Algebraic Invariants of Sandpile Graphs

A Thesis

Presented to

The Division of Mathematics and Natural Sciences

Reed College

In Partial Fulfillment

of the Requirements for the Degree

Bachelor of Arts

John S. Wilmes

May 2010

Approved for the Division
(Mathematics)

_____

David Perkinson

Acknowledgements

I am grateful to Tian Jiang and Josh Hoak for helpful suggestions and comments. I am deeply grateful to Dave Perkinson for suggesting profitable lines of research, offering crucial insight in the proofs of several theorems, guiding my disorganized thoughts into convincing arguments, and in every other respect serving as an enviable thesis adviser. I am profoundly grateful to my parents for support and inspiration, without which this thesis would never have begun.

Table of Contents

Abstract

In this thesis, we define the toppling Betti numbers and minimal free resolutions of toppling ideals. We give a proof of the Riemann-Roch Theorem first stated and proved in (1), and use the theorem to prove that the last Betti number of an undirected graph is the number of minimal recurrent configurations. We describe a minimal complex that we conjecture to be a free resolution for a toppling ideal, and use the conjecture to compute the Betti numbers of several graphs.

Introduction

The abelian sandpile model was originally proposed as a computationally-tractable model of self-organized criticality (3). Since then, several authors have applied the model in studying algebraic graph theory, for example in (8) and (1). The goal of this thesis is to use the sandpile model to make new connections between algebra and graph theory by developing new algebraic graph invariants and studying their connection with graph structure.

The applications of this thesis will be twofold. First, by creating new algebraic graph invariants based on the sandpile model, we may better understand the combinatorial structure of graphs. The flow of sand through a multigraph according the sandpile model intuitively reflects classical ideas from graph theory, such as circulations, flows, and cuts. It is natural to hope that algebraic constructions built on the model might reflect graph structure in useful ways. And in fact, Conjecture 3.28 describes the connected partitions of a graph via the minimal free resolution of the toppling ideal.

The second application of this thesis is to algebra. It was shown in (11) that any lattice ideal defining a finite set of points is the lattice ideal of some sandpile graph. Hence, by giving a graph-theoretic description of toppling ideals, we provide new combinatorial tools for studying general algebraic objects. Again, Conjecture 3.28 provides a simple (though still computationally intractable) combinatorial method for constructing a minimal free resolution of certain ideals. The conjecture also provides immediate access to a minimal generating set for the ideal.

Chapter 1 of this thesis reviews the basic ideas of the sandpile model, along with some of the notation and more recent developments of (1). In Chapter 2, we prove the graph-theoretic version of the Riemann-Roch Theorem first proposed and proved in (1). Our proof is somewhat different, but ultimately founded on the same ideas, and a central goal of the chapter is to show how a number of variously defined objects interrelate. This is accomplished in Theorems 2.7 and 2.13. The main goal of the chapter, however, is simply to develop the foundation for Theorem 3.10. Chapter 3 defines new algebraic graph invariants, in particular the toppling Betti numbers and minimal free resolution of the quotient of the toppling ideal. These invariants were first defined in (11). In Theorem 3.10, we prove that the last Betti number of an undirected graph is the number of minimal recurrent configurations on the graph. Finally, we define a sequence of mappings of free modules built from the connected partitions of a graph and propose in Conjecture 3.28 that this sequence is in fact a minimal free resolution for the quotient of the toppling ideal. We provide a partial

proof, showing that it is at least a minimal complex. The final chapter supposes Conjecture 3.28 holds and uses it to compute minimal free resolutions for trees and complete graphs, and proves a few straightforward corollaries about how the structure of graphs relates to the Betti numbers.

In summary, this thesis attempts to build new bridges between algebra and combinatorics, in the hope that algebraists and graph theorists alike will traverse them.

Chapter 1

Preliminaries

Dhar's abelian sandpile model, first introduced in (3), is the foundation for the algebraic graph theory studied in this thesis. In this chapter, we define the basic ideas of the sandpile model and state a number of well-known results that we will use throughout the thesis. Some familiarity with basic graph theory is assumed, but the reader is referred to Appendix A for a brief review of the fundamentals.

Assumption. Γ is an Eulerian directed multigraph without loops on the vertex set $\{1, \ldots n\}$.

The group $\mathrm{Div}(\Gamma)$ of **divisors** on Γ is the free abelian group on the vertices of the graph. A divisor $D \in \mathrm{Div}(\Gamma)$ is generally written as a formal sum $D = \sum_{v \in V(\Gamma)} a_v v$ for some integers a_v. For such a divisor, we define $D_v = a_v$, so for any $D \in \mathrm{Div}(\Gamma)$, we have $D = \sum_{v \in V(\Gamma)} D_v v$. The group $\mathrm{Script}(\Gamma)$ of **scripts** on Γ is also the free abelian group on the vertices of the graph, with similar notation. We denote the identity element of either group by $\vec{0}$, and define $\vec{1} = \sum_{v \in V(\Gamma)} v$. The distinction between scripts and divisors is made so that we may define the dualized **Laplacian** map $\Delta^t : \mathrm{Script}(\Gamma) \to \mathrm{Div}(\Gamma)$ by

$$\Delta^t \sigma = \sum_{uv \in E(\Gamma)} \sigma_u (u - v)$$

where $\sigma \in \mathrm{Script}(\Gamma)$. In particular,

$$(\Delta^t \sigma)_v = \sigma_v \operatorname{outdeg}(v) - \sum_{vw \in E(\Gamma)} \sigma_w.$$

The divisor D' obtained by **firing** the script σ from the divisor D is

$$D' = D - \Delta^t \sigma.$$

We also say that σ takes D to D'. The following proposition is straightforward.

Proposition 1.1. $\mathrm{Script}(\Gamma)$ *acts on* $\mathrm{Div}(\Gamma)$ *by firing, i.e. by the group action*

$$\sigma \cdot D = D - \Delta^t \sigma.$$

The dualized Laplacian Δ^t is, as one might expect, the dual of the usual Laplacian operator on graphs, as defined in e.g. (6). It is straightforward to verify that $\ker \Delta^t = \mathbb{Z}\vec{1}$. (Certainly $\mathbb{Z}\vec{1} \subset \ker \Delta^t$, but by the matrix-tree theorem, rank $\Delta^t = n - 1$.) This property does not hold in general for non-Eulerian graphs, which is one reason we prefer Eulerian graphs in this thesis.

One intuitive understanding of divisors, scripts, and the firing action involves playing a money-trading game among a group of acquaintances. Each person is identified with a vertex of the graph, with (one-way) friendships corresponding to the edges of the graph in the natural way. We think of a divisor D as assigning to each person a certain amount of money or debt, so that person v has D_v dollars. A script is then a recipe for sharing money within the network. In particular, for a given vertex v, the script $\sigma = v$ causes person v to give one dollar to each of his friends. (When person v is doubly friendly with person w, i.e. when vw appears in $E(\Gamma)$ with multiplicity two, σ causes v to give two dollars to w, etc.) If person v does not have enough money for such a gift, rather than provoking jealousy among his friends, σ compels v to go into debt. Following (3), we will use a different analogy: a divisor corresponds to a certain amount of sand at each vertex of the graph, and firing a script causes a corresponding number of sand grains to be sent along each edge to an adjacent vertex. Though we prefer this analogy for historical reasons, the notion of a negative quantity of sand is less compelling.

For $D \in \operatorname{Div}(\Gamma)$, define the **support** of D by

$$\operatorname{supp}(D) = \{v \in V(\Gamma) : D_v \neq 0\}.$$

We have a partial ordering $\leq$ on $\operatorname{Div}(\Gamma)$ defined component-wise, i.e. $D \leq E$ when $D_v \leq E_v$ for every $v \in V(\Gamma)$. If $D \geq 0$, we say that D is **effective**. The support and effectiveness of scripts are defined identically. We say that two divisors D and E are **linearly equivalent**, and write $D \sim E$, if there exists a script σ taking D to E. The **linear system** $|D|$ of D is defined by

$$|D| = \{E \in \operatorname{Div}(\Gamma) : E \sim D \text{ and } E \geq 0\}.$$

In other words, $|D|$ is the set of effective divisors linearly equivalent to D. We define the **class group** $\operatorname{Cl}(\Gamma)$ of Γ as $\operatorname{Div}(\Gamma)$ modulo linear equivalence. Often, we will identify $[D] \in \operatorname{Cl}(\Gamma)$ with a representative element D. The **degree** $\deg(D)$ of a divisor D is defined as

$$\deg(D) = \sum_{v \in V(\Gamma)} D_v.$$

For a vertex set $U \subseteq V(\Gamma)$, define the divisor or script

$$\chi_U = \sum_{u \in U} u.$$

We will sometimes identify U with χ_U.

For $D \in \operatorname{Div}(\Gamma)$, a script σ is **legal** from D if no vertex v with D_v nonnegative has $(\sigma \cdot D)_v$ negative, i.e. no $\sigma \cdot D$ is nonnegative on $\operatorname{supp}(\sigma)$. If $S = \sigma_1, \ldots, \sigma_k$ is a

sequence of scripts, it is **legal** from D if for $1 \leq i \leq k$, the script σ_i is legal from the divisor $D_{i-1} = \left(\sum_{j=1}^{i-1} \sigma_j \right) \cdot D$ obtained by firing the first $i-1$ scripts in S from D. Often, S will be a sequence of vertices, or vertex subset of Γ, in which case we will consider it to be a sequence of scripts in the natural way. We will sometimes refer to such S as legal vertex firing sequences or legal set firing sequences. Observe that $v \in V(\Gamma)$ is legal as a script from D if $D_v \geq \mathrm{outdeg}(v)$.

Given a legal vertex firing sequence $v_1, \ldots, v_k$ from $D \in \mathrm{Div}(\Gamma)$, note that the script $\sigma = \sum_{i=1}^{k} v_k$ is also legal from D. However, the converse does not hold. Given a script σ legal from D, there is in general no legal vertex firing sequence $v_1, \ldots, v_k$ such that $\sigma = \sum_{i=1}^{k} v_k$. As a counterexample, consider the graph K_3 with one grain of sand on each of two vertices, and zero grains on the remaining vertex. No vertex may be legally fired on its own, but it is possible to fire the support of the divisor.

Remark 1.2. For any linearly equivalent divisors D and E, there exists a unique **minimal effective script** taking D to E, that is, a script $\sigma \geq 0$ such that for any other $\sigma' \geq 0$ taking D to E, $\sigma \leq \sigma'$. Certainly for any $D \sim E$, there is some script σ, possibly not effective, taking D to E. But $\vec{1} \in \ker \Delta^t$, so $\sigma + k\vec{1}$ also takes D to E for all $k \in \mathbb{Z}$. Let k be minimal so that $\sigma + k\vec{1} \geq 0$. Since in fact $\ker \Delta^t = \mathbb{Z}\vec{1}$, we have that $\sigma + k\vec{1}$ is the unique minimal effective script taking D to E. For any minimal effective script σ, there is some vertex v such that $\sigma_v = 0$. Minimal effective scripts are often easier to work with than arbitrary scripts, so we will use them extensively.

Although legal firing scripts are convenient tools, they are functionally equivalent to legal set firing sequences. The following is noted in (6).

Proposition 1.3. *Suppose an effective script σ is legal from $D \in \mathrm{Div}(\Gamma)$. Let $k = \max\{\sigma_v : v \in V(\Gamma)\}$, and for $1 \leq i \leq k$ define $V_i = \{v \in V(\Gamma) : \sigma_v \geq i\}$. Then $S = V_k, V_{k-1}, \ldots, V_1$ is a legal set firing sequence from D such that firing S results in $\sigma \cdot D$.*

A **sink** for Γ is an arbitrary fixed vertex s. Given a sink s, we define the set of non-sink vertices $\tilde{V}(\Gamma) = V(\Gamma) \setminus \{s\}$. A configuration c on Γ is a restriction of a divisor D to the non-sink vertices, $c = D|_{\tilde{V}(\Gamma)} = \sum_{v \in \tilde{V}(\Gamma)} D_v v$. A vertex $v \in V(\Gamma)$ is said to be **unstable** for a divisor D if $D_v \geq \mathrm{outdeg}(v)$, i.e. if v is a legal script from D; a divisor D is stable if every vertex is stable for D, and otherwise D is unstable. A divisor (or a configuration) is **stabilized** relative to s by firing the unstable vertices in $\tilde{V}(\Gamma)$ until a stable divisor is obtained. The stabilization process behaves well, according to Lemmas 2.4 and 2.5 of (6):

Theorem 1.4. *The stabilization process always terminates. Given a divisor D and a sink s, the divisor D° obtained by stabilizing D relative to s is independent of the choice of unstable vertex firings.*

The **maximal stable divisor** $D_{\max}$ is the divisor

$$D_{\max} = \sum_{v \in V(\Gamma)} (\mathrm{outdeg}(v) - 1)v,$$

so called because it is stable and $D_{\max} \geq E$ for any stable divisor E. Given a sink s, the **maximal stable configuration** $c_{\max}$ is the corresponding configuration, $c_{\max} = D_{\max}|_{\tilde{V}(\Gamma)}$. A configuration is **recurrent** if it is obtained by stabilizing a configuration of the form $c_{\max} + c$ for some effective configuration c. The following theorem is straightforward, and is entailed by Lemmas 2.13 and 2.15 of (6).

Theorem 1.5. *Given a sink s, every divisor class $[D] \in \mathrm{Cl}(\Gamma)$ contains a unique divisor E such that $E|_{\tilde{V}(\Gamma)}$ is recurrent.*

The following is given as Theorem 2.23 of (11), but a version appears as early as (3).

Theorem 1.6. *For $D \in \mathrm{Div}(\Gamma)$ and a sink s, define the script $\sigma = s$. Then the following are equivalent:*

(1) $D|_{\tilde{V}(\Gamma)}$ is recurrent.

(2) The stabilization of $\sigma \cdot D$ is D.

(3) In stabilizing $\sigma \cdot D$, every non-sink vertex fires exactly once.

Theorem 1.7 (Corollary 2.16 (6))**.** *Given a sink s, the set of recurrent configurations on Γ form a group under the operation of addition and subsequent stabilization.*

The group defined in Theorem 1.7 is called the **sandile group** of Γ, and is independent up to isomorphism of the choice of sink. It is the best known algebraic graph invariant arising from the sandpile model, but it will not be studied in this thesis.

Since stabilization is the process of firing a maximal sequence of legal vertices from a configuration, it is natural to ask what happens if we instead fire legal sets of non-sink vertices until no legal set firings remain. This process is called **superstabilization**, and similar results hold for superstabilization as for ordinary stabilization. For example, by Corollary 4.6 of (6), we have the following theorem.

Theorem 1.8. *Given a sink s and a divisor D, the superstabilization process terminates with a divisor $D_\circ$ independent of the choice of set firings.*

We say a configuration is **superstable** if it is obtained by superstabilization of an effective configuration. Theorem 4.4 of (6) states that the superstable configurations are dual to the recurrent configurations:

Theorem 1.9. *A configuration c is superstable iff $c_{\max} - c$ is recurrent.*

Chapter 2

The Riemann-Roch Theorem for Graphs

In this chapter, we develop the tools needed to prove the Riemann-Roch Theorem for graphs given by Baker and Norine in (1). Although our proof appears rather different from that given by Baker and Norine, the underlying ideas are identical. In fact, the divisor ν_P for a linear order $<_P$ of $V(\Gamma)$ in that paper is exactly the divisor $D_{\max} - D^P$ defined in this chapter. We attempt to connect these ideas, and others, in Lemma 2.7 and Theorem 2.13. Our methods, however, are slightly more general than those appearing in (1) so that they can be of greater service in subsequent chapters.

If σ is a script, then D is **effective relative to** σ if D is effective and σ is a legal firing from D. Similarly, for $S = \sigma_1, \dots, \sigma_k$ a sequence of scripts on Γ, a divisor D on Γ is **effective relative to** S if D is effective and S is a legal firing sequence from D. A divisor D is a **minimal effective divisor** relative to a sequence S of scripts if it is an effective divisor relative to S, and for any other divisor E effective relative to S, $E \nleq D$.

Proposition 2.1. *Let $S = \sigma_1, \dots, \sigma_k$ be a sequence of scripts on Γ. Then there exists a unique minimal effective divisor relative to S.*

Proof. Define $D_0 = \vec{0}$, and for $1 \leq i \leq k$, define D_i by

$$(D_i)_v = \max\{(D_{i-1})_v, (\Delta^t \sigma_i)_v\} - (\Delta^t \sigma_i)_v.$$

Thus, D_i is obtained from D_{i-1} by adding the minimum amount of sand such that σ_i can be legally fired, and then firing σ_i. Finally, fire in reverse every script in S from D_k, obtaining $D = D_k + \Delta^t \sum_{1 \leq i \leq k} \sigma_i$. We claim that D is effective relative to S. Certainly the result holds when $k = 0$, so suppose for induction that the proposition holds for all sequences of scripts of length $k-1$ with $k \geq 1$. Then, $D' = D_{k-1} + \Delta^t \sum_{1 \leq i \leq k-1} \sigma_i$

is effective relative to $S' = \sigma_1, \ldots, \sigma_{k-1}$. Since $D_k + \Delta^t \sigma_k \geq D_{k-1}$, we have

$$D = D_k + \Delta^t \sum_{1 \leq i \leq k} \sigma_i$$
$$\geq D_{k-1} + \Delta^t \sum_{1 \leq i \leq k-1} \sigma_i$$
$$= D'$$

and so D is also effective relative to S'. Firing S' from D, we obtain $D_k + \Delta^t \sigma_k$, whence we may legally fire σ_k. Thus, D is indeed effective relative to S.

Finally, observe that if E is effective relative to S, we may define effective divisors $E_0 = E$ and $E_i = E_{i-1} - \Delta^t \sigma_i$ for all $1 \leq i \leq k$, so that $(E_{i-1})_v \geq (\Delta^t \sigma_i)_v$. Now $E_0 \geq D_0$, and supposing that $E_{i-1} \geq D_{i-1}$ for $i < k$, we have

$$(E_i)_v = (E_{i-1})_v - (\Delta^t \sigma_i)_v$$
$$= \max\{(E_{i-1})_v, (\Delta^t \sigma_i)_v\} - (\Delta^t \sigma_i)_v$$
$$\geq \max\{(D_{i-1})_v, (\Delta^t \sigma_i)_v\} - (\Delta^t \sigma_i)_v$$
$$= (D_i)_v$$

Hence, by induction on i, $E_k \geq D_k$, and so $E \geq D$. It follows that D is the unique minimal effective divisor relative to S. □

Henceforth, for $S = \sigma_1, \ldots, \sigma_k$ a sequence of scripts, we will denote the minimal effective divisor relative to S by D^S. Occasionally, when S is equally a sequence of scripts for a graph Γ' as for Γ, we will use the notation D_Γ^S to specify the minimal effective divisor on Γ relative to S. An explicit formula for a special case of D^S is given as Lemma 2.4.

Lemma 2.2. *Let $S = \sigma_1, \ldots, \sigma_k$ be a sequence of scripts such that $\sum_{\sigma \in S} \sigma = \ell \vec{1}$ for some integer ℓ. Then for all $1 \leq i \leq k$ the divisor*

$$D_i = \left(\sum_{j \leq i} \sigma_j \right) \cdot D^S$$

obtained by firing the first i scripts of S from D^S is the minimal effective divisor relative to $S_i = \sigma_{i+1}, \ldots, \sigma_k, \sigma_1, \ldots, \sigma_i$.

Proof. Certainly D_i is effective relative to $S' = \sigma_{i+1}, \ldots, \sigma_k$. Observe that by firing S' from D_i, we obtain D^S once again since $\ell \vec{1} \in \ker \Delta^t$. Since D^S is effective relative to $\sigma_1, \ldots, \sigma_i$, it follows that D_i is effective relative to S_i. Now for any divisor E effective

relative to S_i, we have that $(\sum_{j>i} \sigma_j) \cdot E$ is effective relative to S. It follows that

$$
\begin{aligned}
E &\geq \left(-\sum_{j>i} \sigma_j \right) \cdot D^S \\
&= \left(-\vec{1} + \sum_{j \leq i} \sigma_j \right) \cdot D^S \\
&= \left(\sum_{j \leq i} \sigma_j \right) \cdot D^S \\
&= D_i
\end{aligned}
$$

and so D_i is the minimal effective divisor relative to S_i. $\qquad\square$

Definition 2.3 (k-partition). A k-partition of Γ is a partition of $V(\Gamma)$ into k pairwise-disjoint nonempty subsets $V_1, \ldots, V_k$. The k-partition is strongly connected if the induced subgraph on V_i is strongly connected for all i.

By $\mathcal{P}_k(\Gamma)$ we will denote the set of strongly connected k-partitions of Γ. Often it will be useful to order the sets of a k-partition, so by $\mathcal{S}_k(\Gamma)$ we will denote the set of all strongly connected k-partitions with orderings, i.e.

$$
\mathcal{S}_k(\Gamma) = \{V_{\tau(1)}, \ldots, V_{\tau(k)} : \{V_1, \ldots, V_k\} \in \mathcal{P}_k(\Gamma), \tau \in S_k\}.
$$

Thus, every element of $\mathcal{S}_k(\Gamma)$ is also an element of $\mathcal{P}_k(\Gamma)$ in the natural way. We now briefly pause our progress toward the Riemann-Roch theorem to give the following convenient formula.

Lemma 2.4. *Let Γ be undirected and fix $S = V_1, \ldots, V_k \in \mathcal{S}_k(\Gamma)$. Then for $v \in V_i$,*

$$
D_v^S = \sum_{i < j \leq k} \mathrm{wt}(v, V_j).
$$

Proof. Define D by

$$
D_v = \sum_{i < j \leq k} \mathrm{wt}(v, V_j)
$$

so that for $v \in V_1$,

$$
\begin{aligned}
(\chi_{V_1} \cdot D)_v &= \sum_{1 < j \leq k} \mathrm{wt}(v, V_j) - \left((\chi_{V_1})_v d_v - \sum_{vw \in E(\Gamma)} (\chi_{V_1})_w \right) \\
&= \mathrm{wt}(v, V(\Gamma) \setminus V_1) - \left(d_v - \sum_{\substack{vw \in E(\Gamma) \\ w \in V_1}} 1 \right) \\
&= \mathrm{wt}(v, V(\Gamma) \setminus V_1) - \mathrm{wt}(v, V(\Gamma) \setminus V_1) \\
&= 0.
\end{aligned}
$$

It follows that χ_{V_1} is legal from D, but χ_{V_1} is not legal from $D - v$ for $v \in V_1$.

Now let $U_i = V_{i+1}$ for $1 \leq i \leq k - 1$ and $U_k = V_1$. Write $S' = U_1, \ldots, U_k$, so S' is a rotation of S. Define E by

$$E_v = \sum_{i < j \leq k} \mathrm{wt}(v, U_j).$$

For $v \in V_i$ with $i \neq 1$ we have

$$
\begin{aligned}
(\chi_{V_1} \cdot D)_v &= \sum_{i < j \leq k} \mathrm{wt}(v, V_j) - \left((\chi_{V_1})_v d_v - \sum_{vw \in E(\Gamma)} (\chi_{V_1})_w \right) \\
&= \sum_{i < j \leq k} \mathrm{wt}(v, V_j) + \sum_{\substack{vw \in E(\Gamma) \\ w \in V_1}} 1 \\
&= \sum_{i < j \leq k} \mathrm{wt}(v, V_j) + \mathrm{wt}(v, V_1) \\
&= E_v,
\end{aligned}
$$

and certainly by the previous display, $(\chi_{V_1} \cdot D)_v = 0 = E_v$ for $v \in V_1$. The proposition follows by induction and Lemma 2.2. $\qquad\square$

For $S = V_1, \ldots, V_k \in \mathcal{S}_k(\Gamma)$, we have $\mathrm{supp}(\chi_{V_1} \cdot D^S) \cap V_1 = \varnothing$ by the minimality of D^S. Hence, by Lemma 2.2, the divisor obtained by firing the first ℓ sets of S from D^S is not supported on V_ℓ. In fact, we have the following theorem.

Theorem 2.5. *Fix $S = V_1, \ldots, V_k \in \mathcal{S}_k(\Gamma)$. Then for all $E \in |D^S|$, there exists $V_i \in S$ such that the support of E is disjoint from V_i.*

Proof. Suppose $E \in |D^S|$ is such that for all $V_i \in S$, we have $\mathrm{supp}(E) \cap V_i \neq \varnothing$. Let σ be the minimal effective script taking E to D^S, and let ℓ be minimal such that

$$\mathrm{supp}(\sigma) \cup \bigcup_{i \leq \ell} V_i = V(\Gamma),$$

so V_ℓ is the last set of vertices in S to completely fire in the sequence of firings $\sigma, \chi_{V_1}, \ldots, \chi_{V_k}$. Certainly $\ell > 0$ since σ does not have full support. Define

$$\sigma_\ell = \sigma + \sum_{1 \leq i \leq \ell} \chi_{V_i},$$

and note that $\sigma_\ell \geq \vec{1}$. We may then define $\sigma'_\ell = \sigma_\ell - \vec{1}$, so $\sigma'_\ell \geq 0$ acts identically to σ_ℓ on $\mathrm{Div}(\Gamma)$. Finally, define $D_\ell = \sigma_\ell \cdot E$, and note that

$$D_\ell = \left(\sum_{1 \leq i \leq \ell} \chi_{V_i} \right) \cdot D^S,$$

that is, D_ℓ is the divisor obtained by firing the first ℓ sets of S from D. Recall that D_ℓ is not supported on V_ℓ.

Fix $u \in V_\ell \cap \text{supp}(E)$. Since $D_\ell = \sigma'_\ell \cdot E$ is not supported on V_ℓ, u must fire when σ'_ℓ does. However, by the definition of ℓ, there is some $v \in V_\ell$ such that $(\sigma_\ell)_v = 1$, so that $(\sigma'_\ell)_v = 0$. Hence, since V_ℓ is strongly connected and a proper subset of V_ℓ fires when σ'_ℓ does, there are vertices $x, y \in V_\ell$ such that $xy \in E(\Gamma)$, $x \in \text{supp}(\sigma'_\ell)$, and $y \notin \text{supp}(\sigma'_\ell)$. Hence, y gains sand from x when σ'_ℓ fires, so D_ℓ is supported on V_ℓ. We have a contradiction, and so conclude that no such E exists. $\qquad\square$

In particular, if a D is a minimal effective divisor relative to an ordering of the vertices of Γ, then no divisor in $|D|$ has full support.

Definition 2.6 (Alive). A divisor $D \in \text{Div}(\Gamma)$ is alive if $\sigma \cdot D$ is unstable for all scripts σ. A divisor $D \in \text{Div}(\Gamma)$ is minimally alive (or, is a minimal alive divisor), if it is alive but $D - v$ is not alive for all $v \in V(\Gamma)$.

The notion of an alive divisor is equivalent to that of being effective relative to an ordering of $V(\Gamma)$, as will be proved in Lemma 2.7. However, alive divisors more naturally relate to the historically studied question of whether a divisor ever stabilizes, and so we include its definition and the following results for completeness.

Lemma 2.7. *Let Γ have sink s and fix $D \in \text{Div}(\Gamma)$. Then the following are equivalent:*

(1) D is alive.

(2) There exists a recurrent configuration c and an integer $k \geq d_s$ such that $D \sim c + ks$.

(3) There exists an ordering S of $V(\Gamma)$ and $E \sim D$ such that E is effective relative to S.

Proof. [(1) $\implies$ (2)]: Suppose D is alive. Stabilize $D|_{\tilde{V}(\Gamma)}$ to obtain a divisor D'. Now s is unstable, i.e. $D'_s \geq d_s$, so we may fire s from D' and then stabilize. If every non-sink vertex fires in the stabilization, then $D'|_{\tilde{V}(\Gamma)}$ is recurrent by Theorem 1.6. Otherwise, s is again unstable since D is alive, and we may repeat the process. By Theorem 1.6, we obtain the desired result.

[(2) $\implies$ (3)]: Suppose $D \sim E = ks + c$ for some recurrent configuration c and $k \geq d_s$. By Theorem 1.6, after firing s from E, we may legally fire every other vertex in turn, so E is effective relative to some vertex ordering $v_1, \ldots, v_n$ with $s = v_1$.

[(3) $\implies$ (1)]: Suppose $E \sim D$ is effective relative to some ordering $S = v_1, \ldots, v_n$ of $V(\Gamma)$. For any $F \sim D$, let σ be the minimal effective script taking E to F. Let k be minimal such that $\sigma_{v_k} = 0$. Then v_k can fire from F, so F is not stable, and thus D is alive. $\qquad\square$

Not only do divisors effective relative to vertex orderings, alive divisors, and recurrent configurations coincide with each other, but their corresponding minimality conditions also coincide. That is, D is minimally alive iff there exists an ordering S

of $V(\Gamma)$ such that $D \sim D^S$, iff there exists a minimal recurrent configuration c such that $D \sim c + d_s s$. Another consequence of the lemma is that for $D \in \operatorname{Div}(\Gamma)$ alive, there exists $E \sim D$ with E effective.

The following Lemma 2.8 does not hold for Eulerian graphs generally. As a counterexample, consider a directed cycle on three or more vertices. Since Theorem 2.15 entails the lemma, that theorem cannot hold in the directed case.

Assumption. Henceforth, Γ is taken to be undirected.

Lemma 2.8. *Let $D \in \operatorname{Div}(\Gamma)$ be minimally alive. Then $\deg(D) = |E(\Gamma)|$.*

Proof. Let D be minimally alive, so that without loss of generality by Lemma 2.7, we have $D = D^S$ for some ordering $S = v_1, \ldots, v_n$ of $V(\Gamma)$. The result is immediate when $n = 1$, so suppose that it holds for all graphs on $n - 1$ vertices. Let Γ' be the induced subgraph of Γ on $V(\Gamma) \backslash \{v_1\}$—without loss of generality, Γ' is connected, since otherwise by Lemma 2.2 we could choose a different v_1. Observe that $D^S_{v_1} = d_{v_1}$, and furthermore that $D^S|_{\Gamma'}$ is the minimal effective divisor on Γ' relative to $S' = v_2, \ldots, v_n$. Thus,

$$\deg(D^S) = d_{v_1} + \deg(D^S|_{\Gamma'}) = d_{v_1} + |E(\Gamma')| = |E(\Gamma)|,$$

and the lemma follows by induction on n. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

The following is immediate from Lemmas 2.7 and 2.8:

Lemma 2.9. *Let c be a minimal recurrent configuration on Γ. Then*

$$\deg(c) = |E(\Gamma)| - d_s.$$

Hence, the minimal recurrent configurations are exactly the recurrent configurations of minimal degree.

Definition 2.10 (Genus)**.** For a graph Γ, the genus g of Γ is

$$g = |E(\Gamma)| - |V(\Gamma)| + 1.$$

Definition 2.11 (Canonical divisor)**.** The canonical divisor K of Γ is the divisor $D_{\max} - \vec{1}$.

Definition 2.12 (Non-special divisor)**.** For a graph Γ with genus g, define the set of non-special divisors as

$$\mathcal{N} = \{D \in \operatorname{Div}(\Gamma) : \deg(D) = g - 1 \text{ and } |D| = \varnothing\}.$$

Theorem 2.13. *Let Γ have sink s and fix $D \in \operatorname{Div}(\Gamma)$. Then the following are equivalent:*

(1) $D \in \mathcal{N}$

(2) $K - D \in \mathcal{N}$

(3) There exists a maximal superstable configuration c such that $D \sim c - s$

(4) $D_{\max} - D$ is minimally alive.

(5) $|D| = \varnothing$, and $|D + v| \neq \varnothing$ for all $v \in V(\Gamma)$.

Proof. (3) and (4) are equivalent by Lemmas 1.9 and 2.7.

$[(1) \implies (4)]$: Suppose $D \in \mathcal{N}$. Then for all $E \sim D$, there is some $v \in V(\Gamma)$ such that $E_v < 0$, and so $(D_{\max} - E)_v \geq d_v$, and $D_{\max} - D$ is alive. Since

$$\deg(D_{\max} - D) = 2|E(\Gamma)| - |V(\Gamma)| - g + 1 = |E(\Gamma)| - |V(\Gamma)| = |E(\Gamma)|,$$

$D_{\max} - D$ is minimally alive by Lemma 2.8.

$[(4) \implies (2)]$: Suppose $D_{\max} - D$ is minimally alive. Then for all $E \in D_{\max} - D$, by Lemma 2.7 and Theorem 2.5, E does not have full support, so

$$|K - D| = |D_{\max} - D - \vec{1}| = \varnothing.$$

Since $\deg(K - D) = |E| - |V|$, we have $K - D \in \mathcal{N}$.

$[(2) \implies (1)]$: By the above, we have $D \in \mathcal{N}$ implies that $K - D \in \mathcal{N}$. Hence, $K - D \in \mathcal{N}$ implies that $K - (K - D) \in \mathcal{N}$.

$[(4) \implies (5)]$: Let $D_{\max} - D$ be minimally alive, so $|D| = \varnothing$ since $(4) \implies (1)$. For any $v \in V(\Gamma)$, $D_{\max} - D - v$ is not alive, so there is a stable divisor $E \sim D_{\max} - D - v$, and so

$$|D + v| = |D_{\max} - (D_{\max} - D - v)| = |D_{\max} - E| \neq \varnothing.$$

$[(5) \implies (4)]$: Suppose D satisfies (5). Then $D_{\max} - D$ is alive by the argument in $(1) \implies (4)$. On the other hand, there is some $E \in |D+v|$, so $D_{\max} - E \sim D_{\max} - D - v$ is stable. Hence $D_{\max} - D$ is minimally alive. $\square$

Definition 2.14 ($r(D)$)**.** Define $r(D)$ for $D \in \mathrm{Div}(\Gamma)$ by letting $r(D) = -1$ when $|D| = \varnothing$, and otherwise

$$r(D) = \min\{k \in \mathbb{Z} : \exists E \in \mathrm{Div}(\Gamma) \text{ with } \deg(E) = k \text{ such that } |D - E| = \varnothing\} - 1.$$

In other words, $|D| \neq \varnothing$, $r(D) + 1$ is the minimal integer such that there exists $E \in \mathrm{Div}(\Gamma)$ with $\deg(E) = r(D) + 1$ such that $|D - E| = \varnothing$. Equivalently, $r(D)$ is the maximal integer such that for all $E \in \mathrm{Div}(\Gamma)$ with $\deg(D) = r(D)$, we have $|D - E| \neq \varnothing$.

Theorem 2.15 (Riemann-Roch for Graphs)**.** *For $D \in \mathrm{Div}(\Gamma)$,*

$$r(D) - r(K - D) = \deg(D) + 1 - g$$

The following is given as Theorem 2.2 of (1).

Theorem 2.16. *Theorem 2.15 holds iff the following two properties are satisfied:*

(RR1) For every $D \in \mathrm{Div}(\Gamma)$, there exists $\nu \in \mathcal{N}$ such that exactly one of $|D|$ and $|\nu - D|$ is empty.

(RR2) For every $D \in \mathrm{Div}(\Gamma)$ with $\deg(D) = g - 1$, $|D|$ and $|K - D|$ are either both empty or both non-empty.

Proof of Theorem 2.15. Let D be a divisor. If $|D| \neq \varnothing$, then $E \sim D$ is effective, and certainly $|\nu - D| = \varnothing$ since $|\nu| = \varnothing$. If $|D| = \varnothing$, fix a sink s. By Theorem 1.5 and Theorem 1.9, there exists $D' \sim D$ with $D'|_{\tilde{V}(\Gamma)}$ a superstable configuration. Let $c \geq D'|_{\tilde{V}(\Gamma)}$ be a maximal superstable configuration, and define $\nu = c - s$. By Theorem 2.13, we know that $\nu \in \mathcal{N}$, and $\nu - D' \geq 0$ since $D'_s < 0$. Thus, (RR1) holds. On the other hand, (RR2) simply restates the equivalence of (1) and (2) of Theorem 2.13. □

Chapter 3

Algebraic Graph Invariants

The fundamental graph invariant we study in this thesis is the homogeneous toppling ideal, defined shortly hereafter. The inhomogeneous version of the ideal was first defined in (2), and Proposition 4.8 of (11) states that our ideal is the homogenization (with respect to a sink indeterminate) of the original. Translating between the two ideals is generally straightforward. However, as we will see in Proposition 3.6, the quotient by the homogeneous ideal has the advantage of being a graded module and hence its minimal graded free resolution is unique up to isomorphism. This uniqueness allows us to define the Betti numbers of the graph, which are the ranks of the free modules of this minimal free resolution. We show in Theorem 3.10 that the last Betti number is the number of minimal recurrent configurations, and conjecture that the minimal free resolution can be written in terms of the connected partitions of Γ.

The homogeneous toppling ideal, its minimal free resolution, and its Betti numbers have all been defined previously in (11). (In fact, in that paper the free resolution and Betti numbers are defined even for digraphs.) What is novel here are their combinatorial characterizations. Note that Betti numbers have been defined for graphs previously in (7), but these correspond to a different ideal. Where confusion between the two sorts of Betti numbers is possible, we recommend referring to the Betti numbers of this thesis as the toppling Betti numbers of the graph.

Assumption. R is the ring $\mathbb{C}[x_1, \ldots, x_n]$.

For $D \in \mathrm{Div}(\Gamma)$, define
$$x^D = \prod_{v \in V(\Gamma)} x_v^{D_v}.$$

Definition 3.1 (Homogeneous toppling ideal)**.** The homogeneous toppling ideal $I(\Gamma)$ for Γ is

$$I(\Gamma) = \mathrm{span}_{\mathbb{C}} \{x^D - x^E : D, E \in \mathrm{Div}(\Gamma) \text{ such that } D \sim E \text{ and } D, E \geq \vec{0}\} \subset R.$$

Definition 3.2 (Grading)**.** A ring S is graded by an abelian group A if there are subgroups $S_a \subset S$ for $a \in A$ such that

$$S = \bigoplus_{a \in A} S_a$$

as groups, and $S_a S_b \subset S_{a+b}$ for $a, b \in A$. An element $f \in S$ is homogeneous of degree a if $f \in S_a$.

Proposition 3.3. *R is graded by* $\mathrm{Cl}(\Gamma)$.

Proof. Letting $R_D = \mathrm{span}_{\mathbb{C}} \{ x^E : E \in |D| \}$, we have the decomposition

$$R = \bigoplus_{D \in \mathrm{Cl}(\Gamma)} R_D,$$

and for $D_1, D_2 \in \mathrm{Div}(\Gamma)$, we have $x^{D_1} x^{D_2} = x^{D_1 + D_2}$. Thus, $R_{D_1} R_{D_2} \subset R_{D_1 + D_2}$. $\qquad \square$

Definition 3.4 (Twist). For an abelian group A with $a \in A$, the a^{th} twist of the A-graded ring S is the ring $S(a)$ equal to S but with grading given by $S(a)_b = S_{a+b}$ for all $b \in A$.

A twist of a ring simply shifts its grading.

Definition 3.5 (Graded module). Let A be an abelian group, and let S be a ring graded by A. A graded S-module is a module M with subgroups M_a for $a \in A$ such that

$$M = \bigoplus_{a \in A} M_a$$

as groups, and $S_a M_b \in M_{a+b}$ for $a, b \in A$. A graded free S-module is a module M with nonnegative integers β_a for $a \in A$ such that

$$M = \bigoplus_{a \in A} S(-a)^{\beta_a}.$$

(We use the twist $S(-a)$ instead of $S(a)$ when writing M because $S(-a)$ is generated by an element of S of degree a, i.e. an element of S_a.) A mapping $\phi : M \to N$ of graded free S-modules preserves degrees if $\phi(m)$ is homogeneous of degree a when m is.

Proposition 3.6. *The quotient $S/I(\Gamma)$ is a graded R-module.*

Proof. The proposition is obvious since $I(\Gamma)$ is a homogeneous ideal, which is to say it is generated by elements of the groups $R_{[D]}$. $\qquad \square$

Definition 3.7 (Exact sequence). A sequence of S-modules with maps

$$\cdots \longleftarrow F_{-1} \xleftarrow{\phi_0} F_0 \xleftarrow{\phi_1} F_1 \longleftarrow \cdots$$

is exact at F_0 if $\ker \phi_0 = \mathrm{im}\, \phi_1$. The sequence is exact if it is exact at each F_i.

Finally, we are equipped for the crucial definition of this chapter.

Definition 3.8 (Free resolution)**.** A free resolution $\mathcal{F}$ of an S-module M is a collection of free S-modules F_i for $i \geq 0$, with maps $\phi_i : F_i \to F_{i-1}$ for $i \geq 1$ forming the exact sequence

$$\mathcal{F}: \qquad F_0 \xleftarrow{\phi_1} F_1 \xleftarrow{\phi_2} F_2 \cdots \xleftarrow{\phi_{n-1}} F_{n-1} \longleftarrow 0$$

such that $M = \operatorname{coker} \phi_1$. We say $\mathcal{F}$ is a graded free resolution if the F_i are graded and the ϕ_i preserve degrees. When $S = R$, let $P = \langle x_1, \ldots, x_n \rangle \subset R$. The graded free resolution is minimal if $\operatorname{im} \phi_i \subset P F_{i-1}$ for all i.

Of course, minimal graded free resolutions can be defined in greater generality, but for the purposes of this thesis, there is little benefit to the added complexity that would be required. Note that minimal graded free resolutions of finitely-generated R-modules exist. Letting M be such a module, we can construct a minimal graded free resolution for M by building a free module F_0 from a direct sum of copies of R, with one copy for each element of a minimal generating set of M. We then repeat the process, replacing M with the kernel of the map $F_0 \to M$. Hilbert's syzygy theorem (1.13 of (5)) guarantees that the existence of a graded free resolution of length at most n for any finitely-generated R-module M, and so this process eventually stops. Finally, Theorem 20.2 of (5) states that the minimal free resolutions of M are isomorphic. The uniqueness of minimal graded free resolutions up to isomorphism justifies the following definition.

Definition 3.9 (Betti numbers)**.** Let

$$\mathcal{F}: \qquad F_0 \xleftarrow{\phi_1} F_1 \xleftarrow{\phi_1} F_2 \longleftarrow \cdots \xleftarrow{\phi_{n-1}} F_{n-1} \longleftarrow 0$$

be a minimal free resolution of M. If

$$F_i = \bigoplus_{D \in \operatorname{Cl}(\Gamma)} R(-D)^{\beta_{i,D}}$$

then $\beta_{i,D}(M) = \beta_{i,D}$ is the i^{th} Betti number of M in degree D. The i^{th} coarsely-graded Betti number $\beta_i(M)$ of M is

$$\beta_i(M) = \sum_{D \in \operatorname{Cl}(\Gamma)} \beta_{i,D}(M).$$

The Betti numbers of a module convey some of the same information as a free resolution in a much more compact form. Henceforth, we will define the k^{th} Betti number $\beta_{i,D}(\Gamma)$ of Γ in degree D by $\beta_{i,D}(\Gamma) = \beta_{i,D}(S/I(\Gamma))$.

A central goal of this thesis is to relate our new algebraic descriptions of the homogeneous toppling ideal to the graph and sandpile model whence it came. The first step is the following.

Theorem 3.10. *The $(n-1)^{th}$ coarsely-graded Betti number of Γ is the number of minimal recurrent configurations on Γ.*

A proof of the theorem will come shortly, but some additional machinery is needed. Although the Betti numbers of an ideal are a simple numeric invariant, they are nevertheless difficult to compute in general. Hochster's formula (Theorem 3.16) relates the Betti numbers of an ideal to the homology of associated simplicial complexes, and these homologies are sometimes more convenient tools. We now give the essential definitions of simplicial homology so that we can state Hochster's formula in advance of the proof of Theorem 3.10. Our reference for the following is the first chapter of (9).

Definition 3.11 (Simplicial complex). A simplicial complex on the linearly ordered set S is a collection Δ of subsets of S closed under taking subsets. A face σ is a set $\sigma \in \Delta$, and a facet is a maximal face. A k-face σ is a face with $|\sigma| = k + 1$.

Observe that a simplicial complex is uniquely determined by its facets. One usually thinks of a simplicial complex as a set of overlapping simplices embedded in $(|S| - 1)$-dimensional Euclidean space. In our case, the simplicial complexes related to $I(\Gamma)$ are determined by the supports of the divisors in linear systems. (We will see *how* the are related shortly.)

Definition 3.12 (Δ_D). For $D \in \mathrm{Div}(\Gamma)$, define Δ_D as the simplicial complex $V(\Gamma)$ whose facets are the sets in $\{\mathrm{supp}(E) : E \in |D|\}$.

Definition 3.13 (Boundary maps). For a simplicial complex Δ on S, define $F_i(\Delta)$ as the set of i-faces of Δ, and $\mathbb{C}^{F_i(\Delta)}$ as the $\mathbb{C}$-vector space with basis elements e_σ for $\sigma \in F_i(\Delta)$. For $i \geq 0$, the i^{th} boundary map $\partial_i : \mathbb{C}^{F_i(\Delta)} \to \mathbb{C}^{F_{i-1}(\Delta)}$ is defined by

$$\partial_i(e_\sigma) = \sum_{x \in \sigma} \mathrm{sign}(x, \sigma) e_{\sigma \setminus x}$$

where $\mathrm{sign}(x, \sigma) = (-1)^k$ for x the k^{th} element of σ according to the linear order on S.

Thus, given an i-face σ, the vector $\partial_i(\sigma)$ is essentially a formal sum of the faces bounding the i-dimensional simplex corresponding to σ.

Lemma 3.14. *For $i \geq 0$, we have $\partial_{i+1}\partial_i = 0$.*

The lemma is both standard and straightforward, so no proof will be given. It is necessary so that the following is well-defined.

Definition 3.15 (Reduced homology). For a simplicial complex Δ, the k^{th} reduced homology $\tilde{H}_k(\Delta)$ of Δ is

$$\tilde{H}_k(\Delta) = \ker \partial_k / \operatorname{im} \partial_{k+1}.$$

Intuitively, the i^{th} reduced homology of a simplicial complex measures the number of i-dimensional "holes" in the corresponding set of simplices embedded in Euclidean space. Finally, we arrive at the needed theorem. Our version of this theorem is given as Lemma 2.1 of (10).

Theorem 3.16 (Hochster's Formula).

$$\beta_{k,D}(\Gamma) = \dim_{\mathbb{C}} \tilde{H}_{k-1}(\Delta_D).$$

With Hochster's formula in hand, we are equipped to prove Theorem 3.10.

Proof of Theorem 3.10. We wish to find $\beta_{n-1}(\Gamma)$, so by Hochster's formula, we are interested in counting the dimension of $\tilde{H}_{n-2}(\Delta_D)$ as a $\mathbb{C}$-vector space for $D \in \mathrm{Cl}(\Gamma)$. Since Δ_D is a simplicial complex on n vertices, the dimension of $\tilde{H}_{n-2}(\Delta_D)$ is 1 when Δ_D is an empty n-simplex and 0 otherwise, and so we may instead count the number of divisor classes D with the following properties:

(i) No divisor $E \in |D|$ has full support, and

(ii) For every $v \in V(\Gamma)$, there exists $E_v \in |D|$ with $\mathrm{supp}(E_v) = V(\Gamma) \setminus \{v\}$.

Fix a sink s. Let c be a minimal recurrent configuration, so that by Theorem 1.6 for $D = c + d_s s$, there is an ordering S of $V(\Gamma)$ such that $D = D^S$. By Theorem 2.5, no $E \in |D|$ has full support, so we have (i). For (ii), observe E_v is characterized by $E_v - \vec{1} + v \geq \vec{0}$, so it suffices to show that $r(D - \vec{1} + v) \geq 0$. Define $F = D_{\max} - (D + v)$, so that by Lemma 2.8,

$$\deg(F) = 2|E(\Gamma)| - |V(\Gamma)| - |E(\Gamma)| - 1 = g - 2.$$

By Theorem 2.15 we have $r(F) - r(K - F) = -1$. Now $D_{\max} - (F + v) = D$ is minimally alive, so $F + v \in \mathcal{N}$ by the equivalence of (4) and (1) of Theorem 2.13. Thus, $r(F) = -1$, and

$$0 = r(K - F) = r(D_{\max} - \vec{1} - (D_{\max} - D - v)) = r(D - \vec{1} + v).$$

It follows that Δ_D has the desired homology.

To prove the other direction, let D be a divisor satisfying (i) and (ii) above. Then $r(D - \vec{1}) = -1$, but for every $v \in V(\Gamma)$, $r(D - \vec{1} + v) \geq 0$. Then by Theorem 2.13, we have $F = D - \vec{1} \in \mathcal{N}$, so $F \sim c - s$ where c is a maximal superstable configuration. Using Theorem 1.9, we have $D_{\max} - F = d_s s + \bar{c}$ for some minimal recurrent configuration $\bar{c}$.

Thus, the divisor class of D satisfies (i) and (ii) iff its linear system contains a divisor of the form $d_s s + c$ for c a minimal recurrent configuration. $\square$

The theorem now proved relates the minimum recurrent configurations on a graph to its last Betti number. Our goal for the remainder of this chapter is to translate that result to lower Betti numbers. To that end, we now define a graph determined by a connected partition of Γ, and we will see in Theorem 3.18 that the minimum alive divisors on these partition graphs are equivalent to the minimum effective divisors on Γ relative to orderings of the partition.

Definition 3.17 (Partition graph). Fix $P \in \mathcal{P}_k(\Gamma)$. Define the partition graph Γ_P to be the multigraph with vertex set P and an edge $V_i V_j \in E(\Gamma_P)$ with multiplicity $\mathrm{wt}_\Gamma(V_i, V_j)$ whenever there exists $u \in V_i, v \in V_j$ such that $u, v \in E(\Gamma)$. Partition graphs come with projections $\pi_P : \mathrm{Div}(\Gamma) \to \mathrm{Div}(\Gamma_P)$ defined by

$$(\pi_P(D))_{V_i} = \sum_{v \in V_i} D_v.$$

Theorem 3.18. *Fix $P \in \mathcal{P}_k(\Gamma)$ and let S be an ordering of P, so also S is an ordering of $V(\Gamma_P)$. Then $\pi_P(D^S)$ is the minimal effective divisor on Γ_P relative to S.*

Proof. Let $S = V_1, \ldots, V_k$ and let E be the minimal effective divisor on Γ_P relative to S. By Lemma 2.4,

$$
\begin{aligned}
\pi_P(D^S)_{V_i} &= \sum_{v \in V_i} \sum_{i < j \leq k} \mathrm{wt}(v, V_j) \\
&= \sum_{i < j \leq k} \sum_{v \in V_i} \mathrm{wt}(v, V_j) \\
&= \sum_{i < j \leq k} \mathrm{wt}(V_i, V_j) \\
&= E_{V_i}.
\end{aligned}
$$

$\square$

Definition 3.19 ($S \sim T$)**.** For orderings S_1 and S_2 of $P \in \mathcal{P}_k(\Gamma)$, say $S_1 \sim S_2$ if $D^{S_1} \sim D^{S_2}$.

Clearly the relation $\sim$ defined above gives an equivalence relation on $\mathcal{S}_k(\Gamma)$.

Theorem 3.20. *Let $S \sim T$, with $S, T \in \mathcal{S}_n(\Gamma)$. Then there is a sequence of vertex orderings $S = S_1, S_2, \ldots, S_\ell = T$ such that $S_i \sim S_{i-1}$ and S_i is obtained from S_{i-1} for $i \geq 2$ either by a rotation or by transposing two adjacent vertices $u, v \in S_{i-1}$, one of which is unstable in $D^{S_{i-1}}$.*

Proof. Fix a sink s. By Lemma 2.7, there is $D \sim D^S \sim D^T$ such that $D = ks + c$ for some recurrent configuration c and some $k \geq d_s$. By Theorem 1.6, after firing s, every non-sink vertex will fire exactly once to stabilize the resulting configuration, and c will again be obtained. Assume that at each step of this stabilization process, we have fired the lexicographically first unstable vertex (recalling that $V(\Gamma) = \{1, \ldots, n\}$) and let Q be the ordering of $V(\Gamma)$ corresponding to this firing, with s the first vertex of Q. Thus, $D = D^Q$, and it suffices to prove the theorem for $T = Q$.

Without loss of generality, s is the first vertex of S, since otherwise we may rotate S. Suppose $D^S|_{\tilde{V}(\Gamma)}$ is not recurrent. After firing s from D^S, we may legally fire every non-sink vertex once. However, by Theorem 1.6, this does not result in a stable configuration, and so some non-sink vertex v of S is unstable. If S' is the sequence of vertex firings obtained by transposing v with the vertex preceding it in S, then $D^S = D^{S'}$ by Lemma 2.4. Let S' be obtained from S by bringing v to the front of the ordering, so $D^S = D^{S'}$. Let $\tilde{S}$ be obtained by rotating S' so that s is again the first vertex, and v is the last. Observe that $D^{\tilde{S}}$ is obtained from $D^{S'}$ by firing v, so that $D^{\tilde{S}} \sim D^{S'}$. By repeating this procedure, we may fire unstable non-sink vertices until we obtain a stable configuration c, and Theorem 1.6 guarantees that this configuration is recurrent. By Theorem 1.5, in fact $c = D^T|_{\tilde{V}(\Gamma)}$. Thus, it suffices to prove the theorem when S is such that $D^S = D^T$.

If $D^S = D^T$ with $D^T|_{\tilde{V}(\Gamma)}$ recurrent, but $S \neq T$, it must be that at some point while stabilizing after firing s from D^S, we have a choice of two unstable vertices to

fire. In fact, by the characterization of $Q = T$ already given, it must be that in firing S, we have not always fired the lexicographically first vertex when we have a choice of unstable vertices. Thus, by performing a series of rotations on S, we eventually obtain a new ordering $\tilde{S}$ such that two vertices of $D^{\tilde{S}}$ are unstable, but the lexicographically first unstable vertex v comes later in $\tilde{S}$. We have already seen that we may move that v to the front of $\tilde{S}$ to obtain a new ordering S' with $D^{S'} = D^{\tilde{S}}$. By repeating this process, we eventually obtain T, and the proof is complete. $\qquad \square$

Corollary 3.21. *Fix $P \in \mathcal{P}_k(\Gamma)$ and let S_1 and S_2 be orderings of P. Then if $D_{\Gamma_P}^{S_1} \sim D_{\Gamma_P}^{S_2}$, also $D_{\Gamma}^{S_1} \sim D_{\Gamma}^{S_2}$.*

Proof. By Theorem 3.20, it suffices to prove the result when S_2 is obtained from S_1 by either a rotation or swapping two adjacent sets $U, V \in S_1$ with V appearing second, and the vertex $V \in V(\Gamma_P)$ unstable in $D_{\Gamma_P}^{S_1}$. If S_2 is obtained from S_1, then of course both $D_{\Gamma_P}^{S_1} \sim D_{\Gamma_P}^{S_2}$ and $D_{\Gamma}^{S_1} \sim D_{\Gamma}^{S_2}$. On the other hand, if V can be legally fired from $D_{\Gamma_P}^{S_1}$, by examining degrees, using minimality, and invoking Lemma 3.18, it follows that χ_V can be legally fired from $D_{\Gamma_P}^{S_1}$. In that case, $D_{\Gamma}^{S_1} = D_{\Gamma}^{S_2}$. $\qquad \square$

Corollary 3.22. *Let $S \sim T$, with $S, T \in \mathcal{S}_k(\Gamma)$. Then there is a sequence of vertex orderings $S = S_1, S_2, \ldots, S_\ell = T$ such that $S_i \sim S_{i-1}$ and S_i is obtained from S_{i-1} for $i \geq 2$ either by a rotation or by transposing two adjacent sets $U, V \in S_{i-1}$ such that one of $\{\chi_U, \chi_V\}$ is a legal firing from $D^{S_{i-1}}$.*

Proof. Let $P \in \mathcal{P}_k(\Gamma)$ be the partition corresponding to S. By Lemma 3.18, the divisor $\pi_P(D^S)$ is the minimum effective divisor relative to the ordering S of the vertices of Γ_P. Now, if $D_{\Gamma_P}^{S_1} \sim D_{\Gamma_P}^{S_2}$ for some orderings S_1 and S_2 of P, also $D_{\Gamma}^{S_1} \sim D_{\Gamma}^{S_2}$ by Lemma 3.21. Thus, it suffices to prove the corollary for S an ordering of the vertices of a graph, and so the result follows from Theorem 3.20. $\qquad \square$

Corollary 3.23. *Fix $P \in \mathcal{P}_k(\Gamma)$ and let S_1 and S_2 be orderings of P. Then $D_{\Gamma_P}^{S_1} \sim D_{\Gamma_P}^{S_2}$ iff $D_{\Gamma}^{S_1} \sim D_{\Gamma}^{S_2}$.*

Proof. One direction is already given by Corollary 3.21. For the other direction, when $D_{\Gamma}^{S_1} \sim D_{\Gamma}^{S_2}$, we may assume by Corollary 3.22 that S_2 is obtained from S_1 by either a rotation or by swapping two adjacent sets $U, V \in S_1$, the latter of which, V, may be legally fired from D^{S_1}. In the former case, certainly $D_{\Gamma_P}^{S_1} \sim D_{\Gamma_P}^{S_2}$. In the latter case, vertex V of Γ_P is unstable in $D_{\Gamma_P}^{S_1}$, and so $D_{\Gamma_P}^{S_1} = D_{\Gamma_P}^{S_2}$. $\qquad \square$

In order to relate the lower Betti numbers to the higher, we must relate the coarser partitions of Γ to the finer.

Definition 3.24 (Ordered refinement)**.** Let $T = V_1, \ldots, V_k \in \mathcal{S}_k(\Gamma)$. If there exists $S \in \mathcal{S}_{k-1}(\Gamma)$ such that $S = V_1 \cup V_2, V_3, \ldots V_k$, we say T is an ordered refinement of S, and write $T > S$.

Observe that $T = V_1, \ldots, V_k \in \mathcal{S}_k(\Gamma)$ is the ordered refinement of some S iff there exists $u \in V_1$ and $v \in V_2$ such that $uv \in E(\Gamma)$. Furthermore, if T is an ordered refinement, the S such that $S < T$ is uniquely determined.

Lemma 3.25. *Fix $T = V_1, \ldots, V_k \in \mathcal{S}_k(\Gamma)$ such that T has an ordered refinement S.
Then*

$$D^T - D^S = \sum_{v \in V_1} \mathrm{wt}(v, V_2) v \neq \vec{0}$$

Proof. The expression for $D^T - D^S$ is immediate from Lemma 2.4. Certainly we have
$D^T - D^S \geq \vec{0}$, and that the divisor is not equal to $\vec{0}$ follows from the fact that $V_1 \cup V_2$
is connected. $\square$

We define the set of ordered connected k-partition classes by $C\mathcal{S}_k(\Gamma) = \mathcal{S}_k(\Gamma)/ \sim$.
We will henceforth assume that every $C \in C\mathcal{S}_k(\Gamma)$ is endowed with a fixed represen-
tative element $S_C \in C$—the lexicographically least element is a natural choice. For
$S \in C$, if τ is the permutation taking S to S_C, we define $\mathrm{sign}(S) = \mathrm{sign}(\tau)$.

For $0 \leq k \leq n-1$, define the k^{th} $\mathrm{Cl}(\Gamma)$-graded free module of Γ by

$$F_k(\Gamma) = \bigoplus_{C \in C\mathcal{S}_{k+1}(\Gamma)} R(-D^{S_C}).$$

Let e_C be the identity element in $R(-D^{S_C})$. For $1 \leq k \leq n-1$, the free modules
of Γ come with homomorphisms $\bar{\phi}_k : F_k(\Gamma) \to F_{k-1}(\Gamma)$ defined by

$$\phi_k(e_C) = \sum_{\substack{T \in C \\ \exists S < T}} \mathrm{sign}(S)\,\mathrm{sign}(T) x^{D^T - D^S} e_{[S]}.$$

We obtain new homomorphisms ϕ_k from the $\bar{\phi}_k$ by normalizing the coefficients, i.e.
for $\bar{\phi}_k(e_C) = a_1 m_1 + \cdots + a_\ell m_\ell$ for distinct monomials $m_1, \ldots, m_\ell$ and non-zero
$a_1, \ldots, a_\ell \in \mathbb{Z}$, we define

$$\phi_k(e_C) = \frac{a_1}{|a_1|} m_1 + \cdots + \frac{a_\ell}{|a_\ell|} m_\ell.$$

Proposition 3.26. *The map ϕ_k preserves degrees.*

Proof. For any term f of the polynomial in the $[S]$ component of $\phi_k(e_{[T]})$, there exists
$\tilde{T} \sim T$ and $\tilde{S} \sim S$ such that $f = \lambda x^{D^{\tilde{T}} - D^{\tilde{S}}}$ for some $\lambda \in \mathbb{C}$. Then $f \in R_{D^T - D^S}$. Thus,
since $e_{[S]} \in R_{D^S}$, the $[S]$ component of $\phi_k(e_{[T]})$ lies in

$$R_{D^T - D^S} R_{D^S} = R_{D^T}.$$

It follows that ϕ_k preserves degrees. $\square$

Remark 3.27. For $T \in \mathcal{S}_k(\Gamma)$ with $S < T$, the $[S]$ component of $\phi_k(e_{[T]})$ has at most
two terms. This is because for any $\tilde{S} \sim S$ and $\tilde{T} \sim T$ with $\tilde{S} < \tilde{T}$, the first set in
$\tilde{S}$ agrees with the first set in S, and so we have at most two choices for the first two
sets of $\tilde{T}$.

Conjecture 3.28. *Let $F_i = F_i(\Gamma)$ for $0 \leq i \leq n-1$. The sequence*

$$\mathcal{F}(\Gamma) : \qquad F_0 \xleftarrow{\phi_1} F_1 \xleftarrow{\phi_2} F_2 \longleftarrow \cdots \xleftarrow{\phi_{n-1}} F_{n-1} \longleftarrow 0$$

is a minimal graded free resolution of $S/I(\Gamma)$.

Although a complete proof remains elusive, a partial proof is given below. We show that (i) coker $\phi_1 = S/I(\Gamma)$, (ii) for every k we have $\operatorname{im}\phi_k \subseteq \ker\phi_{k-1}$, and (iii) if the sequence is a free resolution, it is minimal. Thus, all that remains in order to prove the conjecture is that $\ker\phi_{k-1} \subseteq \operatorname{im}\phi_k$ for all k.

Note. After the deadline for substantive edits to theses had passed, it was noticed that the definition for ϕ_k was incorrect. The original version of the mapping is given in this revision as $\bar\phi_k$, and the corrected version as ϕ_k. However, the original definition of ϕ_k is retained in the following partial proof. The difference is relevant only in Case (1) of the proof that $\operatorname{im}\phi_k \subseteq \ker\phi_{k-1}$. A completed and corrected proof will appear elsewhere.

Partial proof of Conjecture 3.28. We first show that $\operatorname{im}\phi_k \subseteq \ker\phi_{k-1}$ for all $k \geq 2$. For given k, fix $C \in C\mathcal{S}_k(\Gamma)$. We have

$$
\phi_{k-1}(\phi_k(e_C)) = \phi_{k-1}\left(\sum_{\substack{T \in C \\ \exists S < T}} \operatorname{sign}(S)\operatorname{sign}(T) x^{D^T - D^S} e_{[S]}\right)
$$
$$
= \sum_{\substack{T \in C \\ \exists S < T}} \sum_{\substack{S' \in [S] \\ \exists Q < S'}} \operatorname{sign}(Q)\operatorname{sign}(S')\operatorname{sign}(S)\operatorname{sign}(T) x^{D^T - D^S} x^{D^{S'} - D^Q} e_{[Q]}.
$$

Suppose $Q < S' \sim S < T$ as in each summand. Two scenarios are possible: either (1) there are distinct vertex sets $V_x, V_y, V_z \in T$ such that $V_x \cup V_y \cup V_z \in Q$, or (2) there are distinct vertex sets $V_{x_1}, V_{x_2}, V_{y_1}, V_{y_2} \in T$ such that $V_{x_1} \cup V_{x_2} \in Q$ and $V_{y_1} \cup V_{y_2} \in Q$. In case (2), without loss of generality we have

$$
Q = V_1 \cup V_2, V_3, \ldots, V_{\ell-1}, V_\ell \cup V_{\ell+1}, V_{\ell+2}, \ldots, V_k
$$

and

$$
T = V_\ell, V_{\ell+1} \ldots, V_k, V_1, V_2, \ldots, V_{\ell-1}.
$$

In addition, we have $S_1' > Q$ given by

$$
S_1' = V_1, V_2, V_3, \ldots, V_{\ell-1}, V_\ell \cup V_{\ell+1}, V_{\ell+2}, \ldots, V_k
$$

and $S_1 < T$ given by

$$
S_1 = V_\ell \cup V_{\ell+1}, V_{\ell+2} \ldots, V_k, V_1, V_2, \ldots, V_{\ell-1}.
$$

Since S_1 is obtained from S_1' by a rotation of its sets, by Lemma 2.2 we have $S_1 \sim S_1'$. Now similarly define $\tilde{T} \sim T$ by $\tilde{T} = V_1, \ldots, V_k$, and

$$
S_2 = V_1 \cup V_2, V_3, \ldots, V_k
$$
$$
S_2' = V_\ell, V_{\ell+1}, \ldots, V_k, V_1 \cup V_2, V_3, \ldots, V_{\ell-1}
$$
$$
\tilde{Q} = V_\ell \cup V_{\ell+1}, V_{\ell+2} \ldots, V_k, V_1 \cup V_2, V_3, \ldots, V_{\ell-1}.
$$

Again, $\tilde{Q}$ is obtained from Q by rotation, so $\tilde{Q} \in [Q]$. Observe that

$$D^T - D^{S_1} + D^{S'_1} - D^Q = \sum_{v \in V_\ell} \mathrm{wt}(v, V_{\ell+1})v + \sum_{v \in V_1} \mathrm{wt}(v, V_2)v$$

$$= D^{\tilde{T}} - D^{S_2} + D^{S'_2} - D^{\tilde{Q}}$$

by Lemma 3.25.

Recall that the permutation τ that rotates an ordered set of k elements has sign $(-1)^{k+1}$. Thus, we may combine two terms in the above sum to obtain the summand

$$\mathrm{sign}(Q)\,\mathrm{sign}(S'_1)\,\mathrm{sign}(S_1)\,\mathrm{sign}(T)x^{D^T - D^{S_1}}x^{D^{S'_1} - D^Q}e_{[Q]}$$

$$+ \mathrm{sign}(\tilde{Q})\,\mathrm{sign}(S'_2)\,\mathrm{sign}(S_2)\,\mathrm{sign}(\tilde{T})x^{D^{\tilde{T}} - D^{S_2}}x^{D^{S'_2} - D^{\tilde{Q}}}e_{[\tilde{Q}]}$$

$$= (-1)^{(\ell-1)k}\,\mathrm{sign}(Q)\,(\mathrm{sign}\,(S_1))^2\,\mathrm{sign}(T)x^{D^T - D^{S_1} + D^{S'_1} - D^Q}e_{[Q]}$$

$$+ (-1)^{(\ell-2)k}\,\mathrm{sign}(\tilde{Q})\,(\mathrm{sign}\,(S_2))^2\,\mathrm{sign}(\tilde{T})x^{D^T - D^{S_1} + D^{S'_1} - D^Q}e_{[Q]}$$

$$= (-1)^{(\ell-1)k}\,\mathrm{sign}(Q)\,\mathrm{sign}(T)x^{D^T - D^{S_1} + D^{S'_1} - D^Q}e_{[Q]}$$

$$+ (-1)^{(\ell-2)k}(-1)^{(\ell-2)(k-1)}(-1)^{(\ell-1)(k+1)}\,\mathrm{sign}(Q)\,\mathrm{sign}(T)x^{D^T - D^{S_1} + D^{S'_1} - D^Q}e_{[Q]}$$

$$= \left((-1)^{\ell k + \ell} + (-1)^{\ell k + k + 1}\right)\mathrm{sign}(Q)\,\mathrm{sign}(T)x^{D^T - D^{S_1} + D^{S'_1} - D^Q}e_{[Q]}$$

$$= 0.$$

Case (1) requires no additional ideas but is equally tedious, so it is omitted. We now have $\phi_{k-1} \circ \phi_k = 0$.

To see that $S/I(\Gamma) = \mathrm{coker}\,\phi_1$, we must show that $\mathrm{im}\,\phi_1 = I(\Gamma)$. First, note that $F_0 = R$ since $C\mathcal{S}_1(\Gamma) = \{[V(\Gamma)]\}$ and $D^{V(\Gamma)} = \vec{0}$, so it is reasonable that $\mathrm{im}\,\phi_1 = I(\Gamma)$. Now let vertex n be a sink, so $x_n \in R$ is the corresponding indeterminate. For $P = V_1, V_2 \in \mathcal{P}_2(\Gamma)$, define the polynomials

$$f_P = \prod_{v \in V_1 \setminus \{s\}} x_v^{\mathrm{wt}(v, V_2)} - \prod_{v \in V_2 \setminus \{s\}} x_v^{\mathrm{wt}(v, V_1)}$$

and

$$\tilde{f}_P = \prod_{v \in V_1} x_v^{\mathrm{wt}(v, V_2)} - \prod_{v \in V_2} x_v^{\mathrm{wt}(v, V_1)}.$$

Let $G = \{f_P : P \in \mathcal{P}_2(\Gamma)\}$ and $\tilde{G} = \{\tilde{f}_P : P \in \mathcal{P}_2(\Gamma)\}$. Thus $\tilde{G}$ is the set of polynomials in G homogenized with respect to x_n. Observe that for orderings $S_1 = V_1, V_2$ and $S_2 = V_2, V_1$ of P, we have

$$\tilde{f}_P = x^{D^{S_1}} - x^{D^{S_2}} \in I(\Gamma).$$

In fact, choosing S_1 as the representative of $C = [S_1]$, we have

$$\tilde{f}_P = x^{D^{S_1}} - x^{D^{S_2}} = \sum_{\substack{T \in C \\ V(\Gamma) < T}} \mathrm{sign}(T)x^{D^T}e_R = \phi_1(e_C)$$

and so $\tilde{G} \subset \operatorname{im} \phi_1$. We claim that $\tilde{G}$ is a Gröbner basis for $I(\Gamma)$.

Theorem 3 of (2) states that G is a minimal Gröbner basis for the inhomogeneous version of $I(\Gamma)$ with respect to graded reverse lexicographic order (*grevlex*). (In fact, Theorem 3 uses $\tilde{G}$ and the additional polynomial $x_n - 1$, but the statement is equivalent.) Recall that a Gröbner basis H for an ideal J is characterized by the property that the ideal $\mathcal{I}_J$ given by leading terms of polynomials in J is generated by the leading terms of polynomials in H. Using *grevlex*, no leading term of the polynomials in $I(\Gamma)$ or $\tilde{G}$ contains x_n, so the leading terms of $I(\Gamma)$ and its inhomogeneous version are the same, as are the leading terms of G and $\tilde{G}$. Thus indeed $\tilde{G}$ is a Gröbner basis for $I(\Gamma)$, and we conclude that $\operatorname{coker} \phi_1 = S/I(\Gamma)$.

The minimality of $\mathcal{F}(\Gamma)$ is immediate since for $S < T$ we have $D^T - D^S \neq 0$ by Lemma 3.25, and it follows that the $[S]$ term of $\phi_k(e_{[T]})$ is not a nonzero scalar. $\qquad\square$

Finally, we have the following corollary to Conjecture 3.28, which generalizes Theorem 3.10.

Corollary 3.29. *Let* Γ *have a sink. Then*

$$\beta_k(I_\Gamma) = \sum_{P \in \mathcal{P}_{k+1}(\Gamma)} |\{c : c \text{ is a minimal recurrent sandpile on } \Gamma_P\}|.$$

Proof. Using, in order, Conjecture 3.28, Corollary 3.23, Lemma 3.18, and Lemma 2.7, we have

$$
\begin{aligned}
\beta_k(I_\Gamma) &= |C\mathcal{S}_{k+1}(\Gamma)| \\
&= \sum_{P \in \mathcal{P}_{k+1}(\Gamma)} |\{[S] \in C\mathcal{S}_{k+1}(\Gamma) : S \text{ an ordering of } P\}| \\
&= \sum_{P \in \mathcal{P}_{k+1}(\Gamma)} |\{[D^S] \in \operatorname{Cl}(\Gamma) : S \text{ an ordering of } P\}| \\
&= \sum_{P \in \mathcal{P}_{k+1}(\Gamma)} |\{[\pi_P(D^S)] \in \operatorname{Cl}(\Gamma_P) : S \text{ an ordering of } P\}| \\
&= \sum_{P \in \mathcal{P}_{k+1}(\Gamma)} |\{[D^S] \in \operatorname{Cl}(\Gamma_P) : S \text{ an ordering of } V(\Gamma_P)\}| \\
&= \sum_{P \in \mathcal{P}_{k+1}(\Gamma)} |\{c : c \text{ is a minimal recurrent sandpile on } \Gamma_P\}|.
\end{aligned}
$$

$\qquad\square$

Observe that in fact Theorem 3.10 is a special case.

Chapter 4

Computation and Applications

In this chapter, we examine a few corollaries to Conjecture 3.28, and so every proof will assume that the conjecture holds. We compute the coarse Betti numbers for trees and complete graphs, and discuss how the coarse Betti numbers of a graph are affected when edges are added. As a result, we obtain tight bounds for the coarse Betti numbers of undirected graphs.

Lemma 4.1. *Let Γ be a tree. If $D, E \in \mathrm{Div}(\Gamma)$ are such that $\deg(D) = \deg(E)$, then $D \sim E$.*

Proof. Fix a sink s. Note that if D is not effective, by repeatedly firing the sink and stabilizing, we obtain a divisor D' that is effective away from the sink. Thus, by Theorem 1.8, since D and E have the same degree, it suffices to show that both divisors superstabilize to the same configuration. We claim that the only superstable configuration on Γ is $\vec{0}$.

Suppose c is a superstable configuration that has $c_v > 0$ for some vertex v. Since Γ is a tree, there is a unique path $v_1, \ldots, v_k$ from $v_1 = v$ to $v_k = s$. Let $\Gamma_1, \ldots, \Gamma_t$ be the components of $\Gamma \setminus \{v\}$ that do not contain v_2. Since Γ is a tree, each neighbor of v other than v_2 is a vertex of a unique component Γ_i. Define

$$X = \{v\} \cup \bigcup_{1 \le i \le t} V(\Gamma_i).$$

Since v_2 is not an element of X, neither is s, so X is a subset of $\tilde{V}(\Gamma)$. We claim that X is a legal firing. For $u \in X$, if $u \ne v$, then every neighbor of u is also in X, so the firing is legal at u. On the other hand, every neighbor of v except for v_2 is in X, and $c_v > 0$, so the firing is legal at v as well. It follows that X is a legal firing, and so c is not superstable. $\square$

Theorem 4.2. *Let Γ be a tree on n vertices. Then $\beta_k(\Gamma) = \binom{n-1}{k+1}$.*

Proof. Γ has only one minimal recurrent configuration by Lemma 4.1. Furthermore, for $P \in \mathcal{P}_{k+2}(\Gamma)$, also Γ_P is a tree and so has only one minimal recurrent configuration. Thus, by Corollary 3.29, we have

$$\beta_k(\Gamma) = |\mathcal{P}_{k+2}(\Gamma)|.$$

Let E_k be the collection of subsets of $E(\Gamma)$ of cardinality k. Define $g : \mathcal{P}_k(\Gamma) \to E_{n-k}$ by

$$g(P) = \{vw \in E(\Gamma) : \exists U \in P \text{ such that } v, w \in U\},$$

i.e. $g(P)$ is the set of edges that are contained within a vertex set of P. It is straightforward to verify that g is a bijection. Since $|E(\Gamma)| = n - 1$, it follows that $\beta_k = \binom{n-1}{n-k-2} = \binom{n-1}{k+1}$. $\square$

Theorem 4.3. *Let Γ have an edge uv and let Γ' be obtained from Γ by changing the multiplicity of uv to a different positive value. Then $\beta_k(\Gamma) = \beta_k(\Gamma')$.*

Proof. Fix k. Certainly $\mathcal{P}_{k+2}(\Gamma) = \mathcal{P}_{k+2}(\Gamma')$ since the connected vertex sets of Γ and Γ' are identical. Furthermore, for $P \in \mathcal{P}_{k+2}(\Gamma)$, the partition graphs Γ_P and Γ'_P differ only by the multiplicity of an edge, if they differ at all. Thus, letting $s = u$ be a sink, it suffices by Corollary 3.29 to show that Γ and Γ' have the same number of minimal recurrent configurations. We claim that the minimal recurrent configurations are identical.

By Theorem 1.6, a configuration c is minimally recurrent iff there exists an ordering S of the vertices such that $D^S = d_s s + c$, with s the first vertex of S and D^S stable away from s. On the other hand, for $S \in \mathcal{S}_n(\Gamma)$ with s the first vertex of S, we have $D^S = d_s s + c$ for some recurrent configuration c iff every non-sink vertex is stable in D^S. Let r be such that $\mathrm{wt}_\Gamma(s, v) = \mathrm{wt}_{\Gamma'}(s, v) + r$. We claim that for vertex orderings S with s the first vertex, we have $D_\Gamma^S = D_{\Gamma'}^S + rs$ with D_Γ^S stable on the non-sink vertices iff $D_{\Gamma'}^S + rs$ is. This entails that the minimal recurrent configurations on Γ and Γ' are identical.

Let $S = v_1, \ldots, v_n$ be an ordering of $V(\Gamma)$ with $s = v_1$. By Lemma 2.4, for $i > 1$ we have

$$(D_\Gamma^S)_{v_i} = \sum_{i < j \leq k} \mathrm{wt}_\Gamma(v_i, v_j) = \sum_{i < j \leq k} \mathrm{wt}_{\Gamma'}(v_i, v_j) = (D_{\Gamma'}^S)_{v_i},$$

and also

$$
\begin{aligned}
(D_\Gamma^S)_s &= \sum_{1 < j \leq k} \mathrm{wt}_\Gamma(s, v_j) \\
&= \mathrm{wt}_{\Gamma'}(s, v) + r + \sum_{\substack{1 < j \leq k \\ v_j \neq v}} \mathrm{wt}_{\Gamma'}(s, v_j) \\
&= r + \sum_{1 < j \leq k} \mathrm{wt}_{\Gamma'}(s, v_j) \\
&= (D_{\Gamma'}^S)_s + r.
\end{aligned}
$$

Finally, we need to check that D_Γ^S is stable away from the sink iff $D_{\Gamma'}^S$ is. It suffices to check the condition at v. Suppose $r > 0$. Then certainly if v is stable for D_Γ^S, also v is stable for $D_{\Gamma'}^S$. On the other hand, if v is stable for D_Γ^S with $v = v_i$,

$$(D_\Gamma^S)_v = \sum_{i < j \leq k} \mathrm{wt}_\Gamma(v, v_j) \leq d_{\Gamma'}(v) - (r + 1)$$

since s comes before v in S and $\mathrm{wt}(s,v) \geq r+1$. Since $d_\Gamma(v) = d_{\Gamma'}(v) - r$, also $(D_\Gamma^S)_v \leq d_\Gamma(v) - 1$. The case when $r < 0$ is symmetric, and the proof is complete. $\square$

Theorem 4.4. *Let Γ' be obtained from Γ by adding an edge uv. Then $\beta_k(\Gamma) \leq \beta_k(\Gamma')$.*

Proof. The situation here is almost identical to that of Theorem 4.3. Again fix k. If $P \in \mathcal{P}_{k+2}(\Gamma)$, then also $P \in \mathcal{P}_{k+2}(\Gamma')$ since a connected vertex set of Γ is also connected in Γ'. Furthermore, for $P \in \mathcal{P}_{k+2}(\Gamma)$, the partition graphs Γ_P and Γ'_P differ only by the presence of an additional edge, if they differ at all. Thus, letting $s = u$ be a sink, by Corollary 3.29 it suffices to show that Γ has no more minimal recurrent configurations than Γ'. In fact, we claim that any minimal recurrent configuration of Γ is also a minimal recurrent configuration of Γ'. Using the same argument as in the proof of Theorem 4.3, it suffices to show that for any ordering S with s the first vertex and D_Γ^S stable away from s, we have $D_{\Gamma'}^S = D_\Gamma^S + s$ and $D_{\Gamma'}^S$ stable away from s. Again, an identical argument to that in the proof of Theorem 4.3 shows that $D_{\Gamma'}^S = D_\Gamma^S + s$, and it still holds that for $v_i \neq s$ and $(D_\Gamma^S)_{v_i} < d_\Gamma(v_i)$ we have

$$(D_{\Gamma'}^S)_{v_i} = (D_\Gamma^S)_{v_i} < d_\Gamma(v_i) \leq d_{\Gamma'}(v_i).$$

The proof is complete. $\square$

Definition 4.5 (Chain of sets). Given a set A, an ascending chain of length k of subsets of A is a sequence $S_1 \subseteq S_2 \subseteq \cdots \subseteq S_k$, where $S_i \subseteq A$ for $1 \leq i \leq k$. The chain is strictly ascending if each containment $S_i \subset S_{i+1}$ is proper.

Theorem 4.6. *Let $\Gamma = K_n$ with $n \geq 2$. Then $\beta_k(\Gamma)$ is the number of strictly ascending chains of length k of non-empty subsets of $A = \{1, \ldots, n-1\}$.*

Proof. Let X be the set of strictly ascending chains of length k of non-empty subsets of $A = \{1, \ldots, n-1\}$. Define the map $g : X \to C\mathcal{S}_{k+1}(\Gamma)$ by

$$g(S_1 \subseteq \cdots \subseteq S_{k+1}) = [S_1, \quad S_2 \setminus S_1, \quad S_3 \setminus S_2, \quad \ldots, \quad S_{k+1} \setminus S_k, \quad V(\Gamma) \setminus S_k].$$

Observe that for $P \in \mathcal{P}_k(\Gamma)$, also Γ_P is a complete graph, though possibly with edges of higher multiplicity. By Theorem 4.3, we can ignore the multiplicity of the edges. Thus, by Corollary 3.29, it suffices to show that g is a bijection when $k = n-1$.

Given $C \in C\mathcal{S}_n(\Gamma)$, use Lemma 2.2 to choose $T = v_1, \ldots, v_n \in C$ so that $v_n = n$. Define $S_i = \cup_{j \leq i} v_i$ for $1 \leq i \leq n-2$. Then $S_1 \subseteq \cdots \subseteq S_{n-1}$ is a chain in X, and

$$g(S_1 \subseteq \cdots \subseteq S_{n-1}) = C,$$

so g is surjective.

To show that g is injective, we first prove that for $S = v_1, \ldots, v_n \in \mathcal{S}_n(\Gamma)$ and σ a legal set firing from D^S, there exists $1 \leq \ell \leq n$ such that $\sigma = \sum_{i=1}^\ell v_i$. By Lemma 2.4,

$$D_{v_i}^S = \sum_{i < j \leq k} \mathrm{wt}_\Gamma(v_i, v_j) = |\{v_j \in V(\Gamma) : i < j\}|.$$

Let $\deg(\sigma) = \ell$. Since Γ is complete, when σ fires, each vertex in σ loses $n - \ell$ grains of sand. Thus, the only vertices that can fire, by the characterization above, are those vertices v_i with $i \leq \ell$, and hence $\sigma = \sum_{i=1}^{\ell} v_i$. In other words, firing σ simply rotates S by Lemma 2.2. By Proposition 1.3, any script can be written as a sequence of set firings, and a simple induction shows that each of these set firings is in fact just a rotation of S. Thus, there is a unique element $u_1, \ldots, u_n$ of $[S]$ with $n = u_n$. It follows that g is injective, and the proof is complete. $\qquad\qquad\square$

In particular, for $n \geq 2$ we have $\beta_1(K_n) = 2^(n - 1) - 1$ and $\beta_{n-1}(K_n) = (n - 1)!$. The number $\beta_{k+1}(K_{n+2})$ gives the $T(n, k)$ term of sequence A053440 of the On-Line Encyclopedia of Integer Sequences (12).

Remark 4.7. The combination of Theorems 4.2, 4.3, 4.4, and 4.6 gives tight upper and lower bounds for the coarse Betti numbers of undirected graphs.

Appendix A

Elementary Graph Theory

A **graph** Γ is a tuple (V, E) where V is an arbitrary set, called the **vertices**, and E is a collection of unordered pairs $\{u, v\}$ of distinct elements of V, called the **edges**. Intuitively, a graph is a network of nodes (the vertices), some pairs of which are connected in an undirected fashion (the edges). We will write $V(\Gamma)$ for the set of vertices of Γ and $E(\Gamma)$ for the set of edges. Instead of writing the edge $\{u, v\} \in E(\Gamma)$ as an ordinary set, we will normally use the more concise notation $uv \in E(\Gamma)$. If we permitted edges of the form vv for $v \in V(\Gamma)$, such edges would be called **loops**. Loops are ignored in this thesis because they are occasionally a nuisance, and contribute nothing to the theory. A **multigraph** is a graph Γ whose edge set $E(\Gamma)$ is a multiset. The **degree** d_v of a vertex v is the number of edges in which v participates, $d_v = |\{vw \in E(\Gamma) : w \in V(\Gamma)\}|$.

Theorem A.1.
$$\sum_{v \in V(\Gamma)} d_v = 2|E(\Gamma)|.$$

A **directed graph** (or **digraph**) Γ is a graph whose edge set consists of *ordered* pairs of vertices. Intuitively, a digraph is a network with unidirectional connections between nodes. We use the notation $uv \in E(\Gamma)$ to denote the edge from u to v. The **in-degree** of a vertex v is given by $\text{indeg}(v) = |\{uv \in E(\Gamma)\}|$, and the **out-degree** is given by $\text{outdeg}(v) = |\{vw \in E(\Gamma)\}|$. Every directed graph Γ has a corresponding undirected graph whose vertex set is $V(\Gamma)$ and whose edges set is $E(\Gamma)$ with the ordering ignored. Similarly, every undirected graph Γ can be thought of as the digraph on $V(\Gamma)$ with edges uv and vu for every $\{u, v\} \in E(\Gamma)$.

We define $\text{wt}(v, w) = |\{vw \in E(\Gamma)\}|$ for $v, w \in V(\Gamma)$. For $X, Y \subseteq V(\Gamma)$, we define $\text{wt}(X, Y) = \sum_{x \in X} \sum_{y \in Y} \text{wt}(x, y)$, and for $x, y \in V(\Gamma)$ we use the shorthand $\text{wt}(x, Y) = \text{wt}(\{x\}, Y)$ and $\text{wt}(X, y) = \text{wt}(X, \{y\})$.

A **walk** on a graph is a sequence of vertices $v_1, \ldots, v_n$ such that $v_i v_{i+1} \in E(\Gamma)$. A directed walk is defined identically on digraphs. A **path** is a walk with no repeated vertices—if a walk exists between distinct u and v, then so does a path. A **closed walk** is a walk $v_1, \ldots, v_n$ such that $v_1 = v_n$. A **cycle** is a closed walk on at least three distinct vertices whose only repeated vertices are the first and last. Directed paths, closed directed walks, and directed cycles are defined for digraphs by replacing

the word "walk" with "directed walk" in the corresponding definition for undirected graphs. An **Euler walk** is a closed walk that traverses every edge exactly once.

An (undirected) graph is **connected** if for every pair $u, v \in V(\Gamma)$, there is a path from u to v. A **tree** is a connected undirected graph with no cycles and no edges of multiplicity greater than 1. A **directed tree** is an undirected tree with a single orientation chosen for each edge.

Theorem A.2. *A tree on n vertices has $n - 1$ edges.*

A digraph is **weakly connected** if the corresponding undirected graph (the graph whose edges are the edges of the digraph with their ordering ignored) is connected, and it is **strongly connected** if for every pair $u, v \in V(\Gamma)$, there is a directed path from u to v. A digraph Γ is **Eulerian** if it is weakly connected and if $\operatorname{indeg} u = \operatorname{outdeg} u$ for every $u \in V(\Gamma)$.

Example A.3. Every connected undirected graph is Eulerian. (It is common elsewhere to define undirected Eulerian graphs as those with even degree, in which case there would certainly be non-Eulerian undirected graphs.)

Theorem A.4. *If Γ is Eulerian, it has an Euler walk.*

Corollary A.5. *If Γ is Eulerian, it is strongly connected.*

Two graphs Γ_1 and Γ_2 are **isomorphic** if there exists a map $f : V(\Gamma_1) \to V(\Gamma_2)$ such that $uv \in E(\Gamma_1)$ iff $f(u)f(v) \in E(\Gamma_2)$. For any graph Γ, only the cardinality of $V(\Gamma)$ matters, not the actual elements: replacing the elements of $V(\Gamma)$ while leaving $E(\Gamma)$ intact simply creates an isomorphic graph.

For a positive integer n, the **complete graph** on n vertices is the undirected graph K_n whose vertex set is $\{1, \ldots, n\}$ and whose edge set is

$$E(K_n) = \{ij : i \neq j \in V(K_n)\}.$$

For $n \geq 3$, we define the **cycle graph** on n vertices as the undirected graph C_n with $V(C_n) = \{1, \ldots, n\}$ and

$$E(C_n) = \{\{1, 2\}, \{2, 3\}, \ldots, \{n - 1, n\}, \{n, 1\}\}.$$

The reader is referred to (4) for a thorough treatment of the subject.

References

[1] Matthew Baker and Serguei Norine, *Riemann-Roch and Abel-Jacobi theory on a finite graph*, Advances in Mathematics **215** (2007), no. 2, 766 – 788.

[2] Robert Cori, Dominique Rossin, and Bruno Salvy, *Polynomial ideals for sandpiles and their Gröbner bases*, Theoretical Computer Science **276** (2002), no. 1-2, 1–15.

[3] Deepak Dhar, *Self-organized critical state of sandpile automaton models*, Physical Review Letters **64** (1990), no. 14, 1613 – 1616.

[4] Reinhard Diestel, *Graph theory*, third ed., Graduate Texts in Mathematics, no. 173, Springer, Berlin, 2005.

[5] David Eisenbud, *Commutative algebra with a view toward algebraic geometry*, Graduate Texts in Mathematics, no. 150, Springer-Verlag, New York, 1995.

[6] Alexander E. Holroyd, Lionel Levine, Karola Mszros, Yuyal Peres, James Propp, and David B. Wilson, *Chip-firing and rotor-routing on directed graphs*, In and Out of Equilibrium 2 (Vladas Sidoravicius and Maria Eulália Vares, eds.), Birkhäuser, Basel, 2008, pp. 331–364.

[7] Sean Jacques, *Betti numbers of graph ideals*, Phd thesis, University of Sheffield, 2004.

[8] Criel Merino López, *Chip firing and the Tutte polynomial*, Annals of Combinatorics **1** (1997), no. 1, 253 – 259.

[9] Ezra Miller and Bernd Sturmfels, *Combinatorial commutative algebra*, Graduate Texts in Mathematics, no. 227, Springer, New York, 2005.

[10] Irena Peeva and Bernd Sturmfels, *Syzygies of codimension 2 lattice ideals*, Mathematische Zeitschrift **229** (1998), no. 1, 163–194.

[11] David Perkinson, Jacob Perlman, and John Wilmes, *Primer for the algeraic geometry of sandpiles*, September 2009.

[12] N. J. A. Sloane, *The On-Line Encyclopedia of Integer Sequences*, 1996.