A Structure Theory of the Sandpile Monoid for Directed Graphs

László Babai* Evelin Toumpakari†

June 3, 2010

Abstract

The Abelian Sandpile Model is a diffusion process on (directed) graphs, studied, under various names, in statistical physics, discrete dynamical systems, theoretical computer science, algebraic graph theory, and other fields. The model takes a directed multigraph $\mathfrak{X}$ with a sink accessible from all nodes; associates a configuration space with $\mathfrak{X}$ and defines transition rules between the configurations; and finally, defines a a finite commutative monoid $\mathcal{M}$ (the *sandpile monoid*) on the "stable configurations" and a finite abelian group $\mathcal{G}$ (the *sandpile group*) on the "recurrent configurations." We add the *sandpile semigroup* $\mathcal{S}$ and the Rees quotient $\mathcal{S}/\mathcal{G}$, the *sandpile quotient* to the list.

We study the structure of these algebraic objects and their connection to the combinatorial structure of the underlying directed graphs. We demonstrate that the basic theory follows from elementary facts about commutative monoids. In particular, we point out that $\mathcal{G}$ is both the *unique minimal ideal* and the *universal group quotient* of $\mathcal{M}$. We also note that the *semilattice of idempotents* of a finite commutative monoid $\mathcal{M}$ is also the *universal semilattice quotient* of $\mathcal{M}$ and that this semilattice arises as the meet semilattice of a lattice which, in the context of sandpiles, we call the *sandpile lattice*. Our main result establishes a dual isomorphism between the sandpile lattice and the lattice of ideals of the accessibility poset of cyclic strong components (strong components which contain a cycle) of the underlying digraph. As a consequence, we characterize the sandpile lattices up to isomorphism as finite distributive lattices. Finally we introduce the notion of *transience class* of a sandpile and relate it to the nilpotence class of sandpile quotient. The "transience class" concept offers a new direction of study of the Abelian Sandpile Model.

1 Introduction

1.1 The Abelian Sandpile Model

The Abelian Sandpile Model takes a finite directed multigraph ("digraph" for short) $\mathfrak{X}$ with a *sink* accessible from every vertex along directed paths. The nodes other than the sink are called *sites*. We put a nonnegative finite number of "grains of sand" on each site (such an assigment is a *configuration*) and let a site *topple*, i. e., pass one grain to each of its out-neighbors, when it is in the position to do so (i. e., when it has at least as many grains as its out-degree). The sink swallows all grains it receives and never topples. From any starting configuration, repeated toppling leads to a unique *stable configuration* (configuration where no topplings are possible) in a finite number of steps.

This diffusion process models a phenomenon called "self-organized criticality" in statistical physics (Bak, Tang, Wiesenfeld[7], cf. [21]). It offers a rich structure that has been studied extensively in statistical physics and in a number of other disciplines, including probability theory, algorithms and complexity theory, algebraic graph theory, discrete dynamical systems, cellular automata. The process is a special case of a widely studied "chip-firing game" [31, 12, 13, 26, 1, 23]. A variant of the process was introduced by N. Biggs under the name "dollar-game" [10].

*University of Chicago. Email: `laci@cs.uchicago.edu`

†Oxford International College, Oxford, U.K. Email: `etoumpakari@gmail.com`

The model, for the case of the $n \times n$ square grid augmented with a sink, was introduced in the original paper by Bak et al. [7]. The general case was introduced by Dhar [16] for undirected and by Speer [29] for directed underlying graphs. A number of remarkable properties of the model were discovered by Dhar et al. in the early 90s [16, 17].

One of the concepts Dhar introduced was a commutative monoid structure on the set of stable configurations, the *sandpile monoid*. Dhar also introduced the key distinction between between *transient* and *recurrent* configurations and made the fundamental observation that the recurrent configurations form an abelian group, the *sandpile group*, under the sandpile monoid operation.

We give a brief comparison with two related models, the chip firing game and the dollar game. The sandpile monoid structure does not arise in the chip firing game because the sink is missing. Biggs's dollar game allows negative stacks of chips and therefore all of its configurations are "recurrent;" the model leads directly to the sandpile group but misses the sandpile monoid.

While a large body of literature exists on the sandpile group (both on the general theory, partly for its analogy with the Picard group of an algebraic curve [27, 10, 5, 8] and on the sandpile groups of specific classes of graphs [17, 14, 6, 20, 33, 24]), we could not find a systematic study of the sandpile monoid in the literature.

In this paper we develop a structure theory of the sandpile monoid and its connection to the combinatorial structure of the underlying directed graph. The key organizing concept of this structure theory will be the *sandpile lattice*, the lattice of idempotents of the sandpile monoid. We give a description of this lattice in terms of the accessibility partial order of the cyclic strong components of the underlying digraph $\mathfrak{X}$.

We introduce a random walk on the stable states which we call the *Sandpile Markov Chain*; its recurrent states are exactly the recurrent states of the sandpile. We briefly touch upon mixing issues of this Markov Chain.

We introduce two more associated semigroup structures, the *sandpile semigroup* $\mathcal{S}$, and the *sandpile quotient* $\mathcal{S}/\mathcal{G}$. The sandpile semigroup is defined as the ideal of $\mathcal{M}$ consisting of the stabilizations of the nonzero configurations; the sandpile quotient is the Rees quotient $\mathcal{S}/\mathcal{G}$ (Section 7.1). We have $\mathcal{M} = \mathcal{S} \cup \{0\}$ and characterize the case when $\mathcal{S} = \mathcal{M}$ (Prop. 5.7). We always have $\mathcal{G} \subseteq \mathcal{S} \subseteq \mathcal{M}$. We characterize the important case when $\mathcal{S}$ has a unique idempotent; this happens in particular when the sites span a strongly connected subgraph. We note that $\mathcal{S}$ has a unique idempotent exactly when the sandpile quotient is nilpotent (Sec. 7). Of particular interest is the *nilpotence class* which gave rise to the notion of the *transience class* of a sandpile, the maximum number of grains in a configuration that stabilizes to a transient configuration. This a concept that has already spawned a new direction of study (see Section 7 and the "Open problems" section (Section 8)).

1.2 Structure of the paper: introduction to the theory

Sections 2 and 3 serve as a tutorial; we build up the theory from scratch. We demonstrate that many of the basic facts about the sandpile monoid and the sandpile group are immediate consequences of simple observations about finite commutative monoids. In particular, we point out that $\mathcal{G}$ can be defined as the *unique minimal ideal* of $\mathcal{M}$ (a substructure) and is canonically isomorphic to the *universal group quotient* (a quotient structure) of $\mathcal{M}$, a feature shared by all finite commutative monoids. This observation gives a simple general explanation of the link between the two types of definition of the sandpile group found in the literature: definition through configurations and definition through generators and relations.

We describe the "standard presentation" (in terms of generators and relations) of the sandpile monoid $\mathcal{M}$ and show that it follows naturally from the definition of toppling. Then we show that the standard presentation of the sandpile group $\mathcal{G}$ follows immediately from the fact that $\mathcal{G}$ is the universal group quotient of $\mathcal{M}$. While the standard presentation of the sandpile group has been widely used, we have not seen a mention of the standard presentation of the sandpile monoid in the literature, although the latter seems to be more naturally connected to toppling; accordingly, the derivation of the standard presentation of the sandpile group via the sandpile monoid seems the most natural to us and does not seem to have been pointed out previously.

The standard presentations link both structures to the reduced Laplacian of the underlying digraph $\mathfrak{X}$ and thus to the abstract description of $\mathcal{G}$ introduced in algebraic graph theory (the "Picard group" of the graph). Through the digraph version of Kirchhoff's celebrated Matrix-Tree Theorem, the standard presentation of $\mathcal{G}$ immediately leads to the remarkable observation (due to Dhar [16] in the undirected case) that the order of $\mathcal{G}$ is the same as the number of spanning reverse arborescences (directed spanning trees oriented toward the sink).

We describe the basic theory, essentially due to Dhar, in Section 2. In Section 3 we give simple and generic proofs of the basic facts (with the solutions to some exercises deferred to the Appendix).

1.3 Structure of the paper: new results

The main new results of this paper concern two semilattices (commutative idempotent semigroups) associated with the sandpile monoid: the *semilattice of idempotents* (a substructure) and the *universal semilattice quotient* (a quotient structure). Again, these two are canonically isomorphic for all finite commutative monoids; moreover, they arise as the meet semilattice of a lattice. In the case of sandpiles, we call this lattice the *sandpile lattice.* Our main result, stated in Section 4.3, describes the structure of the sandpile lattice $\mathcal{L}$ in terms of the accessibility poset $\mathcal{P}$ of the cyclic strong components (strong components containing a cycle) of the underlying digraph $\mathfrak{X}$, *viz,* $\mathcal{L}$ is canonically isomorphic to the dual lattice of ideals of $\mathcal{P}$. As a corollary, we characterize the sandpile lattices, up to isomorphism, as finite distributive lattices (Section 4.3).

We answer basic questions about the connections between the algebraic structure of the sandpile monoid and the combinatorial structure of the underlying digraph in Section 5; these results also serve as ingredients of the main proofs to follow in the subsequent sections.

The proofs of the main results (about the sandpile lattice), stated in Section 4.3, are given in Section 6.

In Section 7 we introduce the concept of the *transience class* of a sandpile model and connect it to the nilpotence class of the Rees quotient $\mathcal{S}/\mathcal{G}$ (the "sandpile quotient"). In a subsequent paper we shall characterize the sandpile models with bounded transience class. Motivated by the present work, the asymptotic rate of growth of the transience class has been studied in [2] and [3]; the latter paper found that for some classes of undirected underlying graphs, the transience class grows exponentially, while in [2] we found that the transience class of the square grid sandpile model grows polynomially.

We close this paper with open problems (Section 8). An Appendix includes some of the short and sweet proofs omitted from Section 3.

We warn the reader that the term "lattice" appears in two meanings in this paper: as a discrete subgroup of $\mathbb{R}^n$ (in Fact 2.27 below) and as a poset with join and meet (in the Abstract and in the main results, Sections 4 and 6).

A preliminary version of the material presented in this paper appeared as part of the junior author's thesis [32].

2 The model

2.1 Configurations, toppling

In this paper, we use the term "digraph" to mean a *directed multigraph.* A digraph thus consists of a set V of vertices, a set E of edges, and an incidence function $i : E \to V \times V$. If $i(e) = (u, v)$ then we call $u \in V$ the *tail* and $v \in V$ the *head* of the edge $e \in E$. Let a_{ij} denote the number of $i \to j$ edges. The matrix $A = (a_{ij})$ is the *adjacency matrix* of our digraph. *Undirected graphs* are digraphs with a symmetric adjacency matrix ($a_{ij} = a_{ji}$).

Definition 2.1 Let $\mathfrak{X} = (V, E)$ be a digraph. We say that vertex j is *accessible* from vertex i if there exists a directed path (possibly of length zero) from i to j. So accessibility is a reflexive, transitive relation.

A *sink* is a vertex of out-degree zero.

An *Abelian Sandpile Model* is based on a digraph $\mathfrak{X} = (V, E)$ with a sink accessible from all vertices ("underlying digraph" for short). Note that under this condition the sink is unique. The sink plays a special role; all other vertices are called *sites*. We denote the sink by s, the set of sites by $V_0 = V \setminus \{s\}$, and the subgraph of $\mathfrak{X}$ induced on V_0 by $\mathfrak{X}_0 = (V_0, E_0)$. (By "subgraph" we mean a digraph obtained by taking a subset of the vertices and a subset of the edges; so a subgraph is also a digraph; we avoid the awkward term "sub-digraph.") We call $\mathfrak{X}_0$ the *site-digraph*.

We say that our model is *undirected* if the site-digraph $\mathfrak{X}_0$ is undirected.

We write $\deg(i)$ for the out-degree of vertex i in $\mathfrak{X}$, and $\deg_0(i)$ for the out-degree of site i relative to the site-digraph $\mathfrak{X}_0$. So $\deg(i) = \sum_{j \in V} a_{ij}$, and $\deg(i) - \deg_0(i) = a_{i,s}$.

A *configuration* (or *state*) of the model is a function $\mathbf{x} : V_0 \to \mathbb{N}$ where $\mathbb{N}$ is the set of nonnegative integers. The integer $\mathbf{x}(i)$ is thought of as the number of sandgrains at site i (the *height* of the sandpile). The *configuration space* is $\mathbb{N}^{V_0}$, the set of all configurations.

The site $i \in V_0$ is *stable* if $\mathbf{x}(i) < \deg(i)$. An unstable site can be *toppled*, sending one grain through each edge leaving i. So $\mathbf{x}(i)$ is reduced by $\deg(i)$, and for each site j, the height $\mathbf{x}(j)$ increases by a_{ij}. Grains reaching the sink disappear (those grains "fell off" the "board" consisting of the sites). The sink never topples. A configuration is *stable* if all sites are stable. A *toppling sequence* is a sequence of topplings.

Definition 2.2 The *score* of a toppling sequence S is the configuration $\mathbf{z} \in \mathbb{N}^{V_0}$ which records how many times each site was toppled in S (site i was toppled $\mathbf{z}(i)$ times).

Fact 2.3 *Starting with any configuration and toppling unstable sites in succession, we arrive at a stable configuration in a finite number of steps.*

(See the Appendix for a proof.)

A sequence of topplings that results in a stable configuration is called a *stabilizing sequence* or an *avalanche*.

The theory of the Abelian Sandpile Model is based on the following fundamental observation.

Fact 2.4 *Given an initial configuration* $\mathbf{x}$*, every avalanche leads to the same stable configuration.*

The avalanches also have the same length. In fact, the following stronger statement holds.

Fact 2.5 *Given the configurations* $\mathbf{x}$ *and* $\mathbf{y}$*, every toppling sequence from* $\mathbf{x}$ *to* $\mathbf{y}$ *has the same score. In particular, every avalanche starting at* $\mathbf{x}$ *has the same score.*

This result explains the adjective "abelian" in the name of the model.

The result was found independently by Dhar [16] and Bjørner et al. [13]. For completeness, we include a simple proof in the Appendix.

2.2 The Sandpile Monoid and the Sandpile Semigroup

Our standard reference to semigroup theory is Grillet [18]. A *monoid* is a semigroup with identity.

Let us fix an underlying digraph $\mathfrak{X} = (V, E)$ with n vertices. All definitions below will refer to the configurations of the Abelian Sandpile Model built on $\mathfrak{X}$. We denote the unique stable configuration resulting from stabilizing the configuration $\mathbf{x} \in \mathbb{N}^{V_0}$ by $\sigma(\mathbf{x})$.

Definition 2.6 The *transition digraph* associated with the underlying digraph $\mathfrak{X}$ is the infinite digraph $(\mathbb{N}^{V_0}, G \cup T)$ where $\mathbb{N}^{V_0}$ is the configuration space; for configurations $\mathbf{x}, \mathbf{y}$ we set $(\mathbf{x}, \mathbf{y}) \in G$ if $\mathbf{y}$ is obtained from $\mathbf{x}$ by adding a grain at some site, and $(\mathbf{x}, \mathbf{y}) \in T$ if $\mathbf{y}$ is obtained from $\mathbf{x}$ by toppling some site.

Definition 2.7 We say that a configuration $\mathbf{y} \in \mathbb{N}^{V_0}$ is *accessible* from the configuration $\mathbf{x} \in \mathbb{N}^{V_0}$ if $\mathbf{y}$ is accessible from $\mathbf{x}$ in the transition digraph.

A *monoid* is a semigroup with identity. Our configuration space, $\mathbb{N}^{V_0}$, is the *free commutative monoid* under pointwise addition; we denote the (pointwise) sum of two configurations $\mathbf{x}, \mathbf{y} \in \mathbb{N}^{V_0}$ by $\mathbf{x} + \mathbf{y}$, i. e., for all $i \in V_0$ we set $(\mathbf{x} + \mathbf{y})(i) = \mathbf{x}(i) + \mathbf{y}(i)$. For each site $i \in V_0$, let $\mathbf{t}_i \in \mathbb{N}^{V_0}$ denote the configuration defined by

$$\mathbf{t}_i(j) = \delta_{ij} \tag{1}$$

(a single grain on site i, empty otherwise). The configurations $\mathbf{t}_i$ $(i \in V_0)$ are the free generators of $\mathbb{N}^{V_0}$.

Definition 2.8 (Sandpile monoid) The *sandpile monoid* $\mathcal{M} = \mathcal{M}(\mathfrak{X})$ is the set of stable configurations under the operation of pointwise addition followed by stabilization. We denote this operation by $\oplus$.

So, for stable configurations $\mathbf{x}$ and $\mathbf{y}$ we set

$$\mathbf{x} \oplus \mathbf{y} := \sigma(\mathbf{x} + \mathbf{y}). \tag{2}$$

With this operation, $\mathcal{M}$ is a finite commutative monoid (commutative semigroup with identity); the all-zero configuration 0 is the identity. We note that $\sigma : \mathbb{N}^{V_0} \to \mathcal{M}$ is a homomorphism: for any $\mathbf{x}, \mathbf{y} \in \mathbb{N}^{V_0}$,

$$\sigma(\mathbf{x} + \mathbf{y}) = \sigma(\mathbf{x}) \oplus \sigma(\mathbf{y}). \tag{3}$$

Note that $\sigma(\mathbf{x}) = \mathbf{x}$ if and only if $\mathbf{x} \in \mathcal{M}$. In particular, the map $\sigma : \mathbb{N}^{V_0} \to \mathcal{M}$ is surjective.

We observe that the order of the sandpile monoid is

$$|\mathcal{M}| = \prod_{i \in V_0} \deg(i). \tag{4}$$

We find it useful to give a name to a particular subsemigroup of the sandpile monoid.

Definition 2.9 (Sandpile semigroup) We call the set of stabilizations of the nonzero configurations the *sandpile semigroup*, denoted by $\mathcal{S}$. So $\mathcal{S} = \sigma(\mathbb{N}^{V_0} \setminus \{0\})$.

Clearly, $\mathcal{M} = \mathcal{S} \cup \{0\}$. Note that 0 may or may not belong to $\mathcal{S}$; indeed $0 \in \mathcal{S}$ exactly if 0 is accessible from some nonzero configuration. We shall characterize this circumstance in terms of properties of the underlying digraph $\mathfrak{X}$ (Proposition 5.7). We also note that $\mathcal{S}$ is an *ideal* of $\mathcal{M}$ because $\mathbb{N}^{V_0} \setminus \{0\})$ is an ideal in the free commutative monoid $\mathbb{N}^{V_0}$.

2.3 Recurrence and transience; the Sandpile Markov Chain

Observation 2.10 *If $\mathbf{x}, \mathbf{y} \in \mathcal{M}$ then $\mathbf{y}$ is accessible from $\mathbf{x}$ exactly if there exists $\mathbf{z} \in \mathcal{M}$ such that $\mathbf{x} = \mathbf{y} \oplus \mathbf{z}$.* $\square$

Definition 2.11 We say that a stable configuration $\mathbf{x} \in \mathcal{M}$ is *recurrent* if it is accessible from every configuration; otherwise $\mathbf{x} \in \mathcal{M}$ is *transient*.

Let $\mathbf{x}_{\max}$ denote the *saturated* stable configuration, i. e., $(\forall i \in V_0)(\mathbf{x}_{\max}(i) := \deg(i) - 1)$. Clearly, $\mathbf{x}_{\max}$ is accessible from every configuration, so $\mathbf{x}_{\max}$ is recurrent. The following observation is useful.

Fact 2.12 *A stable configuration is recurrent if and only if it is accessible from $\mathbf{x}_{\max}$.* $\square$

Definition 2.13 Consider the finite Markov Chain on the set $\mathcal{M}$ of stable configurations defined by the following transition rule: add a grain at a random site (chosen uniformly); stabilize. We call this the *Sandpile Markov Chain.*

In the theory of discrete-time Markov Chains, a state $\mathbf{x}$ is called "recurrent" if starting from the state $\mathbf{x}$ the process will almost surely (with probability 1) return to $\mathbf{x}$ in a positive number of steps.

Proposition 2.14 *A stable sandpile configuration* $\mathbf{x}$ *is recurrent in the sense of Def. 2.11 if and only if* $\mathbf{x}$ *is a recurrent state of the Sandpile Markov Chain.*

This is immediate from the following more general observation.

Observation 2.15 *Let* M *be a finite Markov Chain with a set* $\mathcal{M}$ *of states. Let* X *be the transition digraph of* M *(edges indicate positive transition probability). Suppose there is a state* $\mathbf{x}_0$ *that is accessible from all states. Then a state* $\mathbf{x} \in \mathcal{M}$ *is recurrent in the sense of almost sure return if and only if* $\mathbf{x}$ *is accessible from all states.*

Proof: Assume first that $\mathbf{x}$ is not accessible from all states and therefore not accessible from $\mathbf{x}_0$. Then with positive probability, the process, started from $\mathbf{x}$, will reach $\mathbf{x}_0$ without passing through $\mathbf{x}$; and then it can never return to $\mathbf{x}$.

Conversely, if $\mathbf{x}$ is accessible from all states then for some $t > 0$ and $p > 0$, the process, started from any state, will visit $\mathbf{x}$ within t steps with probability $\geq p$. So the probability that no visit occurs within kt steps is $\leq (1-p)^k$ which goes to 0 as $k \to \infty$. □

An alternative definition that appears in the literature would call a stable configuration $\mathbf{x}$ "recurrent" if it is possible to return to $\mathbf{x}$ from $\mathbf{x}$ in a positive number of steps, i. e., if $\mathbf{x}$ is a cyclic vertex of the transition digraph (cf. Defs. 2.6 and 4.9). Below we refer to this as the "wrong definition." Evidently, we do not recommend it. It is not equivalent to ours; under this definition, the set of recurrent configurations would not necessarily form a group. However, in most cases of interest, the two definitions are in fact equivalent.

Proposition 2.16 *Given an underlying digraph* $\mathfrak{X}$*, the following are equivalent.*

(a) For all $\mathbf{x} \in \mathcal{M}$*,* $\mathbf{x}$ *is recurrent under the "wrong definition" if and only if* $\mathbf{x}$ *is recurrent in the sense of Def. 2.11;*

(b) the stable configurations that are recurrent under the "wrong definition" form a group (under the operation in $\mathcal{M}$*);*

(c) the sandpile semigroup $\mathcal{S}$ *has a unique idempotent.*

(d) the transience class of $\mathfrak{X}$ *is finite (see Def. 7.5).*

We shall prove this in Section 3.3.

Circumstance (c) is an important one which occurs in most cases of interest; in particular, (c) holds whenever the site-digraph $\mathfrak{X}_0$ is strongly connected. Thus the "wrong" definition is equivalent to ours in most cases of interest; and to our knowledge, only in such cases has the definition been applied in the literature. We shall characterize circumstance (c) in terms of the structure of the underlying digraph $\mathfrak{X}$ in Theorem 4.15. Its equivalence with (d) (Cor. 7.13) further underscores the interest in this case.

2.4 Sandpile monoid: basic facts

An array of remarkable facts about the sandpile model was discovered by Deepak Dhar around 1990. We list some of these, and provide what we believe are the "proofs from the Book:" we infer them from simple general considerations about commutative monoids. One of the most startling observations is that, for the case of undirected graphs, the number of recurrent configurations is the number of spanning trees. We give the digraph version of this statement.

A *reverse arborescence* in a rooted digraph is a directed spanning tree directed toward the root. The assumption that the sink is accessible from every site is equivalent to saying that $\mathfrak{X}$ has at least one reverse arborescence (where the sink is the root).

Fact 2.17 *The number of recurrent configurations is equal to the number of reverse arborescences.*

For the proof we need to discover, following Dhar, that the recurrent configurations form a group, and find its presentation in terms of generators and relations. This is the subject of the next section.

2.4.1 The Sandpile Group

Fact 2.18 *The recurrent configurations form a group (as a subsemigroup of the sandpile monoid).*

This group is called the *sandpile group*; we denote it by $\mathcal{G} = \mathcal{G}(\mathfrak{X})$. We note that the sandpile group is *not a submonoid* of the sandpile monoid since it has a different identity.

2.4.2 Generators and relations

Definition 2.19 We call a site i *irrelevant* if $\deg(i) = 1$; otherwise i is *relevant.*

For each site $i \in V_0$, let $\mathbf{s}_i = \sigma(\mathbf{t}_i)$ (the configuration $\mathbf{t}_i$ was defined by Equation (1)). Note that $\mathbf{s}_i = \mathbf{t}_i$ for all relevant sites; if i is irrelevant and j is its unique out-neigbor then $\mathbf{s}_i = \mathbf{s}_j$ if j is a site, and $\mathbf{s}_i = 0$ if j is the sink.

The $\mathbf{s}_i$ ($i \in V_0$) generate the sandpile monoid; we refer to them as the *standard generators* of $\mathcal{M}$. By the definition of toppling, they satisfy the diffusion relations

$$\deg(i)\mathbf{s}_i = \sum_{j \in V_0}{}^{\oplus} a_{ij}\mathbf{s}_j \quad (i \in V_0), \tag{5}$$

where $\sum^{\oplus}$ refers to addition in $\mathcal{M}$ (Eq. (2)). (Recall that a_{ij} is the number of edges from vertex i to vertex j.) We shall show that this set of relations is sufficient to define the sandpile monoid (among commutative monoids). We refer to this set of relations as the *standard presentation* of $\mathcal{M}$.

We then show that the same presentation also defines the sandpile group (among Abelian groups). We derive this from general considerations about commutative monoids (notably, that the unique minimal ideal of such a monoid is at the same time its "universal quotient").

Proposition 2.20

(i) The sandpile monoid is the commutative monoid generated by the symbols $\{\mathbf{x}_i : i \in V_0\}$ subject to the set of defining relations $\mathcal{R} = \{\deg(i)\mathbf{x}_i = \sum_{j\in V_0} a_{ij}\mathbf{x}_j : i \in V_0\}$.

(ii) The sandpile group is the abelian group generated by the symbols $\{\mathbf{x}_i : i \in V_0\}$ subject to the set of defining relations $\mathcal{R} = \{\deg(i)\mathbf{x}_i = \sum_{j\in V_0} a_{ij}\mathbf{x}_j : i \in V_0\}$.

We sketch the proof of (i); a simple general observation, given in Section 3, will show that (ii) is an immediate consequence.

Proof of part (i). Let $\widehat{\mathcal{M}}$ denote the commutative monoid defined by the given relations and let $\varphi : \mathbb{N}^{V_0} \to \widehat{\mathcal{M}}$ be the standard epimorphism; so $\varphi(\mathbf{t}_i) = \mathbf{x}_i \in \widehat{\mathcal{M}}$. As noted before, $\sigma : \mathbb{N}^{V_0} \to \mathcal{M}$ is an epimomorphism. The diffusion equations (5) tell us that the generators $\mathbf{s}_i = \sigma(\mathbf{t}_i)$ of $\mathcal{M}$ satisfy the set $\mathcal{R}$ of relations; therefore there is a unique epimorphism $\psi : \widehat{\mathcal{M}} \to \mathcal{M}$ such that $\sigma = \psi\varphi$. We need to show that ψ is an isomorphism. To this end we only need to show that ψ is injective.

Let $\mathbf{x} \in \mathbb{N}^{V_0}$ and let $\mathbf{y} \in \mathbb{N}^{V_0}$ be obtained from $\mathbf{x}$ by toppling site i. The i-th relation in $\mathcal{R}$ guarantees that $\varphi(\mathbf{x}) = \varphi(\mathbf{y})$. Applying this to each toppling in an avalanche we conclude that $\varphi(\mathbf{x}) = \varphi(\sigma(\mathbf{x}))$.

It follows that ψ is injective. Indeed let $\mathbf{u}, \mathbf{v} \in \widehat{\mathcal{M}}$; suppose $\psi(\mathbf{u}) = \psi(\mathbf{v})$. Let $\mathbf{u} = \varphi(\mathbf{x})$ and $\mathbf{v} = \varphi(\mathbf{y})$. Then $\sigma(\mathbf{x}) = \psi(\mathbf{u}) = \psi(\mathbf{v}) = \sigma(\mathbf{y})$; and therefore $\mathbf{u} = \varphi(\mathbf{x}) = \varphi(\sigma(\mathbf{x})) = \varphi(\sigma(\mathbf{y})) = \varphi(\mathbf{y}) = \mathbf{v}$. □

As noted above, irrelevant sites give rise to redundant generators. We can get rid of the irrelevant sites by a simple construction. Let us say that a *path is irrelevant* if all *internal* nodes of the path are irrelevant. In particular, a single edge is always an irrelevant path.

Define the *relevant reduction* $\widetilde{\mathfrak{X}}$ of $\mathfrak{X}$ as follows. The sites of $\widetilde{\mathfrak{X}}$ are the relevant sites of $\mathfrak{X}$; the sink $\widetilde{\mathfrak{X}}$ is the sink of $\mathfrak{X}$. For each maximal irrelevant path $x = x_0 \to x_1 \to \cdots \to x_t = y$ starting at a relevant site x (and ending either at a relevant site or at the sink), we put an edge $x \to y$ in $\widetilde{\mathfrak{X}}$. The multiplicity of the

$x \to y$ edge in $\widetilde{\mathfrak{X}}$ will be the number of irrelevant paths from x to y, so $\deg(x)$ will not change. Note that $\deg_0 x$ can go down if y is the sink.

The following is now immediate.

Proposition 2.21 (Relevant reduction) $\mathcal{M}(\widetilde{\mathfrak{X}}) \cong \mathcal{M}(\mathfrak{X})$. □

We have to comment on the case when $\mathfrak{X}$ is a reverse arborescence; in other words, when every site is irrelevant. In this case, the only stable configuration is the empty configuration, so $|\mathcal{M}(\mathfrak{X})| = 1$. This remains true for $\widetilde{\mathfrak{X}}$ as well; in this case, $\widetilde{\mathfrak{X}}$ has no sites at all, and therefore the only configuration is the empty configuration.

Corollary 2.22 $\mathcal{G}(\widetilde{\mathfrak{X}}) \cong \mathcal{G}(\mathfrak{X})$.

Indeed, isomorphic monoids have isomorphic minimal ideals. □

We have to warn, however, that $\mathcal{S}(\widetilde{\mathfrak{X}}) \cong \mathcal{S}(\mathfrak{X})$ is not true in general (Prop. 2.24).

Definition 2.23 We say that the site i is a *point of no return* if no cycle (and therefore no cyclic strong component) of $\mathfrak{X}$ is accessible from i.

Proposition 2.24 *For the underlying digraph $\mathfrak{X}$, the following are equivalent.*

(a) Either $\mathfrak{X}$ has no points of no return or it has a relevant point of no return.

(b) $\mathcal{S}(\widetilde{\mathfrak{X}}) \cong \mathcal{S}(\mathfrak{X})$.

We shall prove this equivalence at the end of Section 5.3.

This observation shows that $\mathcal{S}$ is *not* a structural invariant of $\mathcal{M}$ (in contrast to $\mathcal{G}$ which is).

2.4.3 The reduced Laplacian

Definition 2.25 The *Laplacian* $L = (L_{ij})_{ij\in V_0}$ of $\mathfrak{X}$ is an $n \times n$ matrix defined by

$$L_{ij} := \begin{cases} \deg(i) - a_{ii} & \text{if } i = j, \\ -a_{ij} & \text{otherwise.} \end{cases} \tag{6}$$

Definition 2.26 The *reduced Laplacian* $\Delta = (\Delta_{ij})_{ij\in V_0}$ of $\mathfrak{X}$ is defined as the $(n-1)\times(n-1)$ matrix obtained from the Laplacian L by deleting the row and column corresponding to the sink.

Let Λ denote the lattice (discrete subgroup) in $\mathbb{R}^{V_0}$ generated by the rows of the reduced Laplacian.

Fact 2.27 (Dhar[16]) *The sandpile group is isomorphic to the quotient $\mathbb{Z}^{V_0}/\Lambda$.*

Proof: Observe that this is just a rephrasing of item (ii) in Proposition 2.20. □

Now we have all we need for the proof of Fact 2.17.

Proof of Fact 2.17. Both numbers in question are equal to the determinant of the reduced Laplacian; the order of the sandpile group because of Fact 2.27; and the number of reverse arborescences because of the directed graph version of Kirchhoff's [22] classical Matrix-Tree Theorem (see Tutte [34], cf. [25]). □

Another immediate corollary to Fact 2.27 describes the structure of the sandpile group. According to the Fundamental Theorem of Finite Abelian Groups, $\mathcal{G}$ can be written as the direct sum of cyclic groups of orders $f_1, \ldots, f_{n_0}$ where $f_i \mid f_{i+1}$. The positive integers f_i are uniquely determined and are called the *invariants* of $\mathcal{G}$.

Fact 2.28 *The invariants of $\mathcal{G}$ are the diagonal elements of the Smith Normal Form of the reduced Laplacian.* □

3 Monoids: general considerations

We shall demonstrate that all facts listed in the preceding section follow from a few simple observations about finite commutative monoids.

3.1 The minimal ideal

Let us recall some notation and terminology. A *semigroup* is a set with an associative binary operation. A *monoid* is a semigroup with identity. If A, B are subsets of a semigroup then AB denotes the subset $\{ab : a \in A, b \in B\}$. A a subset I of a semigroup $\mathcal{S}$ is a *left ideal* if $\mathcal{S}I \subseteq I$; and a *right ideal* if $I\mathcal{S} \subseteq I$. It is an ideal if it is both a left and a right ideal. An ideal I is *minimal* if it is nonempty and the only proper subset of I that is an ideal if $\mathcal{S}$ is the empty set. The *kernel* of a semigroup is the intersection of all of its nonempty ideals. The kernel is either a minimal ideal or it is empty.

Let us say that an element $b \in \mathcal{S}$ is *fully accessible* if $(\forall a \in \mathcal{S})(\exists x \in \mathcal{S})(ax = b)$.

Proposition 3.1 *Let $\mathcal{M}$ be a finite commutative monoid. Then*

(a) the kernel of $\mathcal{M}$ is nonempty, i. e., $\mathcal{M}$ has a unique minimal ideal $\mathcal{G}$;

(b) $\mathcal{G}$ consists of the fully accessible elements of $\mathcal{M}$.

(c) $\mathcal{G}$ is an abelian group;

The proof of these facts is an easy exercise; for easy access, we include a proof in the Appendix.

The following simple consequence does not seem to have been pointed out previously.

Corollary 3.2 *The set $\mathcal{G}$ of recurrent configurations forms the (unique) minimal ideal of the sandpile monoid $\mathcal{M}$; therefore, $\mathcal{G}$ is a group.*

Proof: By definition, the recurrent configurations are precisely the fully accessible elements of the sandpile monoid $\mathcal{M}$, i. e., they form the kernel of $\mathcal{M}$. $\square$

With this we completed the proof of Fact 2.18.

In the light of the foregoing, we propose the following alternative definition of the sandpile group.

Alternative definition. The sandpile group is the (unique) minimal ideal of the sandpile monoid.

Note that this definition is equivalent to the definition given after Fact 2.18; and that definition in turn is equivalent to the definitions appearing in the literature, modulo Proposition 2.16 and the comments preceding it.

3.2 The universal group; generators and relations

Definition 3.3 Let $\mathcal{M}$ be a monoid, let $\mathcal{G}$ be a group and let $\varphi : \mathcal{M} \to \mathcal{G}$ be a homomorphism. We say that $(\varphi, \mathcal{G})$ is the *universal group* of $\mathcal{M}$ if every homomorphism from $\mathcal{M}$ to a group factors through φ.

It is a well known and simple fact that every monoid has a universal group and the universal group is unique up to the natural equivalence. The following correspondence of presentations in terms of generators and relations makes this immediate.

Observation 3.4 *If $\langle A|\mathcal{R}\rangle$ is a presentation of a monoid $\mathcal{M}$ (where A is a set of generators and $\mathcal{R}$ a set of relations) then $\langle A|\mathcal{R}\rangle$ is also a presentation of the universal group $\mathcal{G}$ of $\mathcal{M}$ as a group.* $\square$

We observe that the minimal ideal of a finite commutative monoid is also its universal group, under a natural homomorphism.

Fact 3.5 *Let $\mathcal{M}$ be a finite commutative monoid and let $\mathcal{G}$ be the minimal ideal of $\mathcal{M}$. We write $\mathcal{M}$ additively. Let $e \in \mathcal{G}$ be the identity in $\mathcal{G}$ and let $\varphi : \mathcal{M} \to \mathcal{G}$ be defined by $\varphi(x) := e + x$. Then $(\varphi, \mathcal{G})$ is the universal group of $\mathcal{M}$.*

Proof: Let $\psi : \mathcal{M} \to \mathcal{G}_1$ be a homomorphism to a group $\mathcal{G}_1$. We need to define a homomorphism $\eta : \mathcal{G} \to \mathcal{G}_1$ such that $\psi = \eta\varphi$. Let η be the restriction of ψ to $\mathcal{G}$. Let e_1 denote the identity of $\mathcal{G}_1$. Clearly, ψ maps all idempotents of $\mathcal{M}$ to e_1; in particular, $\psi(e) = \eta(e) = e_1$. Now, for $x \in \mathcal{M}$ we have $\psi(x) = \psi(x) + e_1 = \psi(x + e) = \eta(\varphi(x))$. □

Remark 3.6 There is something remarkable about the fact that the universal group, defined as a "quotient" structure, should also appear as a (natural) substructure (the minimal ideal). While Fact 3.5 must be folklore, we could not find it mentioned in [18]. We shall see another occurrence of this phenomemon in Section 4 (universal semilattice vs. the semilattice of idempotents of a finite commutative monoid, Fact 4.5), another fact we could not locate in [18].

The following consequence of Fact 3.5 does not seem to have been pointed out previously.

Corollary 3.7 *The sandpile group is the universal group of the sandpile monoid (under the homomorphism described in Fact 3.5).* □

This completes the proof of the inference (i) $\Rightarrow$ (ii) in Proposition 2.20. □

3.3 The wrong definition of recurrence

We prove Prop. 2.16. Let $\mathcal{H}$ denote the set of stable configurations that are recurrent under the "wrong definition," i.e., $\mathbf{x} \in \mathcal{H}$ if there is a cycle through $\mathbf{x}$ in the transition digraph. So $\mathbf{x} \in \mathcal{H}$ if and only if $(\exists \mathbf{y} \in \mathcal{S})(\mathbf{x} \oplus \mathbf{y} = \mathbf{x})$.

Let $\mathbf{e}$ denote the identity of $\mathcal{G}$. Obviously, $\mathcal{G} \subseteq \mathcal{H}$ (for $\mathbf{x} \in \mathcal{G}$ we take $\mathbf{y} := \mathbf{e}$). It is also clear that $\mathcal{H}$ is an ideal of $\mathcal{M}$. Indeed, if $\mathbf{x} \oplus \mathbf{y} = \mathbf{x}$ then $(\mathbf{x} \oplus \mathbf{z}) \oplus \mathbf{y} = \mathbf{x} \oplus \mathbf{z}$ for all $\mathbf{z} \in \mathcal{M}$. It follows that $\mathcal{H}$ is a semigroup and $\mathcal{G}$ is an ideal of $\mathcal{H}$; therefore $\mathcal{H}$ is a group if and only if $\mathcal{H} = \mathcal{G}$. This proves the equivalence of (a) and (b).

Let now $\mathbf{x} \in \mathcal{H}$ and $\mathbf{x} \oplus \mathbf{y} = \mathbf{x}$ for some $\mathbf{y} \in \mathcal{S}$. We may assume $\mathbf{y}$ is an idempotent. Indeed, $\mathbf{x} \oplus k\mathbf{y} = \mathbf{x}$ for all $k \geq 1$ by induction, and $k\mathbf{y}$ is an idempotent for some $k \geq 1$.

If now $\mathbf{y} = \mathbf{e}$ then $\mathbf{x} = \mathbf{x} \oplus \mathbf{e} \in \mathcal{G}$. This proves that (c) implies (b) (if $\mathcal{S}$ has a unique idepotent then necessarily $\mathbf{y} = \mathbf{e}$). Conversely, obviously, every idempotent of $\mathcal{S}$ belongs to $\mathcal{H}$, but only $\mathbf{e}$ belongs to $\mathcal{G}$. This proves that (a) implies (c).

The equivalence of (c) and (d) is given in Cor. 7.13. □

4 The Sandpile Lattice

This section contains the main results of the paper.

For the rest of this paper, a *lattice* will be a partially ordered set with join and meet (least upper bound and greatest lower bound).

First we describe elements of a general structure theory of finite commutative monoids (Sections 4.1 and 4.2). In Section 4.3 we shall apply this theory to the sandpile monoid.

4.1 The lattice of idempotents of a finite commutative monoid

An element x of a semigroup is *idempotent* if $x^2 = x$. Note that every nonempty finite semigroup has an idempotent. A *semilattice* is a commutative semigroup in which every element is idempotent. A semilattice defines a partial order on its elements by setting $a \leq b$ if $ab = a$; with this notation, we call the semilattice operation "meet" (greatest lower bound). A finite semilattice always has a smallest element (the meet of all elements); this element z is necessarily a *zero*, i. e., $zx = z$ for every x in the semilattice.

If the semilattice also has a greatest element, we say that the semilattice is *bounded.* The greatest element must be the identity, so a semilattice is bounded if and only if it is a monoid.

A finite bounded semilattice *defines a lattice:* every pair of elements has a "join" (least upper bound), namely, the meet of all upper bounds.

The following is now immediate.

Fact 4.1 *The idempotents of a finite commutative monoid form a submonoid which is a bounded semilattice. The zero of this semilattice is the identity of the minimal ideal of the monoid.* □

Definition 4.2 The *semilattice of idempotents* of the commutative monoid $\mathcal{M}$ is the submonoid consisting of the idempotents of $\mathcal{M}$. We view the operation in $\mathcal{M}$ as the "meet" operation in the semilattice of idempotents. The *lattice of idempotents* of the finite commutative monoid $\mathcal{M}$ is the lattice defined by the semilattice of idempotents.

4.2 The universal semilattice of finite monoids

Definition 4.3 Let $\mathcal{S}$ be a semigroup, let $\mathcal{L}$ be a semilattice, and let $\varphi : \mathcal{S} \to \mathcal{L}$ be a homomorphism. We say that $(\varphi, \mathcal{L})$ is the *universal semilattice* of $\mathcal{S}$ if every homomorphism from $\mathcal{S}$ to a semilattice factors through φ. We say that $\mathcal{S}$ is "a semilattice of the subsemigroups $\varphi^{-1}(u)$ for $u \in \mathcal{L}$."

Fact 4.4 *Every semigroup has a universal semilattice.*

Proof: Take the quotient of the semigroup by the relations $x^2 = x$ and $xy = yx$. □

Once again we find that a universal structure (this time the universal semilattice), defined as a "quotient" structure, also appears as a substructure (the semilattice of idempotents) (cf. Remark 3.6).

Fact 4.5 *For a finite commutative monoid $\mathcal{M}$, the universal semilattice is canonically isomorphic to the semilattice of idempotents of $\mathcal{M}$.*

It follows that $\mathcal{M}$ is a semilattice of semigroups with a unique idempotent.
Proof: Map every element x of the monoid to the unique idempotent in the subsemigroup generated by x. This map is a universal semilattice homomorphism. Moreover, the inverse image of each idempotent is a subsemigroup with a unique idempotent. □

It follows that the universal semilattice of a finite commutative monoid is bounded and therefore defines a lattice. We view the operation of the universal semilattice as the "meet" operation.

Definition 4.6 The *universal lattice* of a finite commutative monoid is the lattice defined by the universal semilattice of the monoid.

4.3 The sandpile lattice: main results

Now we come to the main results of this paper, connecting the combinatorial structure of the underlying digraph to the structure to the algebraic structure of the sandpile monoid.

Definition 4.7 The *sandpile lattice* $\mathcal{L} = \mathcal{L}(\mathfrak{X})$ associated with the underlying digraph $\mathfrak{X}$ is the lattice of idempotents of the sandpile monoid $\mathcal{M}(\mathfrak{X})$.

The following is the main consequence of our structural results.

Theorem 4.8 *For a finite lattice $\mathcal{L}$, the following are equivalent:*

(i) $\mathcal{L}$ is isomorphic to a sandpile lattice.

(ii) $\mathcal{L}$ is distributive.

Note that every finite lattice is isomorphic to the lattice of idempotents of a finite commutative monoid (namely, of its own meet semilattice). In particular, the universal lattice of a finite commutative monoid is *not necessarily distributive*, hence our characterization of the lattices corresponding to sandpile monoids puts a strong restriction on the structure of those lattices.

The proof of Theorem 4.8 goes through a description of the idempotents of the sandpile monoid in terms of the strong components of the underlying digraph $\mathfrak{X}$.

Definition 4.9 A vertex of a digraph is *cyclic* if it belongs to a cycle; and *acyclic* otherwise.

Note that the sink is an acyclic vertex.

Definition 4.10 The *strong components* of a digraph are the equivalence classes of the relation of mutual accessibility. If a strong component contains a cycle, we call it *cyclic*, otherwise *acyclic*.

So, an acyclic strong component consists of a single acyclic vertex. All vertices of a cyclic strong component are cyclic.

The strong components are partially ordered by the accessibility relation.

Definition 4.11 Let $\mathcal{P} = (\Omega, \preccurlyeq)$ be a partially ordered set. A subset $I \subseteq \Omega$ is an *ideal* of $\mathcal{P}$ if $(\forall x, y \in \Omega)($ if $x \in I$ and $y \preccurlyeq x$ then $y \in I)$.

Fact 4.12 *The ideals of a partially ordered set form a distributive lattice under set union and intersection.* □

We can now state the main result of this paper.

Theorem 4.13 *The following lattices are dually isomorphic:*

(i) The sandpile lattice $\mathcal{L}(\mathfrak{X})$.

(ii) The lattice of ideals of the accessibility partial order on the set of cyclic strong components of $\mathfrak{X}$.

Having a single idempotent is a strong structural constraint on a semigroup. For finite commutative semigroups it is equivalent to the nilpotence of the Rees quotient by the minimal ideal (see Section 7.1). Next we give a combinatorial characterization of the underlying digraphs $\mathfrak{X}$ for which $\mathcal{S}$ has a unique idempotent. We shall elaborate on the significance of this characterization in Section 7.

We need two definitions.

Definition 4.14 A DAG (directed acyclic graph) is a digraph without cycles.

So a DAG is a digraph in which every vertex is acyclic.

Theorem 4.15 *The sandpile semigroup $\mathcal{S}$ has a unique idempotent if and only if either $\mathfrak{X}$ is a DAG or $\mathfrak{X}$ has a unique cyclic strong component and no point of no return (i. e., the unique cyclic strong component is accessible from every site, cf. Def. 2.23).*

We note the following corollary.

Corollary 4.16 *If the site-digraph $\mathfrak{X}_0$ is strongly connected then the sandpile semigroup $\mathcal{S}$ has a unique idempotent (namely, the identity of $\mathcal{G}$).* □

5 Basic structure

In this section we describe some basic connections between the algebraic structure of the sandpile monoid and the combinatorial structure of the underlying digraph, to lay some of the groundwork for the proofs of the main results. As before, we fix an underlying digraph $\mathfrak{X}$; the corresponding monoid, semigroup, and group are $\mathcal{M}$, $\mathcal{S}$, $\mathcal{G}$, respectively.

5.1 Simple characterizations

The results in this section answer some natural questions about the possible relations between $\mathcal{M}$, $\mathcal{S}$, $\mathcal{G}$.

Recall that the $\mathbf{t}_i$ ($i \in V_0$) are the standard generators of the free commutative monoid $\mathbb{N}^{V_0}$; and $\mathbf{s}_i = \sigma(\mathbf{t}_i)$ ($i \in V_0$) are the standard generators of the monoid $\mathcal{M}$.

Proposition 5.1 *$\mathcal{S}$ is the subsemigroup of $\mathcal{M}$ generated by the $\mathbf{s}_i$ ($i \in V_0$).*

Given the significance of the Rees quotient $\mathcal{S}/\mathcal{G}$, we note that it is always defined, i. e., $\mathcal{G} \subseteq \mathcal{S}$.

Proposition 5.2 *$\mathcal{S}$ is a non-empty ideal of $\mathcal{M}$ and $\mathcal{G} \subseteq \mathcal{S}$.*

Proposition 5.3 *The following are equivalent:*

(i) $\mathcal{S} = \{0\}$;

(ii) $\mathcal{M} = \{0\}$;

(iii) all sites are irrelevant (have out-degree 1, cf. Def. 2.19);

(iv) $\mathfrak{X}$ *is a directed tree (a reverse arborescence).*

Condition (ii) implies $|\mathcal{G}| = 1$. The converse is not true, however. Let us consider a reverse arborescence T of $\mathfrak{X}$ directed to the sink. We say that an edge $i \to j$ in $\mathfrak{X}$ is a *back edge* if j is an ancestor of i in T, i. e., there is a $j \to \cdots \to i$ path in T.

Proposition 5.4 *The following are equivalent:*

(i) $|\mathcal{G}| = 1$;

(ii) $\mathcal{G} = \{\mathbf{x}_{\max}\}$ *where* $\mathbf{x}_{\max}$ *is the saturated configuration (see Fact 2.12);*

(iii) $\mathfrak{X}$ *has a unique reverse arborescence T; all edges not in T are back edges.*

Next we characterize the case $\mathcal{M} = \mathcal{G}$. A *DAG* is a directed acyclic graph (no directed cycles).

Proposition 5.5 *The following are equivalent:*

(i) $0 \in \mathcal{G}$ *(the zero configuration is recurrent);*

(ii) $\mathcal{M} = \mathcal{G}$ *(all stable configurations are recurrent);*

(iii) $\mathcal{M} = \mathcal{S} = \mathcal{G}$

(iv) $\mathfrak{X}$ *is a DAG.*

This case already covers all abelian groups. We are also interested in getting small DAGS for a given abelian group; it turns out we can tell the exact size of the smallest DAG for a given group.

Proposition 5.6 *(a) Given a finite abelian group G there exists an underlying DAG $\mathfrak{X}$ such that $\mathcal{G}(\mathfrak{X}) \cong G$.*

(b) If $|G| = \prod_{j=1}^{r} p_j$ *where the* p_j *are not necessarily distinct primes then the minimum number of edges of such a DAG is* $\sum_j p_j$*.*

Next we characterize the case $\mathcal{S} = \mathcal{M}$. Note that in all other cases, $\mathcal{S} = \mathcal{M} \setminus \{0\}$.

Proposition 5.7 *The following are equivalent:*

(i) $0 \in \mathcal{S}$

(ii) $\mathcal{M} = \mathcal{S}$*;*

(iii) the zero configuration is accessible from some (not necessarily stable) nonzero configuration;

(iv) $\mathfrak{X}_0$ *has a point of no return (see Def. 2.23).*

Next we study the case $|\mathcal{G}| = |\mathcal{M}| - 1$.

Proposition 5.8 *The following are equivalent:*

(i) $|\mathcal{G}| = |\mathcal{M}| - 1$*;*

(ii) 0 *is transient, all other stable configurations are recurrent*

(iii) the site-digraph $\mathfrak{X}_0$ *contains a unique cycle and all sites not on the cycle are irrelevant (have* $\deg = 1$*).*

Proposition 5.9 *If* $|\mathcal{G}| = |\mathcal{M}| - 1$ *then the sandpile group* $\mathcal{G}$ *is cyclic. The identity element of* $\mathcal{G}$ *is the saturated configuration* $\mathbf{x}_{\max}$*.*

Definition 5.10 Define a *directed wheel graph* as an underlying digraph whose sites induce a directed cycle, and at least one edge goes from each site to the sink.

Note that these are precisely the relevant reductions of the digraphs defined in part (iii) of Prop. 5.8 Therefore it suffices to prove Prop. 5.9 for directed wheel graphs (Prop. 2.21).

This observation foreshadows a far more general situation that characterizes the case when the nilpotence class of the Rees quotient $\mathcal{S}/\mathcal{G}$ (the "sandpile quotient") is bounded (see Theorem 7.15).

5.2 Preliminary lemmas

In this section we introduce some concepts that play a central role in analyising the sandpile monoid.

Throughout the paper, configurations considered are not necessarily stable unless stability is specifically assumed. In particular, in this section, the essence of several results will be in the case of unstable configurations.

Definition 5.11 Let $\mathbf{x} \in \mathbb{N}^{V_0}$ be a (not necessarily stable) configuration. The *weight* of $\mathbf{x}$ is $\text{weight}(\mathbf{x}) = \sum_{i \in V_0} \mathbf{x}(i)$ (total number of grains). For $B \subseteq V_0$ we let $\text{weight}_B(\mathbf{x}) = \sum_{i \in B} \mathbf{x}(i)$.

Definition 5.12

(i) Given a configuration $\mathbf{x} \in \mathbb{N}^{V_0}$, we say that a set $B \subseteq V_0$ of sites is *blank* if $\text{weight}_B(\mathbf{x}) = 0$ (there are no grains on B, see Def. 5.11). We also say $\mathbf{x}$ is blank on B.

(ii) We say that a subgraph of $\mathfrak{X}_0$ is *blank* if the set of sites in the subgraph is blank. We also say that $\mathbf{x}$ is blank on the subgraph.

The following is a useful necessary condition of accessibility (see Def. 2.1).

Lemma 5.13 (No blank cycles) *Let C be a cycle in $\mathfrak{X}_0$. If C is blank with respect to the configuration* $\mathbf{y}$ *but not blank with respect to configuration* $\mathbf{x}$ *then* $\mathbf{y}$ *is not accessible from* $\mathbf{x}$*. In particular, a recurrent configuration cannot have a blank cycle.*

Proof: If C is not blank, adding a grain or toppling a site cannot make C blank. This proves the first statement. For the second statement, we just need to see that C cannot be blank with respect to the saturated stable configuration $\mathbf{x}_{\max}$. Indeed, C contains at least one site of degree $\deg(i) \geq 2$ (because the sink is accessible from C). So $\mathbf{x}_{\max}(i) \geq 1$. □

Definition 5.14 Let $\mathbf{x} \in \mathbb{N}^{V_0}$ be a configuration. The *support* of $\mathbf{x}$, denoted by $\operatorname{supp}(\mathbf{x})$, is defined to be the set of vertices of $\mathfrak{X}_0$ on which $\mathbf{x}$ is non-zero, i.e., $\operatorname{supp}(\mathbf{x}) := \{v \in V_0 : \mathbf{x}(v) > 0\}$. We say that $\mathbf{x}$ is *concentrated* on $A \subseteq V_0$ if $A \supseteq \operatorname{supp}(\mathbf{x})$.

Corollary 5.15 *If the zero configuration is accessible from a configuration* $\mathbf{x}$ *then* $\mathbf{x}$ *is concentrated on the points of no return (see Def. 2.23).*

Proof: Suppose for a contradiction that $\mathbf{x}(i) \geq 1$ for a site $i \in V_0$ which is not a point of no return, i.e., some cycle C is accessible from i.

Let P be a path of length ℓ from i to C. Consider a counterexample with smallest possible ℓ. Then $\ell \neq 0$ since if $\ell = 0$ then $\mathbf{x}$ is not blank on C so 0 is not accessible from $\mathbf{x}$ by Lemma 5.13. For $\ell \geq 1$, consider a toppling sequence leading to the 0 configuration. At some point in this sequence, i must be toppled, putting a grain on the successor of i along P. This new configuration is therefore a counterexample with a smaller value of ℓ, a contradition. □

Definition 5.16 In a digraph, a vertex j is *accessible* from a set B of vertices if there exists $i \in B$ such that j is accessible from i.

Lemma 5.17 (Flushing a DAG) *Let* $\mathbf{x} \in \mathbb{N}^{V_0}$ *be a configuration concentrated on the points of no return (see Def. 2.23). Let A denote the set of sites accessible from* $\operatorname{supp}(\mathbf{x})$*. Set* $k := \prod_{i \in A} \deg(v)$*. Then* $\sigma(k\mathbf{x}) = 0$.

Proof: The induced subgraph $\mathfrak{X}_0[A]$ is a DAG. Let S be the set source vertices (vertices of in-degree 0) of the induced subgraph $\mathfrak{X}_0[A]$. Then $\sigma(\prod_{i \in S} \deg(i)\mathbf{x})$ is concentrated on $A \setminus S$. The result follows by induction on the depth of $\mathfrak{X}_0[A]$ (length of longest path in $\mathfrak{X}_0[A]$). □

Definition 5.18 A configuration $\mathbf{x}$ is *semistable* if for all sites $i \in V_0$ we have $\mathbf{x}(i) \leq 2\deg(i) - 2$. *Semistable stabilization* is a toppling sequence where toppling at $i \in V_0$ can occur only when $\mathbf{x}(i) \geq 2\deg(i) - 1$ and the final configuration is semistable.

Semistable stabilization is useful in establishing accessibility between configurations and in particular in showing that a configuration is recurrent. Variants of this concept have been used extensively in [2].

Lemma 5.19 (Flooding) *Let $u \in V_0$. Then there exists a semistable configuration* $\mathbf{x}$ *and* $k \in \mathbb{N}$ *such that* $\sigma(\mathbf{x}) = \sigma(k\mathbf{t}_u)$ *and* $\mathbf{x}(v) \geq \deg(v) - 1$ *for all vertices v accessible from u.*

Proof: Taking multiples of $\mathbf{t}_u$ means adding grains at u. If we apply semistable stabilization, the grains will eventually flood all vertices accessible from u; semistable toppling of v will never reduce the height below $\deg(v) - 1$. □

Definition 5.20 Let $\mathfrak{Y} = (V, E)$ be a digraph and let $D \subseteq V$. We call $\mathfrak{Y}$ a *functional* digraph with domain D if each vertex in D has out-degree 1 in $\mathfrak{Y}$ and all other vertices have out-degree 0. Such a digraph represents a function $D \to V$.

By a *functional subgraph* of the underlying digraph $\mathfrak{X}$ we mean a subgraph of $\mathfrak{X}$ that is functional with domain V_0. (We exclude the sink from the domain.)

Observation 5.21 *For a functional subgraph $\mathfrak{Y}$ of $\mathfrak{X}$, the following are equivalent:*

(i) $\mathfrak{Y}$ is a reverse arborescence.

(ii) $\mathfrak{Y}$ is acyclic. □

Let $\mathcal{T}$ denote the set of stable transient configurations, i. e., $\mathcal{T} = \mathcal{M} \setminus \mathcal{G}$. The following is immediate from equation (4) and Theorem 2.17, given that the sink is accessible from every vertex.

Corollary 5.22

(i) $|\mathcal{M}|$ is the number of functional spanning subgraphs of $\mathfrak{X}$.

(ii) $|\mathcal{T}|$ is the number of those functional spanning subgraphs of $\mathfrak{X}$ which contain a cycle. □

Corollary 5.23

$$|\mathcal{T}| \geq \prod_{i \in V_0} \deg_0(i). \tag{7}$$

Proof: $\prod_{i \in V_0} \deg_0(i)$ is the number of those functional subgraphs which have no edge pointing to the sink; no such subgraph is an arborescence of $\mathfrak{X}$. □

Definition 5.24 Let $\mathfrak{X}_i = (V_i, E_i)$ be underlying digraphs. The *sink-join* of the $\mathfrak{X}_i$ is the underlying digraph obtained by taking the disjoint union of the $\mathfrak{X}_i$ and then identifying all the sinks to a common sink.

Proposition 5.25 *If $\mathfrak{X}$ is the* sink-join *of $\mathfrak{X}_i$, $i = 1, \ldots, k$, then $\mathcal{M}(\mathfrak{X}) = \mathcal{M}(\mathfrak{X}_1) \times \cdots \times \mathcal{M}(\mathfrak{X}_k)$ and consequently $\mathcal{G}(\mathfrak{X}) = \mathcal{G}(\mathfrak{X}_1) \times \cdots \times \mathcal{G}(\mathfrak{X}_k)$.*

Proof: Clear. □

5.3 Proofs of the simple characterizations

Proof of Proposition 5.1. Immediate from Def. 2.9. □

Proof of Proposition 5.2. $\mathcal{S}$ is an ideal of $\mathcal{M}$ because $\mathbb{N}^{V_0} \setminus \{0\}$ is an ideal of $\mathbb{N}^{V_0}$. Now $\mathcal{G} \subseteq \mathcal{S}$ because $\mathcal{G}$ is the unique minimal ideal. □

Proof of Proposition 5.3. Each link in the following chain of equivalences is obvious: $(i) \iff (ii) \iff (iii) \iff (iv)$. □

Proof of Proposition 5.4. The equivalence $(i) \iff (ii)$ is obvious. The equivalence $(i) \iff (iii)$ follows from Fact 2.17 and an obvious characterization of the digraphs with a unique (reverse) arborescence. □

Proof of Proposition 5.5. The equivalence $(i) \iff (ii)$ is obvious. Given that $\mathcal{G} \subseteq \mathcal{S} \subseteq \mathcal{M}$, the equivalence $(ii) \iff (iii)$ is also obvious.

Next we prove the implication $(i) \implies (iv)$. Assume (i), i. e., $0 \in \mathcal{G}$. This means that 0 is accessible from all configurations. Therefore, by Corollary 5.15, all configurations are concentrated on the points of no return, i. e., all sites are points of no return, so there are no cyclic sites, which means $\mathfrak{X}$ is a DAG.

Assume now (iv), i. e., that $\mathfrak{X}$ is a DAG. Then all sites are points of no return and therefore, by Lemma 5.17, all configurations can be flushed, i. e., the 0 configuration is accessible from any configuration. □

Proof of Proposition 5.6. First we prove the lower bound in (b). If $\mathfrak{X}$ is a DAG then $|\mathcal{G}| = |\mathcal{M}| = \prod_{i \in V_0} \deg(i)$ while the number of edges is $|E| = \sum_{i \in V_0} \deg(i)$. Given the product N of an unspecified number of positive integers N_j, it is clear that the sum of the N_j is minimized when the N_j are the prime factors of N.

Next we construct such a DAG of minimum size for the case when G is cyclic. In this case, let $V_0 = \{1, \ldots, k\}$ and let $s = k+1$ denote the sink. Connect site $j \in V_0$ to $j+1 \in V$ by p_j edges. Now we have the relations $\sigma(\deg(j)\mathbf{t}_j) = \mathbf{t}_{j+1}$ where $\mathbf{t}_{k+1} = 0$. So $\mathbf{t}_1$ generates the cyclic group of order $|G|$.

For the general case, let $G = C(n_i) \times \cdots \times C(n_k)$ where $C(n)$ denotes the cyclic group of order n. Consider for each i the underlying DAG $\mathfrak{X}_i$ with sandpile group $C(n_i)$ constructed above. Take the sink-join of these underlying DAGs (Prop. 5.25). □

Proof of Proposition 5.7. Obviously, $(i) \iff (ii)$. The equivalence $(i) \iff (iii)$ is immediate from Def. 2.9. Next we show $(iii) \implies (iv)$. Let $\mathbf{x}$ be a nonzero configuration (not necessarily stable) from which 0 is accessible. Then, by Corollary 5.15, $\mathbf{x}$ is concentrated on the points of no return.

Conversely, assume (iv). Let i be a site which is a point of no return. Let $\mathbf{t}_i$ be the corresponding standard generator of $\mathbb{N}^{V_0}$ (a single grain on i, none elsewhere). Then by Lemma 5.17, the zero configuration is accessible from $\mathbf{t}_i$. □

Proof of Proposition 5.8. The equivalence $(i) \iff (ii)$ is obvious. Now, by Corollary 5.22 and Fact 2.17, the condition (i) that $|\mathcal{G}| = |\mathcal{M}| - 1$ is equivalent to saying that there is exactly one functional spanning subgraph in $\mathfrak{X}$ that is not a reverse arborescence. This last statement is easily seen to be equivalent to (iii). □

Proof of Proposition 5.9. In the light of the comment after Def. 5.10, we may assume that $\mathfrak{X}$ is a directed wheel graph. Let $0, 1, \ldots, n-1$ be the sites in this cyclic order along the cycle $\mathfrak{X}_0$. Clearly, $\sigma(\deg(i)\mathbf{t}_i) = \mathbf{t}_{i+1}$ (subscript mod n), so any $\mathbf{t}_i$ generates all the $\mathbf{t}_j$ which in turn generate $\mathcal{S}$ as a semigroup. Therefore $\mathcal{S}$ is a cyclic semigroup; so its homomorphic image $\mathcal{G}$ is also cyclic. (We note that in fact, in this case, $\mathcal{G} = \mathcal{S}$.)

We need to show that the saturated configuration $\mathbf{x}_{\max}$ is the identity of $\mathcal{G}$. A brief reflection will show that $\mathbf{t}_i \oplus \mathbf{x}_{\max} = \mathbf{t}_i$ (all sites will topple in succession around the cycle until we return a single grain to the starting site i). □

Proof of Proposition 2.24. The question is, when is $0 \in \mathcal{S}$.

By Proposition 5.7, $0 \in \mathcal{S}(\mathfrak{X})$ exactly if $\mathfrak{X}$ has a point of no return. Since the relevant reduction cannot create such a point, it follows that the only way that $\mathcal{S}(\widetilde{\mathfrak{X}}) \ncong \mathcal{S}(\mathfrak{X})$ happens is if $0 \in \mathcal{S}(\mathfrak{X})$ but $0 \notin \mathcal{S}(\widetilde{\mathfrak{X}})$, i. e., if $\mathfrak{X}$ has a point of no return but $\widetilde{\mathfrak{X}}$ has no such point. In other words, $\mathfrak{X}$ has an irrelevant point of no return and has no relevant points of no return. □

6 The lattice of idempotents: proof of the main results

6.1 Terminology: accessibility, quasiorder, preorder

In this section we introduce additional graph-theoretic and order-theoretic terminology.

A set with a reflexive and transitive relation is called a *preorder*. We say that a set with a transitive relation is a *quasi-order*[1].

Let $X = (\Omega, E)$ be a digraph (i. e., in the usage of this paper, a directed multigraph). Ω is the set of vertices (in some contexts called "states" or "configurations") and $\mathcal{E}$ is the set of edges (called "transitions" in the contexts mentioned). A *walk* of length k from vertex x to vertex y is a sequence $x = v_0, e_1, v_1, e_2, \dots, e_k, v_k = y$ of vertices $v_i \in \Omega$ and edges $e_i \in E$ such that v_{i-1} is the tail and v_i is the head of the edge e_i. A *cycle* of length k is a closed walk of length k ($v_k = v_0$) without repeated vertices (except for $v_k = v_0$. We say that vertex x is *cyclic* if it belongs to a cycle, *acyclic* otherwise.

We say that vertex y is *accessible* from vertex x if there is a walk from x to y. We say that y *properly accessible* from x if there is a walk of positive length from x to y. The only difference is that $x \in \Omega$ is always accessible from itself but x is properly accessible from itself exactly if x is cyclic.

Proper accessibility in X is a quasi-order on Ω; accessibility is a preorder.

For $A, B \subseteq \Omega$ we say that B is accessible from A if $(\exists x \in A)(\exists y \in B)(y \text{ is accessible from } x)$.

We denote by $X[A]$ the subgraph of X induced on A. With some abuse of terminology, we shall often identify the set A with the induced subgraph $X(A)$.

The *strong components* of X are the equivalence classes of the mutual accessibility relation on Ω. We call a strong component *cyclic* if it contains a cycle (under the aforementioned abuse of terminology); in this case all of its vertices are cyclic. The *acyclic* strong components are precisely the sets $\{v\}$ where v is an acyclic vertex.

The strong components of X form a *partially ordered set* with respect to accessibility. An *initial strong component* is a strong component which is not accessible from any other strong component.

Fact 6.1 *In a finite digraph, every vertex is accessible from an initial strong component.* □

For $A \subseteq \Omega$, the *ideal* $I(A)$ generated by A consists of all vertices accessible from A. We say that a subset $B \subseteq \Omega$ is an *ideal* in X if $B = I(B)$.

Definition 6.2 An ideal B is *normal* if all of its initial strong components (i. e., the initial strong components of the induced subgraph $X(B)$) are cyclic. We define the *normal ideal* $N(A)$ generated by A as the ideal generated by the strong components of $I(A)$.

Fact 6.3 *(a)* $I(A \cup B) = I(A) \cup I(B)$. *(b)* $N(A \cup B) = N(A) \cup N(B)$.

6.2 The proof

In this section we fix an underlying digraph $\mathfrak{X} = (V, E)$ and the corresponding sandpile model.

Definition 6.4 Let $\mathbf{x}, \mathbf{y} \in \mathbb{N}^{V_0}$ be configurations. We write $\mathbf{x} \vdash \mathbf{y}$ if $\mathbf{y}$ is obtained from $\mathbf{x}$ by toppling a node. This relation defines the *toppling digraph* on $\mathbb{N}^{V_0}$ (the vertices are the configurations; $\mathbf{x} \to \mathbf{y}$ is an edges if $\mathbf{x} \vdash \mathbf{y}$).

We use $\vdash\vdash$ to denote the reflexive, transitive closure of the $\vdash$ relation, i. e., we write $\mathbf{x} \vdash\vdash \mathbf{y}$ if $\mathbf{y}$ can be obtained from $\mathbf{x}$ by a sequence of topplings. We have $\mathbf{x} \vdash\vdash \mathbf{x}$ by definition, so $\vdash\vdash$ is a preorder on $\mathbb{N}^{V_0}$.

Note that for a stable configuration $\mathbf{y}$ we have $\mathbf{x} \vdash\vdash \mathbf{y}$ exactly if $\mathbf{y} = \sigma(\mathbf{x})$.

Recall from Definition 2.7 that we say that $\mathbf{y}$ is *accessible* from $\mathbf{x}$ if $\mathbf{y}$ is accessible from $\mathbf{x}$ in the *transition digraph* (Def. 2.6), i. e., if $\mathbf{y}$ can be obtained from $\mathbf{x}$ by adding some grains and then performing a sequence of topplings. We denote this circumstance by $\mathbf{x} \models \mathbf{y}$. In other words, $\mathbf{x} \models \mathbf{y}$ exactly if $(\exists \mathbf{z} \in \mathbb{N}^{V_0})(\mathbf{x} + \mathbf{z} \vdash\vdash \mathbf{y})$. We have $\mathbf{x} \models \mathbf{x}$, so $\models$ is a preorder on $\mathbb{N}^{V_0}$.

[1]This terminology is not standard. In the literature the terms "quasi-order" and "preorder" are used interchangeably for a reflexive and transitive relation (and sometimes for an irreflexive and transitive relation). Since the term "quasi-order" is used rather infrequently, we took the liberty of borrowing it for a transitive relation.

Notation 6.5 Let $\mathbf{x} \in \mathbb{N}^{V_0}$ be a configuration. We denote by $\iota(\mathbf{x})$ the ideal in $\mathfrak{X}$ generated by $\mathrm{supp}(\mathbf{x})$ (see Def. 5.14) in $\mathfrak{X}$. We denote by $\nu(\mathbf{x})$ the normal ideal generated by $\mathrm{supp}(\mathbf{x})$ (Def 6.2).

Lemma 6.6 *Let $\mathbf{x}, \mathbf{y} \in \mathbb{N}^{V_0}$ be configuraions.*

(a) If $\mathbf{x} \vdash \mathbf{y}$ then $\nu(\mathbf{x}) = \nu(\mathbf{y})$.

(b) If $\mathbf{x} \models \mathbf{y}$ then and $\nu(\mathbf{x}) \subseteq \nu(\mathbf{y})$.

(c) If $\mathbf{x}, \mathbf{y} \in \mathcal{M}$ (stable configurations) then $\nu(\mathbf{x} \oplus \mathbf{y}) = \nu(\mathbf{x}) \cup \nu(\mathbf{y})$.

Proof: (a) It suffices to prove this in the case when $\mathbf{x} \vdash \mathbf{y}$. The relation $\nu(\mathbf{y}) \subseteq \nu(\mathbf{x})$ is obvious: every vertex in $\mathrm{supp}(\mathbf{y})$ is accessible from $\mathrm{supp}(\mathbf{x})$. The converse is false, but if C is a cycle, accessible from $\mathrm{supp}(\mathbf{x})$ then it remains accessible from $\mathrm{supp}(\mathbf{y})$. This is immediate if $C \cap \mathrm{supp}(\mathbf{x}) = \emptyset$ and it follows from Lemma 5.13 if $C \cap \mathrm{supp}(\mathbf{x}) \neq \emptyset$. An application of this statement to the cycles of the initial strong components yields the inclusion $\nu(\mathbf{x}) \subseteq \nu(\mathbf{y})$.
(b) We have $\mathbf{x} + \mathbf{z} \vdash\vdash \mathbf{y}$ and $\mathrm{supp}(\mathbf{x}) \subseteq \mathrm{supp}(\mathbf{x} + \mathbf{z})$. Therefore $\nu(\mathbf{x}) \subseteq \nu(\mathbf{x} + \mathbf{z}) = \nu(\mathbf{y})$. (We only used the trivial half of (a).)
(c) We have $\mathrm{supp}(\mathbf{x} + \mathbf{y}) = \mathrm{supp}(\mathbf{x}) \cup \mathrm{supp}(\mathbf{y})$ and $\mathbf{x} + \mathbf{y} \vdash\vdash \mathbf{x} \oplus \mathbf{y}$. Now the result follows from (a). □

Lemma 6.7 *Let $\mathbf{x} \in \mathbb{N}^{V_0}$. There exists $n > 0$ such that $\sigma(n\mathbf{x})$ restricted to $\iota(\mathbf{x})$ is recurrent with respect to the underlying digraph $\mathfrak{X}[\iota(\mathbf{x})]$, i. e., $\sigma(n\mathbf{x})$ restricted to $\iota(\mathbf{x})$ belongs to $\mathcal{G}(\mathfrak{X}[\iota(\mathbf{x})])$.*

Proof: Let $C_1, C_2, \ldots, C_m$ be the initial strong components of $\iota(\mathbf{x})$. Note that each C_i intersects $\mathrm{supp}(\mathbf{x})$; let $u_i \in C_i \cap \mathrm{supp}(\mathbf{x})$. Every vertex in $\iota(\mathbf{x})$ is reachable from some u_i. By applying Lemma 5.19 to each u_i we find an $n > 0$ and a (semistable) configuration $\mathbf{y}$ on $\mathfrak{X}[\iota(\mathbf{x})])$ such that $\sigma(n\mathbf{x}) = \sigma(\mathbf{y})$ and $\mathbf{y}(v) \geq \deg(v) - 1$ for all vertices $v \in \iota(\mathbf{x})$. So $\mathbf{y}$ is accessible from the max configuration $\mathbf{x}_{\max}$ and therefore $\sigma(\mathbf{y}) = \sigma(n\mathbf{x})$ is recurrent. □

Recall that $\mathcal{L}$ denotes the lattice of idempotents of $\mathcal{M}$.

Lemma 6.8 *Let $\mathbf{e} \in \mathcal{L}$ be an idempotent in $\mathcal{M}$. Then $\iota(\mathbf{e})$ is a normal ideal in $\mathfrak{X}$, i. e., $\nu(\mathbf{e}) = \iota(\mathbf{e})$.*

Proof: Let C be an initial strong component of $\iota(\mathbf{e})$. We need to show that C is cyclic. Assume, for a contradition, that $C = \{v\}$ for some acyclic vertex v. The idempotent satisfies $\mathbf{e} \oplus \mathbf{e} = \mathbf{e}$ and by assumption $\mathbf{e}(v) \neq 0$. If $2\mathbf{e}(v) < \deg(v)$, then $(\mathbf{e} \oplus \mathbf{e})(v) = 2\mathbf{e}(v) \neq \mathbf{e}(v)$. If $2\mathbf{e}(v) \geq \deg(v)$ then $(\mathbf{e} \oplus \mathbf{e})(v) = 2\mathbf{e}(v) - \deg(v) > \mathbf{e}(v)$, again contradicting $(\mathbf{e} \oplus \mathbf{e})(v) = \mathbf{e}(v)$. □

Lemma 6.9 *Let $\mathbf{e} \in \mathcal{L}$. Then $\mathbf{e}$ restricted to $\nu(\mathbf{e})$ is recurrent with respect to the underlying digraph $\mathfrak{X}[\nu(\mathbf{e})]$, i. e., $\mathbf{e}$ restricted to $\nu(\mathbf{e})$ belongs to $\mathcal{G}(\mathfrak{X}[\nu(\mathbf{e})])$.*

Proof: By Lemma 6.8, $\iota(\mathbf{e}) = \nu(\mathbf{e})$. Now the result follows from Lemma 6.7 and the fact that $(\forall n > 0)(\sigma(n\mathbf{e}) = \mathbf{e})$. □

Corollary 6.10 *Let $\mathbf{e} \in \mathcal{L}$. Then $\mathbf{e}$ restricted to $\nu(\mathbf{e})$ is the identity in $\mathcal{G}(\mathfrak{X}[\nu(\mathbf{e})])$, the sandpile group associated with $\mathfrak{X}[\nu(\mathbf{e})]$.*

Proof: $\mathbf{e}$ is a recurrent idempotent in $\mathcal{M}(\mathfrak{X}[\nu(\mathbf{e})])$ by Lemma 6.9. □

We are now ready to establish our bijection between the set of idempotents of $\mathcal{M}$ and the set of normal ideals of $\mathfrak{X}$.

Proposition 6.11 *The map ν restricted to $\mathcal{L}$ is a bijection between $\mathcal{L}$ and the set of normal ideals of $\mathfrak{X}$.*

Proof: First we prove that ν maps $\mathcal{L}$ onto the set of normal ideals of $\mathfrak{X}$.

Let A be a normal ideal in $\mathfrak{X}$. Let $\mathbf{e}$ be the identity configuration in the sandpile group $\mathcal{G}(\mathfrak{X}[A])$. Let C be an initial strong component of A. Since $\mathbf{e}$ is recurrent in $\mathcal{M}(\mathfrak{X}[A])$, by Lemma 5.13, we have $\text{supp}(\mathbf{e}) \cap C \neq \emptyset$. Let $\mathbf{e}'$ be the extension of $\mathbf{e}$ to V_0 by setting $\mathbf{e}(i) = 0$ for all $i \notin A$. Then $\mathbf{e}'$ is an idempotent in $\mathcal{M}$ and $\nu(\mathbf{e}') = A$.

Next we prove that ν is 1-to-1. Suppose $A = \nu(\mathbf{e})$. We need to show that A determines $\mathbf{e}$. Indeed, by Corollary 6.10, $\mathbf{e}$, when restricted to A, is equal to the identity configuration in $\mathcal{G}(\mathfrak{X}[A])$. For $i \notin A$ we have $\mathbf{e}(i) = 0$ by Lemma 6.8. □

Corollary 6.12 *The map ν restricted to $\mathcal{L}$ is an isomorphism between the semilattice $\mathcal{L}$ and the semilattice of normal ideals of $\mathfrak{X}$ under the union operation.*

Proof: Combine Prop. 6.11 and part (c) of Prop. 6.6. □

Let $\mathcal{C}$ denote the set of cyclic strong components of $\mathfrak{X}$, viewed as a poset under accessibility. Ideals of $\mathcal{C}$ are subsets closed under accessibility.

For an ideal A of $\mathfrak{X}$, define $\mu(A) \subseteq \mathcal{C}$ to be the set of cyclic strong components in A. This is clearly an ideal of $\mathcal{C}$.

Proposition 6.13 *μ provides an isomorphism between the semilattice of normal ideals of $\mathfrak{X}$ under union and the semilattice of ideals of $\mathcal{C}$ under union.*

Proof: Obvious. □

The proof of our main result, Theorem 4.13, now follows. The dual isomorphism of lattices stated in the Theorem translates to the isomorphims of the semilattice of idempotents of $\mathcal{M}$ and the union-semilattice of ideals of the accessibility poset of $\mathcal{C}$. This isomorphism is established by the composition $\mu\nu$, combining Corollary 6.12 and Proposition 6.13. □

Corollary 6.14 *The semilattice $\mathcal{L}$ of idempotents of $\mathcal{M}$ defines a distributive lattice.*

Proof: Suppose A, B are ideals in $\mathcal{C}$. Then $A \cup B$ and $A \cap B$ are ideals as well; these are the "join" and "meet" in the lattice defined by the accessibility poset on $\mathcal{C}$. These operations distribute over each other. □

Remark 6.15 Note that in general, the intersection of normal ideals in $\mathfrak{X}$ is *not* a normal ideal. This is why the extra step after Cor. 6.12 was needed to establish distributivity.

Proposition 6.16 *Every finite distributive lattice $\mathcal{L}$ is isomorphic to the sandpile lattice $\mathcal{L}(\mathfrak{X})$ for some underlying digraph $\mathfrak{X}$.*

Proof: By Birkhoff's representation theorem (see Theorem 15 in [11], Chapter XIV, Section 6, p. 487), $\mathcal{L}$ is isomorphic to the semilattice of ideals of a partial order $\mathcal{P} = (V_0, \preccurlyeq)$. Define the digraph $\mathfrak{X}_0$ to have V_0 as its vertex set and an edge from vertex v to u if $u \preccurlyeq v$. Let V_0 be the set of sites. Define $\mathfrak{X}$ by adding a sink to $\mathfrak{X}_0$ and adding an edge from every site to the sink. □

To prove Theorem 4.8, we combine Proposition 6.16 and Corollary 6.14.

The characterization of underlying digraphs for which there is a unique idempotent in the sandpile semigroup follows.

Proof of Theorem 4.15.

If $\mathcal{S}$ has a unique idempotent then $\mathcal{M}$ has at most two and therefore, by Theorem 4.13, $|\mathcal{C}| \le 1$.

Clearly, $\mathcal{C} = \emptyset \iff \mathfrak{X}$ is a DAG. The latter is equivalent to $\mathcal{M} = \mathcal{S} = \mathcal{G}$ (Prop. 5.5). But a group has a unique idempotent.

$|\mathcal{C}| = 1$ means $\mathfrak{X}$ has a single cyclic strong component. According to Theorem 4.13, this is also equivalent to $\mathcal{M}$ having exactly two idempotents. One of these (0) must not belong to $\mathcal{S}$; the necessary and sufficient condition of this is that $\mathfrak{X}$ has no point of no return (Prop. 5.7). □

7 Transience class and the Sandpile Quotient

In this section we introduce the *sandpile quotient* as the Rees quotient $\mathcal{S}/\mathcal{G}$; the transience class $\mathrm{tcl}(\mathfrak{X})$ as the maximum number of grains in a configuration of which the stabilization is transient; and show that the transience class is one less than the nilpotence class of the sandpile quotient (Theorem 7.12). We also observe that these quantities are finite if and only if $\mathcal{S}$ has a unique idempotent (Prop. 7.3), a situation that was characterized in Theorem 4.15.

7.1 Rees quotient, nilpotence, unique idempotents

Definition 7.1 Let $\mathcal{S}$ be a semigroup. An element $z \in \mathcal{S}$ is a *zero* in $\mathcal{S}$ if $xz = zx = z$ for all $x \in \mathcal{S}$. The zero, if it exists, is clearly unique. We say that $\mathcal{S}$ is *nilpotent* if there exists a positive integer k such that $\mathcal{S}^k = \{z\}$ (all products of length k are zero). The smallest such k is the *nilpotence class* $\mathrm{cl}(\mathcal{S})$. If $\mathcal{S}$ is not nilpotent, we set $\mathrm{cl}(\mathcal{S}) = \infty$.

A detailed structure theory of the finite commutative nilpotent semigroups is given in [19].

Definition 7.2 (Rees quotient) Let $\mathcal{S}$ be a semigroup and $I \subset \mathcal{S}$ an ideal. The *Rees congruence* (denoted by $\sim_I$) is defined as follows: $a \sim_I b$ if $a = b$ or $a, b \in I$ (cf. [18] p. 16). The *Rees quotient* $\mathcal{S}/I$ is the quotient of $\mathcal{S}$ by this congruence, i. e., we contract I to a single element. This element will necessarily be a zero in the Rees quotient. We say that $\mathcal{S}$ is an *ideal extension* of the semigroup I by $\mathcal{S}/I$, a semigroup with zero (cf. [18] p. 54). We say that $\mathcal{S}$ is a *nilpotent ideal extension* of I by $\mathcal{S}/I$ if the Rees quotient $\mathcal{S}/I$ is nilpotent.

Having a unique idempotent is a strong structural constraint on a semigroup.

Proposition 7.3 *For a nonempty finite commutative semigroup $\mathcal{S}$ generated by a set T, the following are equivalent:*
(i) Every element of $\mathcal{S}$ has a power which is fully accessible.
(ii) Every element of T has a power which is fully accessible.
(iii) The Rees quotient of $\mathcal{S}$ by its kernel $\mathcal{G}(\mathcal{S})$ is nilpotent, i. e., $\mathcal{S}$ is a nilpotent ideal extension of a group.
(iv) $\mathcal{S}$ has a unique idempotent.

Proof: Clearly, (i) implies (ii). To see that (ii) implies (iii), recall that the kernel $\mathcal{G} = \mathcal{G}(\mathcal{S})$ consists of the fully accessible elements. So in the Rees quotient $\mathcal{S}/\mathcal{G}$, every generator is nilpotent. By commutativity it follows that every element of $\mathcal{S}/\mathcal{G}$ is nilpotent. Finally, we claim that every product of length $n+1$ in $\mathcal{S}/\mathcal{G}$ is is zero, where $n = |\mathcal{S}/\mathcal{G}|$. Take such a product, $s_0 \dots s_n$ and consider the prefixes $p_i = s_0 \dots s_i$. By the pigeon hole principle, $p_i = p_j$ for some $i < j$. But $p_j = p_i x$ for some $x \in \mathcal{S}/\mathcal{G}$, so by induction $p_i = p_i x^r$ for all $r \ge 0$. Some power x^t is zero hence $p_i = p_i \cdot 0 = 0$ and thus $p_n = p_i y = 0$. Assume (iii) now and let x be an idempotent. Now, by the nilpotence of $\mathcal{S}/\mathcal{G}$, some power of x belongs to $\mathcal{G}$ and therefore x belongs to $\mathcal{G}$. But $\mathcal{G}$ is a group, so x can only be the identity element of $\mathcal{G}$, proving (iv). Finally, let us assume (iv); let y be the unique idempotent in $\mathcal{S}$. Since every element of $\mathcal{S}$ has a power which is an idempotent, y is fully accessible, proving (i). □

Definition 7.4 (Sandpile quotient) The *sandpile quotient* associated with the underlying digraph $\mathfrak{X}$ is the Rees quotient $\mathcal{S}/\mathcal{G}$ where $\mathcal{S} = \mathcal{S}(\mathfrak{X})$ and $\mathcal{G} = \mathcal{G}(\mathfrak{X})$ are the sandpile semigroup and the sandpile group, resp.

The case when the sandpile semigroup has a unique idempotent (and therefore the sandpile quotient is nilpotent) has been characterized in Theorem 4.15.

7.2 Transience class

Definition 7.5 (Transience class) We define the *transience class* $\mathrm{tcl}(\mathfrak{X})$ of the sandpile model (or of the underlying digraph $\mathfrak{X}$) as the greatest integer t such that there exists a (not necessarily stable) configuration $\mathbf{x}$ of weight t (Def. 5.11) such that the stabilization $\sigma(\mathbf{x})$ is transient. Set $\mathrm{tcl}(\mathfrak{X}) = \infty$ if no such integer exists.

We shall compare $\mathrm{tcl}(\mathfrak{X})$ with $\mathrm{cl}(\mathcal{S}/\mathcal{G})$; we shall see that

$$\mathrm{tcl}(\mathfrak{X}) = \mathrm{cl}(\mathcal{S}/\mathcal{G}) - 1, \tag{8}$$

where $\mathcal{S}/\mathcal{G}$ is the sandpile quotient associated with $\mathfrak{X}$ (Def. 7.4) and "cl" denotes the nilpotence class (Def. 7.1).

Proposition 7.6 $\mathrm{tcl}(\mathfrak{X}) = 0$ *if and only if* $\mathfrak{X}$ *is a DAG.*

Proof: $\mathrm{tcl}(\mathfrak{X}) = 0$ if and only if all stable configurations are recurrent. But this happens exactly when $\mathfrak{X}$ is a DAG (Prop. 5.5). □

Proposition 7.7 *If* $\mathfrak{X}$ *is a DAG then* $\mathrm{cl}(\mathcal{S}/\mathcal{G}) = 1$ *and* $\mathrm{tcl}(\mathfrak{X}) = 0$*. In particular, Eqn.* (8) *holds.*

Proof: In this case, $\mathcal{M} = \mathcal{S} = \mathcal{G}$ (Prop. 5.5), therefore $\mathrm{cl}(\mathcal{S}/\mathcal{G}) = 1$. On the other hand, $\mathrm{tcl}(\mathfrak{X}) = 0$ by Prop. 7.6. □

Irrelevant sites will give us some headache, so this is a good time for the reader to review the construction of the *relevant reduction* $\widetilde{\mathfrak{X}}$ of $\mathfrak{X}$ which eliminates irrelevant sites (Prop. 2.21).

Proposition 7.8 *If* $\mathfrak{X}$ *has no irrelevant points of no return then*

$$\mathrm{tcl}(\widetilde{\mathfrak{X}}) = \mathrm{tcl}(\mathfrak{X}), \tag{9}$$

where $\widetilde{\mathfrak{X}}$ *is the relevant reduction of* $\mathfrak{X}$*.*

Proof: The weight from an irrelevant site i can be transfered to the nearest relevant site along a directed path. Such a site exists because i is not a point of no return. □

Proposition 7.9 *If* $\mathfrak{X}$ *has no irrelevant points of no return then*

$$\mathcal{S}(\widetilde{\mathfrak{X}})/\mathcal{G}(\widetilde{\mathfrak{X}}) \cong \mathcal{S}(\mathfrak{X})/\mathcal{G}(\mathfrak{X}). \tag{10}$$

Proof: By Prop. 2.24, if $\mathfrak{X}$ has no irrelevant points of no return then $\mathcal{S}(\widetilde{\mathfrak{X}}) \cong \mathcal{S}(\mathfrak{X})$. The statement follows by noting that $\mathcal{G}$ is the unique minimal ideal of $\mathcal{S}$. □

Next we give a sufficient condition for $\mathrm{tcl}(\mathfrak{X}) = \infty$. As we shall see later, this condition is also necessary (Cor. 7.13).

Proposition 7.10 *If* $\mathfrak{X}$ *has a cycle* C *and a site from which* C *is not accessible then* $\mathrm{cl}(\mathfrak{X}) = \mathrm{tcl}(\mathfrak{X}) = \infty$*.*

Proof: Let i be a site from which C is not accessible. Then C will be blank in the configuration $\sigma(w\mathbf{t}_i)$, and therefore $\sigma(w\mathbf{t}_i)$ will be transient, for all $w \geq 0$. Hence $\mathrm{tcl}(\mathfrak{X}) = \infty$.

On the other hand, the assumption is precisely the negation of the part (a) of Theorem 4.15, so $\mathcal{S}$ has at least two idempotents and therefore $\mathrm{cl}(\mathcal{S}/\mathcal{G}) = \infty$ (Prop. 7.3). □

The essence of Eqn. (8) is in the following statement.

Proposition 7.11 *Suppose $\mathfrak{X}$ has no irrelevant points of no return. Then Equation* (8) *holds.*

Proof: By Propositions 7.8 and 7.9 we may assume $\mathfrak{X} = \widetilde{\mathfrak{X}}$ and therefore $\mathfrak{X}$ has no irrelevant sites, i. e., every site has outdegree ≥ 2.

Suppose there is a configuration $\mathbf{x}$ of weight w such that $\sigma(\mathbf{x})$ is transient. Now $\mathbf{x} = \sum_{i \in V_0} \mathbf{x}(i)\mathbf{t}_i$ (cf. Eqn. (1)). Therefore $\sigma(\mathbf{x})$ is the sum, in $\mathcal{M}$, of $w = \sum_{i \in V_0} \mathbf{x}(i)$ stable nonzero configurations $\mathbf{t}_i \in \mathcal{S}$. ($\mathbf{t}_i$ is stable because the site i is relevant.) This implies that $\mathrm{cl}(\mathcal{S}/\mathcal{G}) \geq w + 1$. So we proved $\mathrm{cl}(\mathcal{S}/\mathcal{G}) \geq \mathrm{tcl}(\mathfrak{X}) + 1$.

Conversely, suppose the nonnegative integer w is less than the nilpotence class $\mathrm{cl}(\mathcal{S}/\mathcal{G})$. This means that there is a sum $\mathbf{x} = \sum_{j=1}^{w} \mathbf{y}_j$ of (not necessarily distinct) nonzero stable configurations $\mathbf{y}_j$ such that $\sigma(\mathbf{x})$ is transient. Since $\mathrm{weight}(\mathbf{y}_j) \geq 1$ for each j, we infer that $\mathrm{weight}(\mathbf{x}) \geq w$. This proves $\mathrm{cl}(\mathcal{S}/\mathcal{G}) \leq \mathrm{tcl}(\mathfrak{X})+1$. □

Now we come to the main result of this section.

Theorem 7.12 *Equation* (8) *holds for all underlying digraphs.*

Proof: If $\mathfrak{X}$ is a DAG then Equation (8) holds by Prop. 7.7. If $\mathfrak{X}$ is not a DAG and has a point of no return then Equation (8) holds by Prop. 7.10.

In the remaining cases, $\mathfrak{X}$ has no points of no return. Equation (8) now follows from Prop. 7.11. □

From the results above we extract the characterization of those underlying digraphs which have finite transience class.

Corollary 7.13 *For the underlying digraph $\mathfrak{X}$, the following are equivalent.*

(a) Either $\mathfrak{X}$ is a DAG or $\mathfrak{X}$ has exactly one cyclic strong component and this strong component is accessible from all sites;

(b) the sandpile semigroup $\mathcal{S}$ has a unique idempotent;

(c) the sandpile quotient is nilpotent;

(d) the transience class of $\mathfrak{X}$ in finite.

Proof: (a) and (b) are equivalent by Theorem 4.15. (b) and (c)are equivalent by Prop. 7.3. (c) and (d) are equivalent by Theorem 7.12. □

We mention an important sufficient condition for the finiteness of the transience class.

Corollary 7.14 *If the site-digraph $\mathfrak{X}_0$ is strongly connected then* $\mathrm{tcl}(\mathfrak{X})$ *is finite.* □

7.3 Bounded transience class

In a subsequent paper [4] we desrcibe the asymptotic structure of the underlying digraphs with bounded transience class. Surprisingly, we find that the transience class is bounded if and only if the number of transient configurations is bounded. As a preview, we state a corollary to the main structure theorem from [4].

Theorem 7.15 *There exist functions ψ_1 and ψ_2 such that if the transience class of the underlying digraph $\mathfrak{X}$ is k then the number of transient configurations is $\leq \psi_1(k)$; and if $\mathfrak{X}$ is not a DAG then the sandpile group $\mathcal{G}$ contains a cyclic subgroup of index $\leq \psi_2(k)$.*

The structure of these underlying digraphs is close to the "directed wheel" described in Def. 5.10.

8 Open questions

8.1 Digraph size for a given sandpile group

For a given finite abelian group G, let $e(G)$ denote the smallest size (number of edges) of an underlying digraph graph $\mathfrak{X}$ such that $\mathcal{G}(\mathfrak{X}) \cong G$.

If G is a direct product, $G = G_1 \times \cdots \times G_k$, then $e(G) \leq \sum_{i=1}^{k} e(G_i)$ by Prop. 5.25.

So it is worth focusing on cyclic groups, even though the decomposition to cyclic groups may not give the best bound on $e(G)$.

Let $e(n)$ denote the value $e(G)$ when G is the cyclic group of order n.

A number is called "smooth" if all its prime divisors are "small." By Prop. 5.6, $e(G) \leq \sum p_j$ where the p_j are the (not necessarily distinct) prime divisors of $|G|$ by Prop. 5.6.

So $e(n)$ will be "small" if n is smooth. Moreover, $e(n)$ will also be small if $n+1$ is smooth, by Propositions 5.8 and 5.9. Specifically, the sandpile group of a directed wheel graph is cyclic of order $\prod_{i \in V_0} \deg(i) - 1$. This number can be a large prime for small degrees. For instance if n is a Mersenne prime, $n = 2^p - 1$, then $e(n) \leq 2p = 2\log_2(n+1)$.

It would be interesting to find lower bounds on $e(n)$. Two specific question:

Question 8.1 (a) Does there exist $c > 0$ such that for infinitely many values of n we have $e(n) \geq n^c$?

(b) Is it true that for every $\epsilon > 0$, for almost all values of n we have $e(n) \leq n^\epsilon$?

By Prop 5.25, w.l.o.g. we can restrict Question (a) to prime values of n.

Question 8.2 Can one characterize those finite abelian groups which arise as the sandpile groups of underlying digraphs with (a) weakly connected (b) strongly connected site digraphs $\mathfrak{X}_0$?

8.2 Transience class

In this section we consider strongly connected site-graphs $\mathfrak{X}_0$; for these, the transience class is finite (Cor. 7.14).

Consider the following process: we start from the empty configuration, add some grains, stabilize, repeat. Since the transience class is finite, eventually we shall reach a recurrent configuration (and then stay recurrent forever).

Interesting questions arise in the study of the rate of growth of the transience class for various families of underlying digraphs. If the transience class is very large, this puts the physical reality of "eventual recurrence" of the system into question. In [3] we exhibit an infinite family of undirected underlying graphs with exponentially large transience class: $\mathrm{tcl}(\mathfrak{X}) = \Theta(\phi^m)$ where m is the number of edges and $\phi = (1+\sqrt{5})/2 \approx 1.618$ is the golden ratio. The question arises, is this the largest possible rate of growth?

Question 8.3 Is there an infinite family of undirected sandpile models $\mathfrak{X}$ and a constant $c > \phi$ such that $\mathrm{tcl}(\mathfrak{X}) > c^m$ where m is the number of edges of $\mathfrak{X}$ and m is unbounded over the given family?

The transience class of the square grid sandpile model, of great interest to statistical physics, has been studied in [2]. The site graph of the model is an $n \times n$ grid; each site on the boundary is connected to the sink by one or two edges so that each site has degree 4 (so each corner is connected to the sink by two edges, all other sites on the boundary by one edge each). We found that the transience class of this model is polynomially bounded ($O(n^C)$) for some large constant C; we expect the actual constant to be less than 5.

For a site i, let $\mathrm{tcl}_i(\mathfrak{X})$ denote the largest k such that $\sigma(k\mathbf{t}_i)$ is transient, i. e., putting k grains on site i and stabilizing we obtain a transient configuration. Call this the transience class of site i. Call $\mathrm{ltcl}(\mathfrak{X}) := \max_{i\in V_0} \mathrm{tcl}_i(\mathfrak{X})$ the *local transience class*; call the transience class "global" to distinguish it from the local transience class. It is clear that

$$\mathrm{ltcl}(\mathfrak{X}) \leq \mathrm{tcl}(\mathfrak{X}) \leq \sum_{i\in V_0} \mathrm{tcl}_i(\mathfrak{X}) \leq |V_0|\,\mathrm{ltcl}(\mathfrak{X}). \tag{11}$$

Question 8.4 Let $\mathfrak{X}$ be the $n \times n$ grid sandpile model. Is it true that $\mathrm{tcl}(\mathfrak{X}) = \mathrm{tcl}_i(\mathfrak{X})$ where i is a corner of the square?

It is natural to suspect (as the senior author of this paper did) that for undirected models, the lower bound in inequality (11) should be tight. This was disproved by Zoran Sunic who showed that there are undirected sandpile models for which the ratio of the global and local transience classes is 2. We include his results and comments with his kind permission [30].

Proposition 8.5 (Sunic) *For every k there exists an undirected sandpile model with $\mathrm{ltcl}(\mathfrak{X}) = k$ and $\mathrm{tcl}(\mathfrak{X}) = 2k$.*

Proof: Let $V = \{0, 1, 2\}$ with 0 the sink. Put $k + 1$ undirected edges between the two sites, and add one edge from site 1 to the sink. Now the configuration that achieves the transience class $2k$ is $k(\mathbf{t}_1 + \mathbf{t}_2)$; but $\sigma((k + 1)\mathbf{t}_i)$ is recurrent for $i = 1, 2$, so $\mathrm{tcl}_i = k$. □

Question 8.6 (Sunic) Is it the case that for all undirected sandpile models $\mathfrak{X}$ we have $\mathrm{tcl}(\mathfrak{X}) \leq 2\,\mathrm{ltcl}(\mathfrak{X})$?

The answer is negative if we omit the condition that the model be undirected; in that case the ratio $\mathrm{tcl}(\mathfrak{X})/\max_{i\in V_0} \mathrm{tcl}_i(\mathfrak{X})$ is unbounded. In fact, things get as bad as possible: for arbitrarily large V_0, the global/local ration can be as large as $|V_0|$.

Proposition 8.7 (Sunic) *For every k and n there exists an underlying digraph $\mathfrak{X}$ with strongly connected site digraph such that $\mathfrak{X}$ has n sites; the transience class of each site is k; and the global transience class is kn.*

Proof: Sunic's example is a "fat cycle" with a link to the sink. The sites are $\{0, 1, \ldots, n-1\}$; there are $k+1$ edges directed from site i to site $i + 1 \bmod n$ (i. e., $k + 1$ edges from $n - 1$ to 0), and one edge goes from site 0 to the sink. Now $\mathrm{tcl}(\mathfrak{X}) = kn$, attained by the configuration $k \sum_{i=0}^{n-1} \mathbf{t}_i$, while for each i, the configuration $\sigma((k + 1)\mathbf{t}_i)$ is recurrent, so $\mathrm{tcl}(i) = k$. □

9 Appendix

For completeness and easy access, we include some of the simple and elegant proofs skipped in Sections 1 and 3.

9.1 Avalanches terminate

Proof of Fact 2.3. Let r_i denote the directed distance from vertex i to the sink. Let $D = \max_{i\in V_0} \deg(i)$. Let us start from configuration $\mathbf{x}_0$; let m be the initial number of grains of sand, i. e., $m = \sum_{i\in V_0} \mathbf{x}_0(i)$.

Only for the purposes of this proof, for any configuration $\mathbf{x}$ that has at most m grains of sand, set $\mathbf{x}(s) = m - \sum_{i\in V_0} \mathbf{x}(i)$, where s is the sink. Associate the following "potential" with $\mathbf{x}$:

$$P(\mathbf{x}) = \sum_{i\in V} \mathbf{x}(i) D^{n-r_i}. \tag{12}$$

Note that the summation extends over *all vertices*, including the sink.

It is clear that every toppling strictly increases the potential: at least one grain of sand gets strictly closer to the sink (including possibly falling into it), and the D-fold increase of that grain's contribution to the potential cannot be compensated by the decrease caused by grains that go farther from the sink. On the other hand, the potential is an integer bounded by mD^n (the potential of the empty configuration). So the sequence of topplings must terminate. □

9.2 Avalanches terminate in the same configuration

In this section we prove Facts 2.4 and 2.5. The proof is based on a "Diamond Lemma" argument that follows the lines of the standard proof of the Jordan-Hölder theorem in group theory.

We say that a walk from a node x in a digraph is *maximal* if it cannot be continued (forward). Such a walk may be finite or infinite; if it is finite, it ends at a sink (vertex of outdegree 0). We permit directed cycles; an infinite walk can enter a directed cycle and keep cycling in it forever.

Let us call a digraph *confluent* if for every node v either all maximal walks starting at v are infinite or all of them terminate at the same node in the same number of steps.

Let us say that a digraph satisfies the *diamond condition* if for every triple x, y, z of nodes, if there are edges $x \to y$ and $x \to z$ then there exists a node w and edges $y \to w$ and $z \to w$.

Lemma 9.1 *If a digraph satisfies the diamond condition then it is confluent.*

Proof: Let k be the length of the shortest maximal walk starting a v. If $k = \infty$, we are done. For finite k we prove the statement by induction on k.

For $k = 0$ there is nothing to prove; and for $k = 1$, the diamond condition implies that the outdegree of v is 1 and the one-step walk reaches a sink.

Assume $k \geq 2$. Consider two maximal walks X and Y starting at v; let $X = (x_0, x_1, \dots, x_k)$ and $Y = (y_0, y_1, \dots)$ where $x_0 = y_0 = v$; the walk X has k steps and ends in a sink; Y may *a priori* be finite or infinite.

If $x_1 = y_1$ then we apply the inductive hypothesis to the portions of the two walks starting at x_1.

Assume $x_1 \neq y_1$. Apply the diamond condition to find a vertex z such that there are edges $x_1 \to z$ and $y_1 \to z$. Let $Z = (z_0, z_1, \dots)$ be a maximal walk, where $z_0 = z$.

Now consider the two maximal walks $(x_1, x_2, \dots, x_k)$ and $(x_1, z_0, z_1, \dots)$. Both walk start from x_1. The first of these has length $k - 1$, so by the inductive hypothesis, Z ends at $z_{k-2} = x_k$. Now consider the walks $(y_1, z_0, z_1, \dots, z_{k-2})$ and $(y_1, y_2, \dots)$. By the same argument, Y terminates at $y_k = z_{k-2}$. □

Proof of Fact 2.4. Apply Lemma 9.1 to the toppling graph (see Def. 6.4). The toppling graph obviously satisfies the diamond condition, and all maximal walks are finite by Fact 2.3 (avalanches terminate). □

Proof of Fact 2.5. Define the *layered toppling graph* $\mathcal{LT}$ on the vertex set $\Omega = \mathbb{N}^{V_0} \times \mathbb{N}^{V_0}$ as follows. We put an edge from $(\mathbf{y}, \mathbf{z}) \in \Omega$ to $(\mathbf{y}', \mathbf{z}') \in \Omega$ if $\mathbf{y}'$ is obtained from $\mathbf{y}$ by toppling a site, say $i \in V_0$, and $\mathbf{z}' = \mathbf{z} + \mathbf{t}_i$ (we add 1 in the i-th coordinate). It is clear that the diamond condition holds in $\mathcal{LT}$ and all walks terminate. Now a walk from $(\mathbf{x}, 0)$ to $(\mathbf{y}, \mathbf{z})$ corresponds to a toppling sequence from $\mathbf{x}$ to $\mathbf{y}$ with score $\mathbf{z}$. Confluence of the walks starting at $(\mathbf{x}, 0)$ thus means that all avalanches starting at $\mathbf{x}$ have the same score. It follows that all toppling sequences from $\mathbf{x}$ to any (not necessarily stable) configuration $\mathbf{y}$ also has the same score since such a sequence can be completed to an avalanche by adding an avalanche starting at $\mathbf{y}$. □

Note that the same "layered digraph" trick can be used to deduce the Jordan-Hölder Theorem from Lemma 9.1.

We remark that the same proof works for the more general "chip firing" model where infinite toppling sequences are possible; the result then says that if, from a given initial configuration, there exists a terminating maximal toppling sequence, then all maximal toppling sequences terminate in the same number of steps and in the same final configuration and have the same score ([13]).

9.3 Commutative monoids

We prove Proposition 3.1. First we observe that the intersection of two nonempty ideals I and J in a semigroup $\mathcal{S}$ is a nonempty ideal since $I \cap J \supseteq IJ$. The following is now immediate.

Fact 9.2 *If $\mathcal{S}$ is a nonempty finite semigroup then the kernel of $\mathcal{S}$ is nonempty and therefore it is the unique minimal ideal. (Commutativity was not needed here.)* □

Now we come to the key observation underlying Fact 2.18.

Fact 9.3 . *If the kernel of a (finite or infinite) commutative semigroup $\mathcal{S}$ is not empty then it is a group (under the operation in $\mathcal{S}$).*

Proof: It is well known that a nonempty semigroup in which all linear equations $ax = b$ and $ya = b$ are solvable is a group. By commutativity we only need to consider one of these. We claim $ax = b$ is always solvable in the kernel $I \neq \emptyset$ $(a, b \in I)$. Solvability of $ax = b$ for all b amounts to saying that $aI = I$. But aI is a nonempty (right) ideal of $\mathcal{S}$ contained in I, so by commutativity and the minimality of I we have $aI = I$. □

We have completed the proofs of parts (a) and (c) of Proposition 3.1.

Proof of Proposition 3.1, part (b). By definition, the set A of fully accessible elements is precisely $\bigcap_{a \in \mathcal{S}} a\mathcal{S}$. All the sets $a\mathcal{S}$ are ideals, therefore A is an ideal and it contains the kernel. Now for all $a \in A$ we have $A \subseteq a\mathcal{S} \subseteq A$, therefore A is a minimal ideal, so A is the kernel. □

References

[1] R. Anderson, L. Lovász, P. Shor, J. Spencer, É. Tardos, S. Winograd: Disks, balls, and walls: the analysis of a combinatorial game. *Amer. Math. Monthly* **96/6** (1989) 481–493.

[2] L. Babai, I. Gorodezky: Sandpile transience on the grid is polynomially bounded. In: Proc. 18th Ann. Symp. on Discrete Algorithms (SODA'07), ACM–SIAM 2007, pp. 627-636.

[3] L. Babai, I. Gorodezky, A. Shapiro: It's a long way to recurrence: exponential transience class in the Abelian Sandpile Model. In preparation.

[4] L. Babai, E. Toumpakari: Abelian sandpiles with bounded transience class. In preparation.

[5] R. Bachem, P. de la Harpe, T. Nagnibeda: The lattice of integral flows and the lattice of integral cuts of a finite graph. *Bull. Soc. Math. France* **125/2** (1997) 167–198

[6] H. Bai: On the critical group of the n-cube. *Linear Algebra and its Applications* **369** (2003) 251–261

[7] P. Bak, C. Tang, K. Wiesenfeld: Self-organized Criticality. *Phys. Rev. A* **38** (1988) 364–374

[8] M. Baker, S. Norine: Riemann-Roch and Abel-Jacobi theory on a finite graph. *Adv. Math.* **215/2** (2007) 766–793

[9] N. Biggs: Algebraic Potential Theory on Graphs. *Bull. London Math. Soc.* **29** (1997) 641–682

[10] N. Biggs: Chip-Firing and the Critical Group of a Graph. *Journal of Algebraic Combinatorics* **9** (1999) 25–45

[11] G. Birkhoff, S. MacLane: *Algebra.* 3rd Edition, Chelsea Publ. Company, 1988.

[12] A. Bjørner, L. Lovász: Chip-firing Games on Directed Graphs. *Journal of Algebraic Combinatorics* **1** (1992) 305–328

[13] A. Bjørner, L. Lovász, P. Shor: Chip-firing Games on Graphs. *Europ. J. Combinatorics* **12** (1991) 283–291

[14] R. Cori, D. Rossin: On the Sandpile Group of Dual Graphs. *Europ. J. Combinatorics* **21** (2000) 447–459

[15] M. Creutz: Abelian Sandpile. *Computers in Physics* **5** (1991) 198–203

[16] D. Dhar: Self-organized Critical State of Sandpile Automaton Models. *Phys. Rev. Lett.* **64** (1990) 1613–1616

[17] D. Dhar, P. Ruelle, S. Sen, D. Verma: Algebraic Aspects of Abelian Sandpile Models. *J. Phys. A* **28/4** (1995) 805–831

[18] P.A. Grillet: *Semigroups.* Pure and Applied Mathematics **193**, M. Dekker, 1995

[19] P. A. Grillet: Nilsemigroups on trees. *Semigroup Forum* **43** (1991) 187–201

[20] B. Jacobson, A. Niedermaier, V. Reiner: Critical groups for complete multipartite graphs and Cartesian products of complete graphs. *J. Graph Theory* **44/3** (2003) 231–250

[21] H. Jensen: *Self-Organized Criticality.* Cambridge Lecture Notes in Physics Vol. 10, Cambridge University Press, 1998

[22] G. Kirchhoff: Über die Auflösung der Gleichungen, auf welche man bei der Untersuchung der linearen Verteilung galvanischer Ströme geführt wird. *Ann. Phys. Chem.* **72** (1847) 497-508

[23] Michael Kleber: Goldbug variations. *The Mathematical Intelligencer.* **27/1** (2005) 55–63

[24] L. Levine: The sandpile group of a tree. *European J. Combinatorics* **30/4** (2009) 1026–1035

[25] L. Lovász: *Combinatorial Problems and Exercises.* 2nd Edition (1993) Akadémiai Kiadó, Budapest, and Elsevier

[26] L. Lovász, P. Winkler: Mixing of random walks and other diffusions on a graph. In: *Surveys in Combinatorics*, 1995, P. Rowlinson, ed., London Math.Soc. Lecture Notes 218, 1995, pp. 119–154.

[27] D. J. Lorenzini: A finite group attached to the Laplacian of a graph. *Discrete Math.* **91/3** (1991) 277–282.

[28] M. H. A. Newmann: On theories with a combinatorial notion of "equivalence." *Ann. Math.* **43** (1942) 211–264

[29] E. Speer: Asymmetric Abelian Sandpile Models. *Journal of Statistical Physics* **71** nos.1/2 (1993) 61–74

[30] Z. Sunic, private communication, May 2007.

[31] G. Tardos: Polynomial bound for a chip firing game on graphs. *SIAM J. Disc. Math.* **1** (1988) 397–398

[32] E. Toumpakari: *On the Abelian Sandpile Model.* Ph. D. Thesis. Department of Mathematics, University of Chicago, June 2005. http://people.cs.uchicago.edu/~laci/students

[33] E. Toumpakari: On the sandpile group of regular trees. *Europ. J. Combinatorics* **28/2** (2007) 822–842

[34] W. T. Tutte: The dissection of equilateral triangles into equilateral triangles. *Proc. Cambridge Phil. Soc.* **44** (1948) 463–482