^V$, $\sigma + \Delta v$ is stabilizable.

d. If $\sigma \leq \sigma'$ (the inequality is coordinate-wise) and σ' is stabilizable, so is σ.

Proof. a. By adding or subtracting vertices at the sink, we know σ is z-equivalent to some configuration in $\sigma' \in \mathbb{Z}_0^V$. Let $m \in \mathbb{Z}$ be the integer such that $\sigma = \sigma' + m\delta_z$

Next, $(\sigma' + \Delta(\mathbb{Z}^V)) \in S(G)$ is some sandpile group element. Thus there is a unique recurrent configuration ρ in $\sigma' + \Delta(\mathbb{Z}^V)$. Since ρ is recurrent, $\rho \in Rec(z)$. Thus since $\rho \in (\sigma' + \Delta(\mathbb{Z}^V))$, there is some $u \in \mathbb{Z}^V$ such that $\rho = \sigma' + \Delta u$. We conclude that $\sigma = \rho + m\delta_z + \Delta u$.

b. For the forward direction, if σ is stabilizable, there exists a sequence of firings stabilizing σ. Suppose that $i \in V$ is fired u_i times. Then $\sigma + \Delta u$ represents the stabilized configuration, and since each chip is fired a non-negative number of times, $u \in \mathbb{N}^V$.

Conversely, if such a u exists, simply fire each vertex u_i times. The resulting configuration must be stable since all v have fewer than d_v chips.

c. We know there exists $u \in \mathbb{N}^V$ such that $\sigma + \Delta u$ is stabilizable. Then let $c = \max_{i \in V} |v_i|$. Let $\vec{1}$ denote the all-ones vector. Notice that $(v + c\vec{1})$ is non-negative in every component. Then

$$\sigma + \Delta u = \sigma - \Delta v + \Delta v + \Delta(c\vec{1}) + \Delta u \tag{46}$$
$$= \sigma - \Delta v + \Delta(v + c\vec{1} + u) \tag{47}$$

Thus since $v + c\vec{1} + u \geq 0$ component-wise, and $\sigma + \Delta u$ is stable, by 8(b) we conclude that $\sigma - \Delta v$ is stabilizable.

d. Suppose σ' is stabilizable, and $u \in \mathbb{N}^V$ the firings of its stabilizaton. Then $\sigma + \Delta u \leq \sigma' + \Delta u \leq deg - 1$ (entry-wise), so by 8(b) σ is stabilizable. □

Lemma 3.3.7. Suppose $\sigma \in \mathbb{Z}^V$ and σ can be stabilized in the sink-free model. Denote its unique stabilization by $S(\sigma)$. Then there must be some vertex which never fires in the stabilization.

Proof. Suppose for contradiction that every vertex fires at least once. Let v be the first vertex to fire for the last time. Then right after v has fired for the last time, it has at least zero chips. Since v is the first vertex to stop firing, all of its neighbors fire at least once after v is done firing. So v receives at least d_v additional chips. But then v will have at least d_v chips in the final configuration. Thus the final configuration will be unstable, which is a contradiction. □

We are ready to prove the existence and uniqueness of the z-recurrent decomposition, which states that every configuration is z-equivalent to a unique element of $Rec(z)$.

Theorem 3.3.8. ***(Existence and Uniqueness of z-Recurrent decomposition)**: Given $\sigma \in \mathbb{Z}^V$ and sink $z \in V$, there exist unique $\rho \in Rec(z)$, $m \in \mathbb{Z}$, and $u \in \mathbb{Z}^V$ such that $u_z = 0$, and*

$$\sigma = \rho + m\delta_z + \Delta u$$

Moreover, σ is stabilizable in the sink-free model iff $m < 0$.

Proof. By 3.3.6(a), there exists a unique $\rho \in Rec(z)$ such that σ is z-equivalent to ρ, meaning that there are $c \in \mathbb{Z}$, $w \in \mathbb{Z}^V$ such that $s = \rho + c\delta_z + \Delta w$.

We wish to show the uniqueness of c, w. Since G is connected, we know the kernel of Δ is exactly the multiples of $\vec{1}$. Therefore, let $u = w - (w_z)\vec{1}$. Then $\Delta u = \Delta w$ and $u_z = 0$. Since the kernel is exactly the multiples of $\vec{1}$, u is uniquely determined by the fact that it is linearly equivalent to w and that $u_z = 0$.

Next, since $|\Delta u| = 0$, it follows $c = |s| - |\rho|$, so c is uniquely determined. Let $m = c$. Then there are unique ρ, m, v such that $s = \rho + m\delta_z + \Delta u$.

Finally, we want to show that σ is stabilizable (with respect to the sink-free model) iff $m < 0$. If $m < 0$, then $\rho + m\delta_z$ is stable. Then by 3.3.6(c), $\rho + m\delta_z + \Delta v$ is stabilizable, so σ is stabilizable.

If $m = 0$, then $\rho + m\delta_z = \rho$ is unstable, since z will have at least d_z chips.

To show it is not stabilizable, note that by definition, every vertex except z will fire exactly once in the stabilization of $(\rho + \Delta\delta_z)$ with respect to z. Further, by the abelian property, if $S(\rho + \Delta\delta_z)$ exists, then the order of firings does not matter. So suppose we re-order the firings so all firings of z occur at the end. Then since every vertex besides z fires once in the stabilization with respect to z, every vertex will fire at least once if we re-order so z fires at the end. But then z receives at least d_z chips, so it must fire at least once. Then every vertex of G fires in the stabilization (with respect to the non-sink model) of $\rho + \Delta\delta_z$. By the lemma above, it follows that $S(\rho + \Delta\delta_z)$ does not exist.

Thus $\sigma = \rho + \Delta u$ cannot be stabilized, since otherwise we could translate u by $\vec{1}$ and contradict 3.3.6(b).

Thus since σ cannot be stabilized for $m = 0$, by 3.3.6(d) neither can σ for $m > 0$. □

3.4 Cuts and Flows, Again

In this section, we give a final characterization of the sandpile group, in terms of the cut and flow spaces defined in section 2.2. Our exposition largely follows Biggs ([Big97], [Big99]), who studied a version of the sandpile model called the *dollar game*.

Throughout this section, let $G = (V, E)$ denote a connected, undirected graph with arbitrary edge orientation. Let $D \in \mathbb{Z}^{V \times E}$ denote the incidence matrix, and B_I, Z_I denote integral cut and flow space respectively.

Definition 3.4.1. The **Picard group** of a graph G with arbitrary orientation is the quotient group

$$Pic(G) = D(\mathbb{Z}^E)/D(B_I)$$

We will show that the Picard group is isomorphic to the sandpile group.

Proposition 3.4.2. Let Δ be the graph Laplacian, and let $\sigma : \mathbb{Z}^V \to \mathbb{Z}$ be $\sigma(\vec{v}) = \vec{1}^T\vec{v}$. Then:

$$S(G) \cong Ker(\sigma)/\Delta(\mathbb{Z}^V)$$

Proof. Recall that the definition of the sandpile group is just $\mathbb{Z}_0^V/\Delta(\mathbb{Z}^V)$. We are left to show $Ker(\sigma) = \mathbb{Z}_0^V$. This is clear from definitions, as $\mathbb{Z}_0^V = \{v \in \mathbb{Z}^V : \vec{1}^T\vec{v} = 0\} = \{v \in \mathbb{Z}^V : \sigma(v) = 0\} = Ker(\sigma)$. □

Proposition 3.4.3. ([Big99], 7.1) $D(\mathbb{Z}^E) = Ker(\sigma)$.

Proof. $\subseteq$: Clearly if D_e is the column of D corresponding to some $e \in E$, then $\sigma(D_e) = 0$, since $D_{h(e),e} = 1$, $D_{t(e),e} = -1$, and all other entries are zero. Thus since $D(\mathbb{Z}^V)$ is just integer linear combinations of columns of D, we obtain $D(\mathbb{Z}^E) \subseteq Ker(\sigma)$.

$\supseteq$: Let $f \in Ker(\sigma)$. Fix $x \in V$. Then

$$f = \sum_{v \in V} f_v \delta_v \tag{48}$$
$$= (\sum_{v \in V} f_v \delta_v) - \delta_x(\vec{1}^T f) \tag{49}$$
$$= (\sum_{v \in V} f_v \delta_v) - \delta_x(\sum_{v \in V} f_v) \tag{50}$$
$$= \sum_{v \in V} f_v(\delta_v - \delta_x) \tag{51}$$
$$= \sum_{v \in V \backslash \{x\}} f_v(\delta_v - \delta_x) \tag{52}$$
$$\tag{53}$$

Thus, since $f \in \mathbb{Z}^E$, it suffices to show that for all $v \neq x$ that $(\delta_v - \delta_x) \in D(\mathbb{Z}^E)$.

Fix $v \in V \backslash \{x\}$. Since G is connected, there must exist a path from x to v, say $v_0, e_1, v_1, e_2, v_2, ..., v_{r-1}, e_{r-1}, v_r$, where $v_0 = x$ and $v_r = v$. Then

$$\delta_v - \delta_x = \delta_{v_r} - \delta_{v_0} \tag{54}$$
$$= \delta_{v_r} - \delta_{v_{r-1}} + \delta_{v_{r_1}} - \delta_{v_{r-2}} + ... + \delta_{v_1} - \delta_{v_0} \tag{55}$$
$$= D(\pm\delta_{e_r}) + D(\pm\delta_{e_{r-1}}) + ... + D(\pm\delta_{e_1}) \tag{56}$$

Where $\delta_e \in \mathbb{Z}^E$ is the indicator vector at the edge $e \in E$. Clearly each $D(\pm\delta_i) \in D(\mathbb{Z}^E)$. Thus by linearity, it follows $\delta_v - \delta_x \in D(\mathbb{Z}^E)$, and we conclude that $f \in D(\mathbb{Z}^E)$. □

Proposition 3.4.4. ([Big99], 7.2) $D^T(\mathbb{Z}^V) = B_I$.

Proof. $\subseteq$: Let $x \in \mathbb{Z}^V$ and $z \in Z$. Then

$$\langle z, D^T x \rangle = z^T D^T x \tag{57}$$
$$= (D^T x)^T z \tag{58}$$
$$= x^T D z \tag{59}$$
$$= x^T \vec{0} = 0 \tag{60}$$

Since z was arbitrary, it follows that $(D^T x) \in Z^\perp$, so $D^T x \in B$. Further, since D^T is an integer matrix and x is an integer vector, $D^T x \in \mathbb{Z}^E$. Thus $(D^T x) \in \mathbb{Z}^E \cap B = B_I$.

$\supseteq$: If $b \in B_I$, then notice $b = \sum_{v \in V} a_v b_{\{v\}} = \sum_{v \in V} a_v D^T \delta_v$ where $a_v \in \mathbb{Z}$ for all v. Thus since b is an integer linear combination of the rows of D^T, $b \in D^T(\mathbb{Z}^V)$. □

Combining these three propositions gives the result easily.

Theorem 3.4.5. $S(G) \cong Pic(G)$

Proof. We simply perform a series of substitutions.

$$S(G) \cong Ker(\sigma)/\Delta(\mathbb{Z}^V) \quad \text{Proposition 3.4.2} \tag{61}$$
$$\cong D(\mathbb{Z}^E)/\Delta(\mathbb{Z}^V) \quad \text{Proposition 3.4.4} \tag{62}$$
$$\cong D(\mathbb{Z}^E)/(DD^T(\mathbb{Z}^V)) \quad \Delta = DD^T \tag{63}$$
$$\cong D(\mathbb{Z}^E)/(D(D^T\mathbb{Z}^V)) \tag{64}$$
$$\cong D(\mathbb{Z}^E)/D(B_I) \quad \text{Proposition 3.4.3} \tag{65}$$
$$\cong Pic(G) \tag{66}$$

Notice that $S(G)$, $Pic(G)$, and all the quotient groups listed in the substitutions have the same group operation: namely coset addition. So the substitutions give group isomorphisms. □

Remark 3.4.6. Notice that under both the usual 1-norm and 2-norm for Euclidean vectors, the least nonzero vector in B_I corresponds to the minimum cut of the graph. Given the characterization of $S(G)$ as $D(\mathbb{Z}^E)/D(B_I)$, we might speculate that the minimum cut vector in B_I might correspond to some recurrent configuration of $S(G)$, with algorithmic implications to follow. However, it is unclear how to make use of this identification, since B_I contains *all* cuts, and each element of the sandpile group corresponds to a coset of $D(B_I)$. Thus, the isomorphism presented here seems too coarse-grained to give information about any particular cut. Instead, it characterizers the entire collection of integral cuts.

4 Smith Normal Forms and Sandpiles

For any mathematical object, it is natural to ask what its simpler constituent parts are, and whether it can be described completely in terms of these simpler components. For finite abelian groups, the basic building blocks are cyclic groups, whose multiplicities and orders are described by the *invariant factors* of the group. Up to isomorphism, every finite abelian group is completely captured by these invariant factors, as we describe in 4.1.

We are thus interested in the invariant factors of sandpile groups. Towards this end, in 4.2 we will introduce the *Smith Normal Form* of an integer matrix, which is a convenient tool to compute the invariant factors of its cokernel group. Then in 4.3, we will prove that the invariant factors of the sandpile group are indeed given by the cokernel of the reduced Laplacian matrix. This will allow us to prove several remarkable facts about sandpile groups, including the fact that every finite abelian group is the sandpile group of some graph.

Finally, in section 4.4 we will turn our attention to the *trivial count*, which counts the number of Smith factors equal to 1. We demonstrate lower bounds based on diameter of the graph, and in terms of product graphs. Applied to the hypercube graph, the resulting lower bound is tight. Finally, we present the results of numerical experiments concerning the Smith factors of grid graphs and certain expander graphs, and formulate conjectures on the basis of these experiments.

4.1 Invariant Factors of Finite Abelian Groups

In this section, we review the Sylow theorems and the fundamental theorem of finitely generated abelian groups. Together, these results give a classification of finite abelian groups in terms of their *invariant factors*, which are simply the multiplicities and orders of their cyclic subgroups of the form $\mathbb{Z}/n\mathbb{Z}$.

Theorem 4.1.1. *(Sylow): Let G be a (not necessarily commutative) group of order n. If $n = p^\alpha m$ for some prime p, and integer $\alpha \geq 1$, and $gcd(p, m) = 1$, then*

(i) There exists a subgroup of G order order p^α. We call this a Sylow p-subgroup of G.

(ii) If $P, Q \leq G$ are p-subgroups and P is a sylow p-subgroup, then Q is contained in some conjugation of P. That is, there exists $g \in G$ such that

$$Q \leq gPg^{-1}$$

(iii) The number of sylow p-subgroups of G is $(1 + kp)$ for some $k \in \mathbb{N}$.

(iv) Let P be a Sylow p-subgroup. Then P is the unique Sylow p-subgroup of G iff P is conjugate in G. In particular, if G is abelian, then there exists a unique Sylow p-subgroup for all p dividing $|G|$.

Proof. See, for example, [DF04] section 4.5. □

Theorem 4.1.2. *(Fundamental Theorem of Finitely Generated Abelian Groups):Let G be a finitely generated abelian group. Then there exist unique $r, s \in \mathbb{N}$ and $n_1, ..., n_s \in \mathbb{N}$ such that*

$$G \cong \mathbb{Z}^r \times (\mathbb{Z}/n_1\mathbb{Z}) \times \cdots \times (\mathbb{Z}/n_s\mathbb{Z})$$

Subject to:
(i) For all j, $n_j \geq 2$
(ii) For $1 \leq i \leq s - 1$, $n_{i+1} | n_i$

Proof. See, for example, [DF04] section 5.2. □

Corollary 4.1.3. (i) G is finite iff $r = 0$

(ii) If $r = 0$, then $|G| = n_1 \cdot n_2 \cdots n_s$

(iii) If p is a prime and $p \mid |G|$, then there exists $i \in [s]$ such that $p|n_i$. Consequently, $p|n_{i-1}, p|n_{i-2}, ..., p|n_1$. In particular, $p|n_1$ for all prime p dividing $|G|$.

Putting these results together, we obtain a simple classification of finite abelian groups.

Theorem 4.1.4. *(Classification of Finite Abelian Groups): Let G be a finite abelian group of order n. Let n have prime decomposition*

$$n = p_1^{\alpha_1} p_2^{\alpha_2} \cdots p_k^{\alpha_k}$$

(i) Let A_i be the unique Sylow p_i-subgroup of G. Then

$$G \cong A_1 \times A_2 \cdots \times A_k$$

(ii) For each A_i, there exist $\beta_1, ..., \beta_t \in \mathbb{N}^+$ which are a partition of α_i, and such that

$$A_i \cong \mathbb{Z}/p_i^{\beta_1}\mathbb{Z} \times \cdots \times \mathbb{Z}/p_i^{\beta_t}\mathbb{Z}$$

These $\beta_1 ..., \beta_t$ are precisely the invariant factors of the group A_i, which exist and are unique by the fundamental theorem of finitely generated abelian groups.

Remark 4.1.5. Theorem 4.1.4 gives a simple way of counting and enumerating all abelian groups of order n, up to isomorphism. We can summarize the steps as follows.

Input: Abelian group of order n.
Step 1: Find the prime decomposition of n, say $n = p_1^{\alpha_1} p_2^{\alpha_2} \cdots p_k^{\alpha_k}$
Step 2: For each α_i, count the number of ways to partition the positive α_i into a sum of positive integers.
Step 3: For each partition $\beta_1, ..., \beta_t$ of α_i, obtain a group of order $p_i^{\alpha_i}$, namely

$$\mathbb{Z}/p_i^{\beta_1}\mathbb{Z} \times \cdots \times \mathbb{Z}/p_i^{\beta_t}\mathbb{Z}$$

Output: Conclude that there are $q_1 \cdots q_k$ unique groups of order n, by taking the product of k groups of order $p_1^{\alpha_1}, ..., p_k^{\alpha_k}$ respectively via step 3.

For example, since $36 = 2^2 3^2$ there are four abelian groups of order 36, namely:

- $(\mathbb{Z}/2\mathbb{Z})^2 \times (\mathbb{Z}/3\mathbb{Z})^2$
- $(\mathbb{Z}/4\mathbb{Z}) \times (\mathbb{Z}/3\mathbb{Z})^2$
- $(\mathbb{Z}/2\mathbb{Z})^2 \times (\mathbb{Z}/9\mathbb{Z})$
- $(\mathbb{Z}/4\mathbb{Z}) \times (\mathbb{Z}/9\mathbb{Z})$

4.2 Existence and Uniqueness of Smith Normal Form

Since the sandpile group is a finite abelian group, it is natural to ask what the invariant factors of the sandpile group are, for families of interesting graphs. Here, the definition of the sandpile group in terms of the cokernel of the Laplacian matrix is useful. Given any integer matrix, its cokernel is an abelian group with a finite number of generators (namely, the columns of the matrix).

The Smith Normal Form of a (square) integer matrix M is a unique diagonal matrix whose entries give the invariant factor decomposition of $coker(M)$. Studying the Smith Normal Forms is thus a valuable tool for understanding the sandpile group, and cokernel groups more generally.

Definition 4.2.1. For an integer matrix $M \in \mathbb{Z}^{n\times n}$ and $k \in [n]$, the k^{th} **determinantal divisor** of M is the g.c.d. (greatest common divisor) of all $(k \times k)$ minors of M. Denoting this quantity by $d_k(M)$, we have

$$d_k(M) = |gcd\{det(M_{I,J}) : I, J \subset [n], |I| = |J| = k\}|$$

Define $d_0(M) = 1$.

We can now define the Smith Normal Form of a square integer matrix.[10]

Definition 4.2.2. A square integer matrix $M \in \mathbb{Z}^{n\times n}$ is in **Smith Normal Form** if its only nonzero entries are on its main diagonal, and if its diagonal entries $M_{1,1}, ..., M_{n,n}$ are given by

$$M_{i,i} = \frac{d_i(M)}{d_{i-1}(M)}$$

That is, the diagonal entries are ratios of successive determinantal divisors of M. The nonzero diagonal entries $M_{1,1}, ..., M_{r,r}$ are called the **Smith invariant factors** of M.

There is a relatively simple algorithm for computing the Smith Normal Form, which we state in the box below. By proving that the algorithm converges to a unique solution for every input, we will have proved existence and uniqueness.

Smith Normal Form Algorithm (Adapted from [CP18], 2.33)

Input: Integer matrix $M \in \mathbb{Z}^{n\times n}$

Step 1: If $n = 1$, return $|M_{1,1}|$. If M contains all zeroes, return M.

Step 2: Permute the rows and columns of M so that $M_{1,1}$ is the least nonzero entry in absolute value. If $M_{1,1}$ is negative, set $M_{1,1}$ to $-M_{1,1}$.

Step 3: Add integer multiples of row 1 to other rows, and column 1 to the other columns, until every entry in the first row and column of M besides $M_{1,1}$ is zero. If at any point a nonzero entry in the matrix is encountered which is smaller than $M_{1,1}$ in absolute value, go back to step 2.

Step 4: Let M' be the submatrix of M consisting of rows $2, ..., n$ and columns $2, ..., n$. If any entry in M', say $M_{i,j}$, is not divisible by $M_{1,1}$, add column j to column 1. Go back to step 2.

Step 5: Let M' be the submatrix of M consisting of rows $2, ..., n$ and columns $2, ..., n$. Compute the Smith Normal Form of M' by returning to step 1.

Output: Return a matrix whose first diagonal entry is $M_{1,1}$ and whose subsequent diagonal entries are from the Smith Normal Form of M'.

Theorem 4.2.3. *For any integer matrix $M \in \mathbb{Z}^{n\times n}$, the Smith Normal Form of M, denoted $S \in \mathbb{Z}^{n\times n}$ exists and is unique. It is the unique output of the algorithm described above.*

Further, there exist integer matrices $A, B \in \mathbb{Z}^{n\times n}$ such that $det(A) = det(B) = \pm 1$ and

$$AMB = S$$

Proof. Induction on n.

Base case: For $n = 1$, the algorithm just returns the sole entry of M. Since there is only one choice of (1×1) submatrix, $d_1(M) = |M_{1,1}|$. Since $d_0(M) = 1$ by definition, it follows that $|M_{1,1}| = \frac{d_1(M)}{d_0(M)}$, so it is indeed the sole Smith invariant factor.

[10] As it turns out, the Smith Normal Form exists for for any matrix with entries in a principal ideal domain. To simplify presentation, we include only the necessary results about square matrices with entries in $\mathbb{Z}$.

Inductive step: Suppose the algorithm is correct up to some $k \in \mathbb{N}$. Let $n = k + 1$.

Observe that permuting rows and columns, and adding integer multiples of rows and columns to each other, are all elementary row operations. So they do not change the determinantal divisors of M. Moreover, each such operation corresponds to either left or right-multiplying M by a suitable elementary row matrix. Taking the product of these matrices gives the unimodular matrices A, B in the statement of the theorem.

Thus if step 5 is reached, then by inductive hypothesis the last k Smith Invariant Factors of M are correctly computed, by inductive hypothesis. We are left to prove correctness and convergence of steps 1 through 4. We proceed by cases:

Case 1: The algorithm halts on step 1. Then M consists of all zeroes. Thus all determinantal divisors are zero, and so all Smith Invariant Factors are zero.

Case 2: M does not halt on step 1. Then M does not consist of all zeros. After step 2, M will be permuted so that $M_{1,1}$ is minimal among the nonzero entries of M in absolute value.

Case 2(a): If $M_{1,1}$ is the greatest common divisor of all entries of M, then steps 3 and 4 will proceed without returning to step 2, and $M_{1,1}$ will correctly be returned as the first Smith invariant factor.

Case 2(b): Suppose instead that $M_{1,1}$ is not the greatest common divisor of all entries of M. Then in either step 3 or 4, the algorithm will return to step 2 after some row/column permutations and additions. As argued above, none of these operations change the determinantal divisors.

We claim that regardless of whether step 3 or 4 triggers the return to step 2, the $(1, 1)$ entry of M will strictly decrease in absolute value. This guarantees that step 5 is reached in finitely many iterations of steps (2-4).

Case 2(b)(i): Step 3 redirects to step 2. Then some $M_{i,j} \neq M_{1,1}$ is encountered such that $0 < |M_{i,j}| < |M_{1,1}|$. Thus the new $(1, 1)$ entry is strictly smaller in absolute value.

Case 2(b)(ii): Step 4 redirects to step 2, after step 3 completes an iteration without returning to step 2. Then the $M_{i,j}$ from step 4 is added to the zero entry in $M_{i,1}$. Thus $M_{i,1} = M_{i,j}$. Notice the only nonzero entries in the first row and column of M are $M_{1,1}, M_{i,1}$. Since $M_{1,1}$ does not divide $M_{i,1}$, there is some $a \in \mathbb{Z}$ such that $0 < |M_{i,1} - aM_{1,1}| < |M_{1,1}|$. Thus subtracting a multiples of the first row from the i^{th} row will result in the $(i, 1)$ entry being less than the $(1, 1)$ entry in absolute value. Thus, the first and i^{th} rows of M to be swapped, and the $(1, 1)$ entry will be strictly smaller in absolute values.

Thus, in either case the $(1, 1)$ entry of M strictly decreases. So after a finite number of steps, Step 5 will be reached.

Since the value of $M_{1,1}$ at step 5 is guaranteed to divide all other entries of the matrix, and determinantal divisors are unchanged by row/column permutations and additions, we conclude that $M_{1,1}$ is the correct first Smith invariant factor. Thus by inductive hypothesis, we are done. □

Having shown that the Smith Normal Form exists and is unique, we will next relate the structure of a sandpile group of a graph to the Smith Normal Form of its reduced Laplacian matrix.

4.3 Smith Invariant Factors of Sandpile Groups

Having established the relevant background and motivation, in this section we can use the Smith Normal Form of the reduced Laplacian matrix to completely describe the sandpile group of a graph. As we then show, the Smith Normal Form allows one to easily deduce several powerful facts about sandpiles as mere corollaries.

Theorem 4.3.1. *Given a finite, connected graph G with Laplacian matrix $\Delta \in \mathbb{Z}^{n\times n}$, let $z \in V$ be a fixed sink vertex. Then if $\Delta_z \in \mathbb{Z}^{(n-1)\times(n-1)}$ denotes the reduced Laplacian along the row/column*

corresponding to z, let $S \in \mathbb{Z}^{(n-1)\times(n-1)}$ denote the unique Smith Normal Form of Δ_z. Then for some unimodular matrices A, B,

$$\Delta_z = ASB = A \begin{bmatrix} \alpha_1 & 0 & \cdots & 0 \\ 0 & \alpha_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \alpha_{n-1} \end{bmatrix} B$$

Here, $\alpha_1|\alpha_2$, $\alpha_2|\alpha_3$, and so on. Each $\alpha_i \geq 0$.

Then if we let $\mathbb{Z}/0\mathbb{Z}$ denote the trivial group, the sandpile group is given by:

$$S(G) \cong \prod_{i=1}^{n-1} \mathbb{Z}/\alpha_i\mathbb{Z}$$

Proof. First, notice that since G is connnected, its kernel is one-dimensional and spanned by $\vec{1}$. Thus Δ_z is invertible, and we can write the sandpile group as the cokernel of Δ_z. That is, $S(G) = \mathbb{Z}_0^V/\Delta(\mathbb{Z}_V)$. Since $\mathbb{Z}^{V-z} \cong \mathbb{Z}_0^V$ and $\Delta(\mathbb{Z}_V) \cong \Delta_z(\mathbb{Z}^{V-z})$, it follows that $S(G) \cong \mathbb{Z}^{n-1}/\Delta_z(\mathbb{Z}^{n-1})$.

Next, let the Smith Normal Form of Δ_z be $\Delta_z = ASB$. Notice that A is a finite product of row operation matrices, and B is a finite product of column operation matrices. In fact, the row/column operations are given by:

i. Swapping two rows (column)

ii. Negating a row (column)

iii. Adding one row (column) to another

All of these operations do not change the span of the lattice [11] generated by the rows (columns) of S.

Thus, we can write $A^{-1}\Delta_z B^{-1} = S$, and $coker(\Delta_z) = coker(S)$.

Next, the generating set of the lattice spanned by the columns of S is simply $\alpha_1 e_1, ..., \alpha_{n-1}e_{n-1}$. So

$$S(G) \cong coker(\Delta_z) = coker(S) \cong \mathbb{Z}^{n-1}/(\prod_{j=1}^{n-1} \alpha_j\mathbb{Z}) \cong \prod_{j=1}^{n-1} (\mathbb{Z}/\alpha_j\mathbb{Z})$$

□

We can now derive two amazing facts as easy corollaries.

Corollary 4.3.2. The order of the sandpile group of G is equal to the number of labeled spanning trees of G, or its *tree number* $\kappa(G)$.

$$|S(G)| = \kappa(G)$$

Proof. By the matrix-tree theorem, the number of labeled spanning trees is equal to any cofactor of the Laplacian matrix. Thus since $\Delta_z = ASB$ and A, B are unimodular, it follows that:

$$\kappa(G) = det(\Delta_z) = det(A)det(S)det(B) = det(S) = \prod_{i=1}^{n-1} \alpha_i = |S(G)|$$

□

[11] By *lattice* we mean the set of all linear combinations with integer coefficients.

Corollary 4.3.3. For a finite, connected graph G, the sandpile group is independent of choice of sink.

Proof. Let $y, z \in V$ be distinct choices of sink. Let M_y be the matrix obtained by swapping rows $1, y$ and columns $1, y$ of Δ. Let M_z be analogously defined. Since Smith invariant factors are unchanged by row/column swaps, the output of the Smith Normal form algorithm on M_y, M_z will be equal. Since $det(\Delta) = 0$, the n^{th} Smith invariant factor will be zero. The factors $1, ..., n-1$ will be equal, and correspond to the Smith factors of the reduced Laplacians Δ_y, Δ_z respectively. Thus the sandpile groups with y, z as sink have the same Smith invariant factors, and thus must be isomorphic. □

Remark 4.3.4. If G_1, G_2 are isomorphic graphs, then a permutation of the rows/columns of the Laplacian of G_1 gives G_2. Thus, the Smith normal forms of Δ_1, Δ_2 are equal. In particular, we obtain that the sandpile group is an *isomorphism invariant* of its graph, so $G_1 \cong G_2 \Rightarrow S(G_1) \cong S(G_2)$. However, the converse is not true. Consider the path graph and star graph on n vertices. It can be shown that any tree has a trivial sandpile group. So both the path and star graph have isomorphic sandpile groups, but the graphs are clearly not isomorphic for $n > 3$.

In general, not much is known about the exact Smith invariant factors of various classes of graphs. We gather some known results in the following table.

Graph Notation	Description	$(n-1)$ Smith Invariant Factors
K_n	Complete graph on n vertices	$(1, n, ..., n)$ (see [Big97], Sec. 30)
T_n	Tree on n vertices	$(1, 1, ..., 1)$
C_n	Cycle graph on n vertices	$(1, 1, ..., 1, n)$ (see Prop 4.3.6)
$K_{n_1,n_2,...,n_k}$	Complete multipartite graph on $\sum_{i=1}^{k} n_i$ vertices	Known (see [JNR03])
H_n	Boolean hypercube on 2^n vertices	Partial Results (see [Bai03])
$\prod_{i=1}^{k} K_{n_i}$	Product of complete graphs on $n_1, ..., n_k$ vertices	Partial results (see [JNR03])

In the case of trees, since the number of labeled spanning trees is exactly 1 (namely, the graph itself), we can exactly determine the sole element of sandpile group.

Proposition 4.3.5. Let G be an arbitrary connected graph with sink z. Let $s \in \mathbb{Z}^V$ be the configuration where for $v \in V \setminus \{z\}$, $s_v = deg(v) - 1$, and s_z is arbitrary. Then s is a critical configuration.

Proof. We claim that s satisfies Dhar's burning algorithm.

Fire the sink once. Then every neighbor of the sink is unstable, so fire each of them once. Then each of the neighbors of those vertices, which have not already been fired, are unstable, so fire those. Inductively, since G is connected, every vertex will fire at least once.

Suppose that during the stabilization, some $v \neq z$ fires twice. Then it received at least $deg(v)+1$ chips from its neighbors throughout the stabilization. Thus, at least one of its neighbors fired twice. Suppose $w \neq z$ is the first vertex to fire twice. Then upon its second firing, it will have at most $deg(w) - 1$ chips, since it lost $deg(w)$ chips during its first firing and then gained back at most 1 chip from each neighbor. But then w will not fire twice. Thus no vertex fires twice.

Since every vertex fires exactly once, by Dhar's burning test s is critical. □

As previously shown, every (finite, connected, undirected) graph has a corresponding finite abelian in the form of its sandpile group. Remarkably, a partial converse to this statement is also true: that is, every finite abelian group is the sandpile group of some graph!

First, we need to find a graph whose sandpile group is cyclic. As one might guess, the cycle graph has this property.

Proposition 4.3.6. Let C_n be the undirected cycle graph on n vertices. Then

$$S(C_n) \cong \mathbb{Z}/n\mathbb{Z}$$

Proof. Observe that there are n labeled spanning trees of G, so the order of $S(C_n)$ is n. Next, notice the reduced Laplacian of C_n along any single vertex is of the form

$$\begin{bmatrix} 1 & -1 & 0 & 0 & \cdots & 0 \\ -1 & 1 & -1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & -1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ -1 & 0 & 0 & \cdots & -1 & 1 \end{bmatrix}$$

Thus, deleting the leftmost column and bottom row gives a lower-triangular sub-matrix, with -1 along its main diagonal. Thus its determinant thus ± 1. Deleting successive rows and rows preserves the upper-triangularity of the sub-matrix. Thus $d_0(\Delta_z) = d_1(\Delta_z) = ... = d_{n-2}(\Delta_z) = 1$. Thus every Smith factor but the largest equals 1.

Finally, the product of the Smith Invariant factors gives the order of G. Thus the largest Smith factor must equal n, and we conclude that $S(C_n) \cong (\mathbb{Z})^{n-1} \times \mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/n\mathbb{Z}$. □

Next, we define and prove a more general statement about the sandpile group of a "glued graph."

Definition 4.3.7. Let G_1, G_2 be undirected graphs. Let G be the graph formed by identifying G_1, G_2 along a single vertex v. That is, select some $v_1 \in V(G_1)$ and $v_2 \in V(G_2)$. Then $V(G) = \{v\} \cup (V(G_1) \setminus \{v_1\}) \cup (V(G_2) \setminus \{v_2\})$.

For $w \in V(G_1)$, $(v, w) \in E(G) \iff (v_1, w) \in E(G_1)$, and for $w \in V(G_2)$, $(v, w) \in E(G) \iff (v_2, w) \in E(G_2)$. All edges not incident to v_1, v_2 in G_1, G_2 are included as well. Then v is the **articulation point** of G, and G_1, G_2 are the articulated components.

Proposition 4.3.8. ([Kli18] 4.5.9): Let G be the graph with articulation point v and articulated components $G_1, ..., G_k$. Then the sandpile group of G is the product of the sandpile groups of its articulated components.

$$S(G) \cong S(G_1) \times S(G_2) \cdots \times S(G_k)$$

Proof. Consider the reduced Laplacian L_v of G, which has the row and column corresponding to v deleted. Evidently L_v can have its rows and columns permuted to become block-diagonal, where each block is the Laplacian of some articulated component G_i reduced along its row and column corresponding to v.

Let S be the Smith Normal Form of L_v. Observe that the blocks of S give the smith normal forms for the reduced Laplacian of each G_i. Thus, the Smith invariant factors of the i^{th} diagonal block give the invariant factors of $S(G_i)$. The result follows. □

The desired result then follows from the previous two facts.

Corollary 4.3.9. ([Kli18] 4.5.9): Every finite abelian group is the sandpile group of some graph.

Proof. Let G be a finite abelian group. Take its elementary divisor decomposition. For each cyclic factor $\mathbb{Z}/m\mathbb{Z}$, take the cycle graph on m vertices. Glue all of the cycle graphs on a single vertex. □

4.4 Trivial Counts of Products and Expanders: Experiments and Conjectures

In general, it is difficult to determine the Smith invariant factors for the sandpile group of a family of graphs. The proof of 4.3.6, for example, is instructive, as it relies entirely on the particularly simple structure of the Laplacian of C_n. While we might hope to apply similar methods to other graphs with "nice" Laplacians (and indeed, this is exactly what others have done - see e.g. [Big97], [Bai03]), it is unclear what to do absent some clever argument about the Laplacian structure.

Given this difficulty, we are interested in relaxations and easier versions of the broader question, which is understanding the entire invariant factor decomposition of a sandpile group. This section will study one such relaxation, which asks how many Smith invariant factors are *trivial* (equal to 1). These factors are called trivial since they correspond to the trivial group $\mathbb{Z}/\mathbb{Z}$ in the invariant factor decomposition of the sandpile group, and thus make no contribution to its structure.

Definition 4.4.1. If G is a connected, undirected graph, let $b(G)$ denote the multiplicity of 1 as a smith factor of $S(G)$. We call $b(G)$ the **trivial count** of $S(G)$. Notice that $b(G)$ is the largest value of k such that $s_k(G) = 1$.

Recall that $b(C_n) = n - 2$, $b(K_n) = 1$, and $b(T_n) = n - 1$, where T_n is any tree on n vertices. Thus, heuristically we can say that if G has large $b(G)$, then the structure of the sandpile group is "closer" to a tree or cycle graph than a complete graph.

Of course, this begs the question: What exact properties of trees, cycles, and complete graphs explain this discrepancy in their trivial counts? It is difficult to say, and in this section we will try to better understand the meaning and value of the trivial count in various graphs.

First, we review some known lower bounds on the trivial count.

Proposition 4.4.2. ([RMW93], Theorem 1): Let G be a connected, undirected graph. Let $S \subset V$, $|S| = k$, and suppose that $G[S]$ (the induced subgraph) is a forest on w components. Suppose each component is a path on at least two vertices. Then $d_{k-w}(G) = 1$ (the $(k-w)^{th}$ determinantal divisor of $\Delta(G)$), so $b(G) \geq k - w$.

Proof. Order the vertices of V so that S gives the first k vertices. Order the vertices of S according to the forests in $G[S]$. For each $(r \times r)$ submatrix in the top left $(k \times k)$ corresponding to a tree in $G[S]$ forest, removing the first row and last column gives an upper triangular matrix with -1 on the main diagonal. Therefore the determinant of this submatrix is ± 1.

We can repeat this procedure for every path inside $G[S]$. Thus since there are w components, there exists a $(k-w) \times (k-w)$ submatrix of determinant ± 1, and thus $d_{k-w}(G) = 1$. □

Remark 4.4.3. Notice that this argument implicitly proves that the path graph P_n has $b(G) \geq n - 2$.

Corollary 4.4.4. $b(G) \geq diam(G)$ for connected G.

Proof. If G has diameter d, then there exist $u, v \in V$ such that $u = v_1, v = v_{d+1}$, and $v_1, ..., v_{d+1}$ is a path through G forming the shortest path from u to v. Then let $S = \{v_1, ..., v_{d+1}\}$. Observe that $G[S]$ is a path, since if there were any edges outside the $u \to v$ path, these would give a shorter $u \to v$ path.

Therefore $b(G) \geq (d+1) - 1 = d$. □

Corollary 4.4.5. If G is connected, then $b(G) \geq 1$, with equality iff $G = K_n$.

Proof. First, $diam(G) \geq 1$, so $b(G) \geq diam(G) \geq 1$.

Second, if $G = K_n$ then we know $b(K_n) = 1$. If $G \neq K_n$ then $diam(G) \geq 2$, so $b(G) \geq 2 > 1$. □

Next, we turn our attention to products of graphs, whose Laplacian matrices can be easily analyzed with block-matrix methods. We need a few lemmas first.

Definition 4.4.6. Given a graph G on n vertices, and $k \leq n$, let $d_k(G) = d_k(\Delta(G))$. That is, $d_k(G)$ denotes the k^{th} determinantal divisor of $\Delta(G)$.

Proposition 4.4.7. If G is a connected, undirected graph, and $d_k(G) = \pm 1$, then $b(G) \geq k$.

Proof. Recall that if s_m is the m^{th} Smith invariant factor, that $s_m = \frac{d_m(G)}{d_{m-1}(G)}$.

If $d_k(G) \pm 1$ then there is some $(k \times k)$ minor of $\Delta(G)$ equal to ± 1. Thus since $\Delta(G)$ is an integer matrix, all submatrices of this minor have determinant ± 1. Thus $d_0(G) = d_1(G) = ... = d_k(G) = \pm 1$.

We conclude that $s_1 = \left|\frac{d_1(G)}{d_0(G)}\right| = 1$, and $s_2 = \left|\frac{d_2(G)}{d_1(G)}\right| = 1$, and so on up to s_k. Thus $b(G) \geq k$. □

Lemma 4.4.8. Let A be a $(2n \times 2n)$ block-diagonal matrix of the form

$$A = \begin{bmatrix} M_1 & O \\ N & M_2 \end{bmatrix}$$

Where O is an $(n \times n)$ matrix of zeroes. Then

$$det(A) = det(M_1)det(M_2)$$

Proof. Let S_k denote the set of permutations of a set of k elements.

The Leibniz formula for determinants gives

$$det(A) = \sum_{\sigma \in S_{2n}} sgn(\sigma) \prod_{i \in [2n]} A_{i,\sigma(i)}$$

Let $S' \subset S_{2n}$ be all $\pi \in S_{2n}$ such that $\pi([n]) = [n]$ (so π sends $[n]$ to $[n]$, and $\{n+1, ..., 2n\}$ to $\{n+1, ..., 2n\}$)

Suppose $\sigma \in S_{2n} \setminus S'$. Then there exists some $j \in [n]$ such that $\sigma(j) > n$. Then $A_{j,\sigma(j)} = 0$, so this permutation contributes nothing to the summation in the Leibniz formula for $det(A)$. It follows that

$$det(A) = \sum_{\sigma \in S'} sgn(\sigma) \prod_{i \in [2n]} A_{i,\sigma(i)}$$

Evidently each $\sigma \in S'$ is unique determined by a pair of elements $\pi, \pi' \in S_n$, where π determines how σ acts on $[n]$, and π' determines how σ acts on $[2n] \setminus [n]$. Thus

$$sgn(\sigma) \prod_{i \in [2n]} A_{i,\sigma(i)} = sgn(\pi)sgn(\pi') \prod_{i \in [n]} A_{i,\pi(i)} A_{i+n,\pi(i)+n}$$

Taking the summation over all $\sigma \in S_{2n}$, it follows that $det(A) = det(M_1)det(M_2)$. □

An almost identical argument gives the generalization to block-triangular matrices.

Corollary 4.4.9. Let A be a $(kn \times kn)$ block-diagonal matrix of the form

$$A = \begin{bmatrix} M_1 & O & O & \cdots & O \\ * & M_2 & O & \cdots & O \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ * & * & * & \cdots & M_k \end{bmatrix}$$

Where O is an $(n \times n)$ matrix of zeroes, and $*$ is any $(n \times n)$ matrix. Then

$$det(A) = det(M_1)det(M_2)\cdots det(M_k)$$

Now, we can prove a lower bound on any product involving a cycle or path graphs. Our method is a generalization of the ideas in [Bai03].

Proposition 4.4.10. Let G_n be any graph on n vertices, P_m denote the path graph on m vertices, and C_m denote the cycle graph on m vertices. Then

i. $b(G_n \times P_m) \geq n(m-1)$

ii. $b(G_n \times C_m) \geq n(m-2)$

Proof. i. Let $L_n = L(G_n)$ and let $L_{n,k} = L_n + kI$. Let I denote the $(n \times n)$ identity matrix and O the $(n \times n)$ matrix of all zeros. Then notice that

$$L(G_n \times P_m) = \begin{bmatrix} L_{n,1} & -I & O & \cdots & O \\ -I & L_{n,2} & -I & \cdots & O \\ O & -I & L_{n,2} & \cdots & O \\ \vdots & \vdots & \ddots & \vdots & \\ O & O & \cdots & -I & L_{n,1} \end{bmatrix}$$

Deleting the first column and last row (in block form) gives a lower triangular matrix whose main diagonal has $(m-1)$ copies of $-I_n$. Therefore $d_{n(m-1)} = 1$, since there is an $n(m-1) \times n(m-1)$ block of the diagonal matrix whose determinant is ± 1. Deleting additional rows and columns gives that $d_1 = d_2 = ... = d_{n(m-1)} = 1$. Thus the first $n(m-1)$ Smith invariant factors equal 1.

ii. With the same notation, the Laplacian in block form is given by:

$$L(G_n \times C_m) = \begin{bmatrix} L_{n,1} & -I & O & \cdots & -I \\ -I & L_{n,2} & -I & \cdots & O \\ O & -I & L_{n,2} & \cdots & O \\ \vdots & \vdots & \ddots & \vdots & \\ -I & O & \cdots & -I & L_{n,1} \end{bmatrix}$$

Deleting the first row, last row, first column, and last column gives a lower triangular matrix whose main diagonal has $(m-2)$ copies of $-I_n$. Therefore $d_{n(m-2)} = 1$, since there is an $n(m-2) \times n(m-2)$ block of the diagonal matrix whose determinant is ± 1. So the first $n(m-2)$ Smith invariant factors equal 1. □

Writing certain families of graphs as products with a path or cycle, we can give lower bounds on the trivial count.

Corollary 4.4.11. i. Let H_n denote the hypercube graph on 2^n vertices. Then $b(H_n) \geq 2^{n-1}$.

ii. Let $Gr_{n,m} := P_n \times P_m$ denote the grid graph on $n \cdot m$ vertices, with no boundary. Then $b(Gr_{n,m}) \geq \max\{n(m-1), m(n-1)\}$.

iii. Let $Gr'_{n,m} := C_n \times C_m$ denote the grid graph on $n \cdot m$ vertices, with boundary. Then $b(Gr'_{n,m}) \geq \max\{n(m-2), m(n-2)\}$.

Proof. i. Notice that $H_n = H_{n-1} \times P_2$. Thus since H_{n-1} has 2^{n-1} vertices, $b(H_n) = b(H_{n-1} \times P_2) \geq 2^{n-1}$.

ii. The graph products commute in this case - that is, $P_n \times P_m \cong P_m \times P_n$. Apply the proposition twice.

iii. The graph products commute in this case - that is, $C_n \times C_m \cong C_m \times C_n$. Apply the proposition twice. □

As it turns out, the bound in (i) is tight - see [Bai03]. We conjecture that the bound in (ii) is tight when $n = m$, and the bound in (iii) is tight when $n = m$ and both are even. Numerical experiments support these conjectures for small n, m - see Figure 1.

Conjecture 1: (a) $b(P_n \times P_n) = n(n-1)$.
(b) If n is even, then $b(C_n \times C_n) = n(n-2)$.

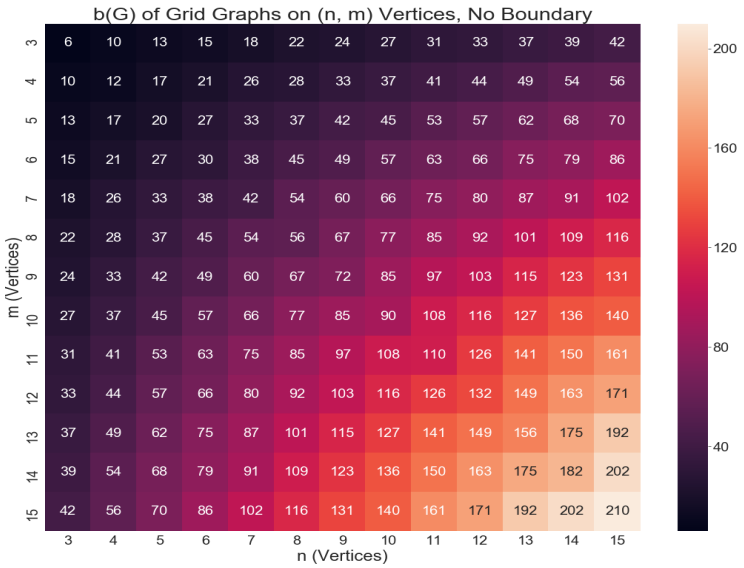

(a) Grid graphs without boundary. Entry (n, m) corresponds to $Gr_{n,m} = P_n \times P_m$.

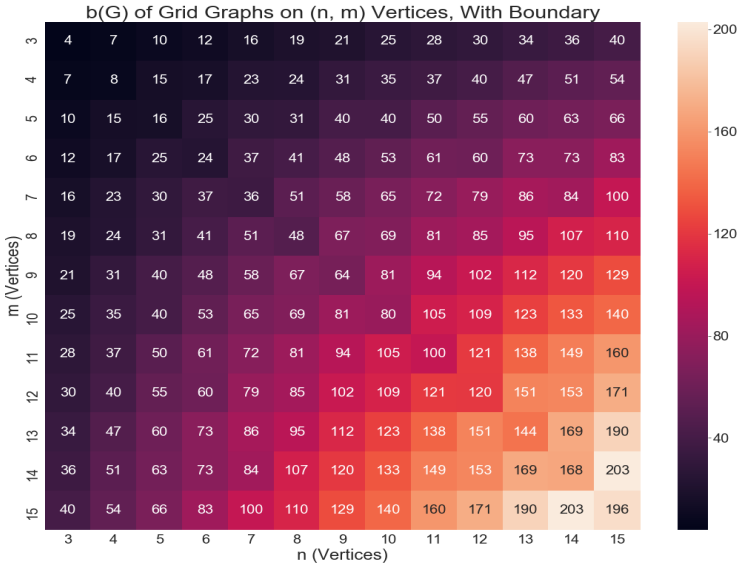

(b) Grid graphs with boundary. Entry (n, m) corresponds to $Gr'_{n,m} = C_n \times C_m$.

Figure 1: Each table shows $b(G)$ for the grid graphs. Bottom: Grid graphs without boundary. Top: Grid graphs with boundary. We conjecture that the bounds in 4.4.11 are tight for the grid graph without boundary when $n = m$ (see the main diagonal on the top). We also conjecture that the bounds in 4.4.11 are tight for the grid graph with boundary when $n = m$ and both are even (see the main diagonal on the bottom). The counts were computed by finding the Smith Normal Form of each graph Laplacian matrix, using Mathematica and Python.

Finally, we turn to the trivial counts for expander graphs. Recall that in 2.3, we presented two families of expander graphs. The chordal cycle graphs are 3-regular and have a prime number of vertices, while the the Margulis, Gabber Galil graphs (denoted MGG graphs) are 8-regular and have $n \cdot m$ vertices for positive integers n, m.

Both graphs exhibit a rapid, exponential growth in the value of their largest Smith invariant factor - see Figure 2. Recall that the product of the factors gives the tree number of the graph, so we would expect the *product* of the Smith Factors to grow rapidly with the number of vertices. However, this does not explain the fact that a *single* Smith factor makes so large a contribution to the order of the sandpile group.

A possible explanation comes from a recent work on the sandpile groups of Erdos-Renyi graphs [Woo17]. It shows that is much more likely for the Sylow p-subgroups of $G(n, p)$ to be cyclic than a product of smaller cyclic subgroups. For example, if 7 divides the order of the group, it is much more likely that $\mathbb{Z}/49\mathbb{Z}$ occurs in the invariant factor decomposition than $\mathbb{Z}/7\mathbb{Z} \times \mathbb{Z}/7\mathbb{Z}$. Given that expanders are known to look like random graphs, this route is worth exploring.

Conjecture 2: Let MGG_n be the Margulis, Gabber, Galil graph on n vertices, and CC_p be the chordal cycle graph on p vertices. If $v(G)$ denotes the value of the largest Smith invariant factor of $S(G)$, then as $n \to \infty$,

$$v(MGG_n) = O(2^n) \text{ and } v(CC_p) = O(2^p)$$

Finally, a complementary set of results concern the fraction of Smith factors equal to 1. If G is a graph on n vertices, this is just $\frac{b(G)}{n-1}$. We call this the trivial ratio.

Definition 4.4.12. Let G be a graph on n vertices. The **trivial ratio** of G is the fraction of its Smith invariant factors equal to 1, denoted by $r(G)$. So

$$r(G) = \frac{b(G)}{n-1}$$

If we are interested in asymptotic behavior of $b(G)$ as $n \to \infty$, then the trivial ratio is a more appropriate measure. In particular, we might ask whether the trivial ratio approaches 1 as $n \to \infty$. For certain families of graphs it seems this is the case, including the chordal cycle graphs - but not, interestingly, the MGG graphs - see Figure 3. Thus we suspect that the trivial ratio does not depend on expansion, since some expanders have trivial ratio approaching 1, while others do not.

Conjecture 3: Let MGG_n be the Margulis, Gabber, Galil graph on n vertices, and CC_p be the chordal cycle graph on p vertices. As $n, p \to \infty$

$$(n-1) - r(MGG_n) = O(n)$$
$$(p-1) - r(CC_p) = O(1)$$

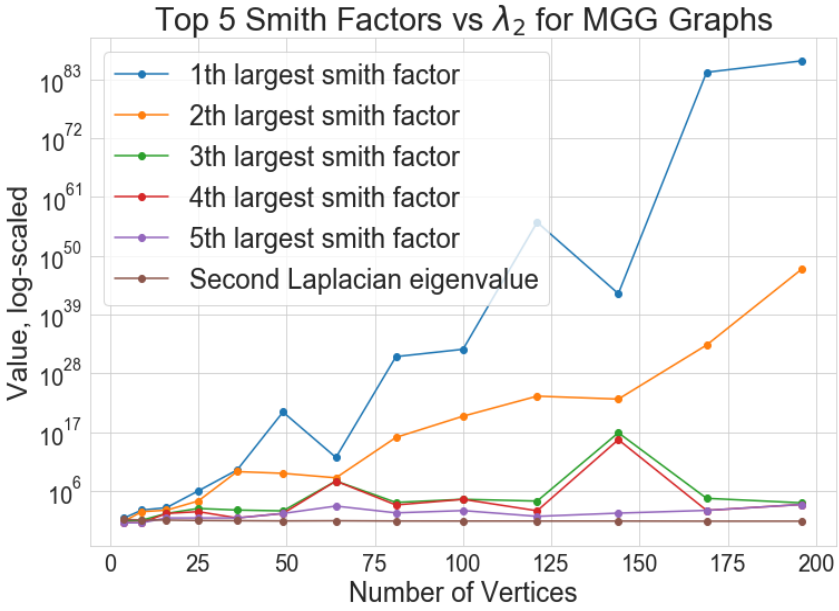

(a) MGG (Margulis, Gabber, Galil) graphs.

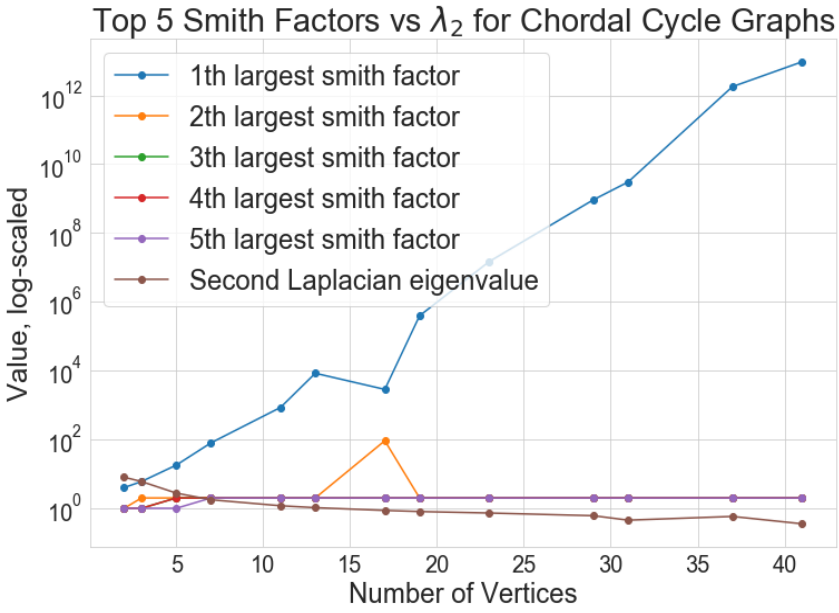

(b) The chordal cycle graphs.

Figure 2: Growth of the 5 largest Smith factors, as well as the second Laplacian eigenvalue, for two different expander families. The x-axis corresponds to the number of vertices for some member of the graph family, and the y-axis gives the value of the smith factor or eigenvalue, log-scaled. Notice that for both families, the largest Smith factor shows a rapid, exponential growth. For the MGG graphs (top), the second-largest smith factor seems to also grow quite rapidly, while the others seem stabilize in the limit. For the chordal cycle graphs (bottom), all but the largest Smith factor seem to stabilize in the limit. We include the second Laplacian eigenvalue as a proxy for the Cheeger constant of both graphs. Notice that λ_2 always stabilizes, showing that there cannot be a correlation between the Cheeger constant and the value of the largest Smith factor. The Smith invariant factor values were computed by finding the Smith Normal Form of each graph Laplacian matrix, using Mathematica and Python. Eigenvalues were computed in NumPy.

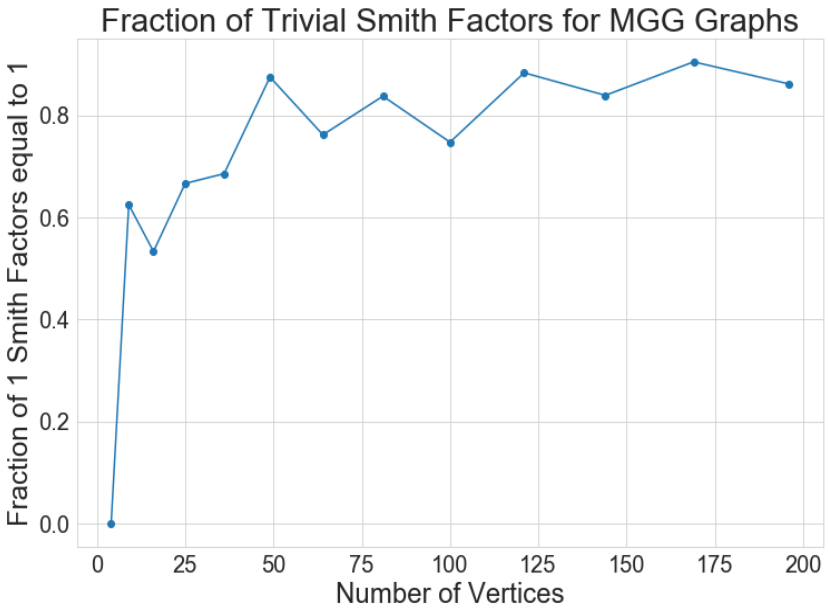

(a) MGG (Margulis, Gabber, Galil) graphs.

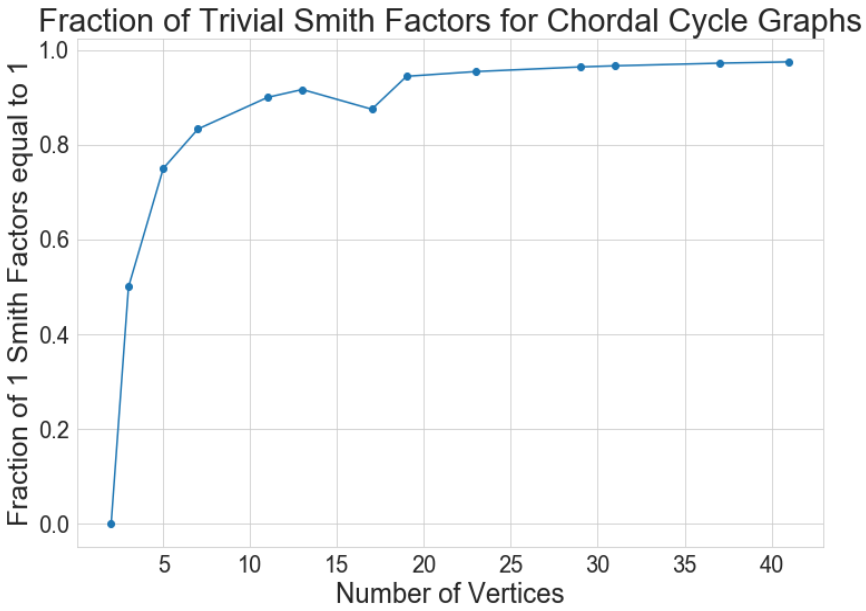

(b) The chordal cycle graphs.

Figure 3: The growth of $r(G)$ as number of vertices grows, for two different expander families. We conjecture $r(G)$ approaches 1 for the chordal cycle graphs (bottom) but not for the MGG graphs (top). The Smith invariant factor values were computed by finding the Smith Normal Form of each graph Laplacian matrix, using Mathematica and Python.

References

[ARV09] Arora, Sanjeev, Satish Rao, and Umesh Vazirani. "Expander flows, geometric embeddings and graph partitioning." Journal of the ACM (JACM) 56.2 (2009): 5.

[Bai03] Bai, Hua. "On the critical group of the n-cube." Linear algebra and its applications 369 (2003): 251-261.

[BTW87] Bak, Per, Chao Tang, and Kurt Wiesenfeld. "Self-organized criticality: An explanation of the 1/f noise." Physical review letters 59.4 (1987): 381.

[BN07] Baker, Matthew, and Serguei Norine. "Riemann–Roch and Abel–Jacobi theory on a finite graph." Advances in Mathematics 215.2 (2007): 766-788.

[Big97] Biggs, Norman. "Algebraic potential theory on graphs." Bulletin of the London Mathematical Society 29.6 (1997): 641-682.

[Big99] Biggs, Norman L. "Chip-firing and the critical group of a graph." Journal of Algebraic Combinatorics 9.1 (1999): 25-45.

[BL16] Bond, Benjamin, and Lionel Levine. "Abelian networks I. Foundations and examples." SIAM Journal on Discrete Mathematics 30.2 (2016): 856-874.

[CK78] Chaiken, Seth, and Daniel J. Kleitman. "Matrix tree theorems." Journal of combinatorial theory, Series A 24.3 (1978): 377-381.

[CC96] Chung, Fan RK. "Laplacians of graphs and Cheeger's inequalities." Combinatorics, Paul Erdos is Eighty 2.157-172 (1996): 13-2.

[CP18] Corry, Scott, and David Perkinson. Divisors and Sandpiles: An Introduction to Chip-Firing. Vol. 114. American Mathematical Soc., 2018.

[Dha90] Dhar, Deepak. "Self-organized critical state of sandpile automaton models." Physical Review Letters 64.14 (1990): 1613.

[DF04] Dummit, David Steven, and Richard M. Foote. Abstract algebra. Vol. 3. Hoboken: Wiley, 2004.

[FL16] Farrell, Matthew, and Lionel Levine. "CoEulerian graphs." Proceedings of the American Mathematical Society 144.7 (2016): 2847-2860.

[Hol08] Holroyd, Alexander E., et al. "Chip-firing and rotor-routing on directed graphs." In and Out of Equilibrium 2. Birkhäuser Basel, 2008. 331-364.

[HLW06] Hoory, Shlomo, Nathan Linial, and Avi Wigderson. "Expander graphs and their applications." Bulletin of the American Mathematical Society 43.4 (2006): 439-561.

[JNR03] Jacobson, Brian, Andrew Niedermaier, and Victor Reiner. "Critical groups for complete multipartite graphs and Cartesian products of complete graphs." Journal of Graph Theory 44.3 (2003): 231-250.

[JLP15] Jerison, Daniel C., Lionel Levine, and John Pike. "Mixing time and eigenvalues of the abelian sandpile Markov chain." arXiv preprint arXiv:1511.00666 (2015).

[Kli18] Klivans, Caroline J. The Mathematics of Chip-Firing. Chapman and Hall/CRC, 2018.

[Lev15] Levine, Lionel. "Threshold state and a conjecture of Poghosyan, Poghosyan, Priezzhev and Ruelle." Communications in Mathematical Physics 335.2 (2015): 1003-1017.

[Lor08] Lorenzini, Dino. "Smith normal form and Laplacians." Journal of combinatorial theory, series b 98.6 (2008): 1271-1300.

[RS17] Ramachandran, Akshay, and Aaron Schild. "Sandpile prediction on a tree in near linear time." Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms. Society for Industrial and Applied Mathematics, 2017.

[RMW93] Grone, Robert, Russell Merris, and William Watkins. "Laplacian unimodular equivalence of graphs." Combinatorial and graph-theoretical problems in linear algebra. Springer, New York, NY, 1993. 175-180.

[Woo17] Wood, Melanie. "The distribution of sandpile groups of random graphs." Journal of the American Mathematical Society 30.4 (2017): 915-958.