Sandpile transience on the grid is polynomially bounded

László Babai* Igor Gorodezky†

Abstract

We study the process of adding one grain at a time in the Abelian Sandpile Model (ASM) based on the $n \times n$ grid (Bak et al., 1988, Dhar, 1990, Dhar et al., 1995). We prove that within a polynomial (in n) number of steps, we necessarily reach a recurrent configuration regardless of the choice of sites where the grains are dropped. This adds credence to the common notion that "recurrence" represents "long-term behavior" of the system.

1 Introduction

The Abelian Sandpile Model (ASM) is a diffusion process on graphs devised in statistical physics ([BTW]) as a model of a phenomenon called *self-organized criticality.* It is related to a "chip firing game" studied in combinatorics and the theory of computing ([Ta, BLS]). The ASM, first studied in depth by Dhar et al. [Dh, MD, DRSV], provides a rich mathematical structure reaching across numerous disciplines. Probabilistic/statistical, game theoretic, combinatorial, algebraic, number theoretic, algorithmic, dynamical systems aspects of the ASM, its fractal behavior, its connection to certain cellular automata and to random walks and other diffusion processes have been studied by a diverse group of researchers. (Several of these connections are illuminated in the delightful article [Kl].)

In an ASM, grains of sand are added at "sites" (vertices of a graph); when the height of the pile at a site reaches the degree of the vertex, one grain is passed to a neighbor through each edge incident with the vertex ("toppling"). A special vertex, the *sink*, swallows all grains it receives and is never toppled. A configuration is *stable* if no vertex can be toppled.

The ASM based on the $n \times n$ grid with a sink added (the sink is connected to each vertex on the perimeter of the grid) was introduced by [BTW] and studied in detail in the pioneering work of Dhar et al. [DRSV]. For its plausible significance to statistical physics, this is the most extensively studied family of ASMs; it is also the focus of our work.

We study the process of adding one grain at a time and toppling as necessary. In the course of this process, a sequence of configurations is visited. Certain configurations are *recurrent* (they can be revisited), all others are *transient* (cannot be seen twice in such a sequence).

One of the striking facts of the general theory is that *the number of stable recurrent configurations is the number of spanning trees* of the underlying graph [MD]. We note that the number of transient configurations tends to be larger by an exponential factor.

The transient/recurrent distinction is fundamental to the evolution of the system; recurrent configurations are perceived as representing "long-term behavior." Our concern in this paper is, how long it takes for the system to reach a recurrent configuration. If it takes exponentially long, that brings the physical reality of the identification of "long-term behavior" with "recurrence" into question.

So we can think of the goal being to reach a recurrent configuration quickly (starting from the empty configuration $\mathbf{c} = 0$). After a brief comment on the average case, we shall turn to the worst case (grain sites selected adversarially).

The average case is easily settled; for an arbitrary graph, a "coupon collector" argument shows that if the grains are added at random sites, a recurrent configuration is reached in an expected polynomial number of steps (polynomial in the number of edges).

The worst case is defined by the notion of the *transience class* of the model, introduced in [BT]. The transience class is the maximum number of grains that can be added (by an adversary) before the configuration necessarily becomes recurrent.

While for some families of graphs the transience class is exponentially large [BGS], our main result establishes a *polynomial upper bound on the transience class for the grid-based ASMs* (Theorem 3.1).

To us, this outcome seemed surprising in light of the close association of the grids with the examples we found for exponential transience class. (These examples consist of a path, with each vertex adjacent to the sink. This graph is very much like the perimeter of the grid, together with the sink.) Motivated by that experience

*Department of Computer Science, University of Chicago. Email: laci@cs.uchicago.edu.

†Department of Mathematics, University of Chicago. Partially supported by an NSF VIGRE grant.
Email: igor@jfi.uchicago.edu.

we spent considerable effort trying to prove the opposite of the eventual result and developed techniques, based on Chebyshev polynomials, toward that failed goal. We believe that those techniques were not wasted and can be used to study the genesis of natural families of transient states. We plan to return to this project in a separate paper.

Our polynomial bound has a ridiculous exponent (around 30); we believe the true exponent to be closer to 3.

As a corollary to our main result, we prove that most transient configurations on the $n \times n$ grid cannot be reached by any sequence of moves (grain additions) that are limited to a set of $\le cn^2/\log n$ sites for certain constant $c > 0$. In contrast, all recurrent configurations can be reached by dropping grains at a set of n sites.

2 The Abelian Sandpile Model

2.1 Configurations over the ambient space

We will only consider undirected graphs. Parallel edges are permitted.

Definition 2.1. *A* graph G *is an ordered triple* $(V(G), E(G), f_G)$*, where* $V(G)$ *and* $E(G)$ *are finite sets and* f_G *is a function from* $E(G)$ *to the set of unordered pairs of not necessarily distinct elements of* $V(G)$*. The elements of* $V(G)$ *are called* vertices *and the elements of* $E(G)$ *are called* edges*.*

(This concept is often called a *multigraph*; we omit the prefix "multi.") We say that a vertex u is an *endvertex* of an edge e if $u \in f_G(e)$. Two vertices u, v are *adjacent*, written $u \sim v$, if there exists an edge e such that $f_G(e) = \{u, v\}$. As usual, $\mathbb{N} = \{0, 1, 2, \dots\}$.

The Abelian Sandpile Model takes a connected graph G with a special vertex $s \in V(G)$ designated as the *sink*. Following [BT], we refer to the pair $\mathcal{X} = (G, s)$ as the *ambient space* for our ASM. We write $V = V(\mathcal{X}) = V(G)$. Non-sink vertices are referred to as *ordinary*; the set of ordinary vertices is denoted by $V_0 = V_0(\mathcal{X}) = V \setminus \{s\}$.

Definition 2.2. *A configuration on* $\mathcal{X}$ *is a map* $\mathbf{c} : V_0 \longrightarrow \mathbb{N}$*. The* weight *of* $\mathbf{c}$ *is* $\sum_{v \in V_0} \mathbf{c}(v)$*.*

We think of the value $\mathbf{c}(v)$ as the number of "grains of sand" deposited at the "site" v; we will adopt this terminology for the remainder of the paper. We shall say that vertex v "holds $\mathbf{c}(v)$ grains" under configuration $\mathbf{c}$. Vertex v is *empty* if $\mathbf{c}(v) = 0$.

Definition 2.3. *The* support *of a configuration* $\mathbf{c}$ *is the set*

$$\operatorname{supp}(\mathbf{c}) = \{\, v \in V_0 : \mathbf{c}(v) > 0\}. \tag{2.1}$$

2.2 Stability, toppling

Definition 2.4. *A vertex* v *is* unstable *under* $\mathbf{c}$ *if* $\mathbf{c}(v) \ge \deg(v)$*; otherwise it is* stable*. A configuration* $\mathbf{c}$ *is* stable *if all vertices are stable under* $\mathbf{c}$*; otherwise it is unstable.*

If a vertex v is unstable it *topples* by passing one grain along each edge of which it is an endvertex. That is, when v is toppled, it loses $\deg(v)$ grains and each neighbor w gains $a_{v,w}$ grains where $a_{v,w}$ is the number of edges joining v and w. The sink never topples; all grains it receives are removed from the board.

The model thus has two basic operations: adding a grain at a site, and toppling an unstable site. A *toppling sequence* is a sequence of topplings while no grains are added. If a toppling sequence results in a stable configuration, it is called an *avalanche*. *Stabilization* is the process of performing an avalanche.

Notation 2.5. *Let* $\mathbf{c}_1$ *and* $\mathbf{c}_2$ *be configurations. We write* $\mathbf{c}_1 \ge \mathbf{c}_2$ *if* $(\forall v \in V_0)(\mathbf{c}_1(v) \ge \mathbf{c}_2(v))$ *and* $\mathbf{c}_1 \vdash \mathbf{c}_2$ *if a (possibly empty) toppling sequence leads from* $\mathbf{c}_1$ *to* $\mathbf{c}_2$*. We write* $\mathbf{c}_1 \succcurlyeq \mathbf{c}_2$ *if* $(\exists \mathbf{c}_3)(\mathbf{c}_1 \vdash \mathbf{c}_3 \ge \mathbf{c}_2)$*; we say in this case that* $\mathbf{c}_1$ dominates $\mathbf{c}_2$*.*

The following observation will be useful.

Proposition 2.1. *If* $\mathbf{c}_1 \succcurlyeq \mathbf{c}_2$ *and* $\mathbf{c}_1' \succcurlyeq \mathbf{c}_2'$ *then* $\mathbf{c}_1 + \mathbf{c}_1' \succcurlyeq \mathbf{c}_2 + \mathbf{c}_2'$*. In particular, if* $\mathbf{c}_1 \succcurlyeq \mathbf{c}_2$ *and* $k \ge 0$ *then* $k\mathbf{c}_1 \succcurlyeq k\mathbf{c}_2$*.*

It is easy to see that every toppling sequence is finite so we eventually reach a stable configuration. The following result, provable by a simple "diamond lemma" argument familiar from group theory (Jordan-Hölder theorem), is the key to the algebraic structure of the ASMs.

Theorem 2.2. ([BLS, Dh]) *Given a configuration* $\mathbf{c}$*, there exists a unique stable configuration* $\sigma(\mathbf{c})$ *such that* $\mathbf{c} \vdash \sigma(\mathbf{c})$*, i. e., every avalanche leads to* $\sigma(\mathbf{c})$*.*

Definition 2.6. *Suppose* $\mathbf{c}_1 \vdash \mathbf{c}_2$*. The* score vector $\mathbf{z}^{\mathbf{c}_1,\mathbf{c}_2} : V_0 \to \mathbb{N}$ *records how many times each* $v \in V_0$ *was toppled during a toppling sequence from* $\mathbf{c}_1$ *to* $\mathbf{c}_2$*. Notation:* $\mathbf{z}^{\mathbf{c}} := \mathbf{z}^{\mathbf{c},\sigma(\mathbf{c})}$*.*

Proposition 2.3. *Assuming* $\mathbf{c}_1 \vdash \mathbf{c}_2$*, the score vector* $\mathbf{z}^{\mathbf{c}_1,\mathbf{c}_2}$ *is well defined (depends on* $\mathbf{c}_1$ *and* $\mathbf{c}_2$ *only, not on the particular toppling sequence chosen).*

2.3 Recurrence

Definition 2.7. *The configuration* $\mathbf{c}$ *is* recurrent *if given any configuration* $\mathbf{d}$*, there exists a configuration* $\mathbf{d}' \ge \mathbf{d}$ *such that* $\mathbf{d}' \vdash \sigma(\mathbf{c})$*.*

A configuration is thus recurrent if its stabilization is "reachable" from any configuration by adding grains and toppling. Configurations that are not recurrent are called *transient*.

Let $\mathbf{c}_{\max}$ denote the stable configuration defined by $\mathbf{c}_{\max}(v) = \deg(v) - 1$ for all $v \in V_0$.

PROPOSITION 2.4. (a) *If* $\mathbf{c}_1 \succcurlyeq \mathbf{c}_2$ *and* $\mathbf{c}_2$ *is recurrent then so is* $\mathbf{c}_1$. (b) $\mathbf{c}_{\max}$ *is recurrent.*

DEFINITION 2.8. *Define the* transience class $\mathrm{tcl}(\mathcal{X})$ *as the weight of the heaviest transient configuration on* $\mathcal{X}$. *If no such maximum exists, set* $\mathrm{tcl}(\mathcal{X}) = \infty$.

It is easy to see (cf. [BT]) that $\mathrm{tcl}(\mathcal{X})$ is finite if and only if the subgraph of $\mathcal{X}$ induced on V_0 is connected.

REMARK 2.1. Let $\mathcal{M} = \mathcal{M}(\mathcal{X})$ be the set of stable configurations. Define the binary operation $\oplus$ on $\mathcal{M}$ by pointwise addition and stabilization: $\mathbf{c} \oplus \mathbf{d} = \sigma(\mathbf{c} + \mathbf{d})$. With this operation, $\mathcal{M}$ is a commutative monoid, called the *sandpile monoid*.

It is known [Dh, DRSV] that the stable recurrent configurations form a subgroup of the sandpile monoid. Let $\mathcal{G}$ denote this abelian group, called the *sandpile group*. It is also known that $\mathcal{G}$ is the unique minimal ideal of $\mathcal{M}$ [BT].

Let $\mathcal{S}$ be the subsemigroup of $\mathcal{M}$ generated by the nonzero stable configurations, called the *sandpile semigroup* (so $\mathcal{M} = \mathcal{S} \cup \{0\}$). Assume that the induced subgraph of $\mathcal{X}$ on V_0 is connected and has at least one edge. Then $\mathcal{G}$ is a subset and therefore an ideal of $\mathcal{S}$ and we can consider the *Rees quotient* $\mathcal{S}/\mathcal{G}$ obtained by contracting $\mathcal{G}$ to a single element (which becomes a zero of the quotient). Now $\mathcal{S}/\mathcal{G}$ is a nilpotent semigroup, called the *sandpile quotient* [BT]. It is easy to see that the nilpotence class of $\mathcal{S}/\mathcal{G}$ is $\mathrm{tcl}(\mathcal{X}) + 1$.

2.4 The augmented grid as an ambient space

The ambient space associated with the $n \times n$ grid is of particular importance to statistical physics (cf. [Dh, DRSV]).

DEFINITION 2.9. *Given the* $n \times n$ *grid graph (with* n^2 *vertices), associate with it an ambient space* $\mathcal{X}_n$ *as follows. Add a sink to the grid; attach each non-corner boundary vertex to the sink with a single edge, and attach each corner of the grid to the sink with two edges. (So there will be a total of* $n^2 + 1$ *vertices; all vertices other than the sink will have degree 4; and the sink will have degree* $4n$*.)*

Note that by the above, $\mathrm{tcl}(\mathcal{X}_n)$ is finite.

We will often refer to $\mathcal{X}_n$ as "the $n \times n$ grid;" in the context of ambient spaces this will cause no confusion.

We will also consider the infinite grid, which will serve as a model of local behavior on large finite grids. We regard the infinite grid as an ambient space with no sink; a configuration will by definition have *finite weight*. Stabilization continues to be well-defined under this assumption; Theorem 2.2 remains valid.

3 Main results

Our main result is a polynomial upper bound on the transience class of $\mathcal{X}_n$, the ambient space associated with the $n \times n$ grid.

THEOREM 3.1. $\mathrm{tcl}(\mathcal{X}_n) = O(n^{30})$.

A corollary on the genesis of almost all transient configurations follows.

We define the *rank* of a stable configuration $\mathbf{c}$ to be the smallest number of sites required to "generate" $\mathbf{c}$.

DEFINITION 3.1. *For a stable configuration* $\mathbf{c}$ *we set*

$$\mathrm{rank}(\mathbf{c}) = \min\{|\operatorname{supp}(\mathbf{d})| : \sigma(\mathbf{d}) = \mathbf{c}\} \tag{3.2}$$

Here $\mathbf{d}$ ranges over all (not necessarily stable) configurations. We view the "rank" as a cost measure for reaching a given configuration. We shall see that all recurrent configurations over the grid are "inexpensive" while most transient configurations are "expensive."

PROPOSITION 3.2. *Over the* $n \times n$ *grid, every recurrent configuration has rank* $\le n$.

COROLLARY 3.3. *For all but a* $(1 + c)^{-n^2}$ *fraction of the stable transient configurations* $\mathbf{c}$*,*

$$\mathrm{rank}(\mathbf{c}) = \Omega(n^2/\log n). \tag{3.3}$$

We don't know whether the log factor could be removed in this result (see Conjecture 8.3). Some sparse transient configurations have rank $\Omega(n^2)$. (Put a grain on every third vertex in every third row.)

4 Preliminaries

4.1 Single-step parallel avalanche

DEFINITION 4.1. *Let* $U(\mathbf{c})$ *denote the set of vertices that are unstable under* $\mathbf{c}$. *The* single-step, parallel toppling *of* $\mathbf{c}$ *is the toppling of each* $v \in U(\mathbf{c})$ *exactly once. Let* $\mathrm{st}(\mathbf{c})$ *denote the result of this operation.*

Next we define the "single-step parallel toppling" process; we refer to the parameter t as "time."

DEFINITION 4.2. *Let* $\mathrm{st}^0(\mathbf{c}) = \mathbf{c}$ *and for* $t \ge 0$*, let* $\mathrm{st}^{t+1}(\mathbf{c}) = \mathrm{st}(\mathrm{st}^t(\mathbf{c}))$. *Let* $T_0(\mathbf{c}) = \operatorname{supp}(\mathbf{c})$ *and* $T_t(\mathbf{c}) = \bigcup_{i=0}^{t} \operatorname{supp}(\mathrm{st}^i(\mathbf{c}))$. *We call* $T_t(\mathbf{c})$ *the set of vertices* touched by time t. *Let* $\mathbf{z}_t^{\mathbf{c}} = \mathbf{z}^{\mathbf{c},\mathrm{st}^t(\mathbf{c})}$.

In other words, $v \in V_0$ belongs to $T_t(\mathbf{c})$ if v received a grain by time t (including the possiblity $\mathbf{c}(v) > 0$). Note that for $t \geq 1$ we have $v \in T_t(\mathbf{c})$ if and only if $v \in \operatorname{supp}(\mathbf{c})$ or v has a neighbor in $\operatorname{supp}(\mathbf{z}^{\mathbf{c}}_{t-1})$.

DEFINITION 4.3. *Let τ be first time that $\mathrm{st}^\tau(\mathbf{c}) = \sigma(\mathbf{c})$. At this point the process stabilizes, and we use the subscript or superscript ∞ to describe our variables at any time $t \geq \tau$. So $\mathrm{st}^\infty(\mathbf{c}) = \sigma(\mathbf{c})$. We set $T(\mathbf{c}) = T_\infty(\mathbf{c})$. Note that $\mathbf{z}^{\mathbf{c}}_\infty = \mathbf{z}^{\mathbf{c}}$.*

We will omit the superscript from $\mathbf{z}^{\mathbf{c}}_t$ when $\mathbf{c}$ is clear from the context.

It is clear from the definitions that $\mathbf{z}_0(v) = 0$ and $\mathbf{z}_1(v) = \min\{\lfloor \mathbf{c}(v)/\deg(v)\rfloor, 1\}$. In general, for $t \geq 0$, one can prove by induction that

(4.4)
$$\mathbf{z}_{t+1}(v) = \min\left\{ \left\lfloor \frac{\mathbf{c}(v) + \sum_{u\in V_0} a_{u,v}\mathbf{z}_t(u)}{\deg(v)} \right\rfloor, \mathbf{z}_t(v)+1 \right\}.$$

4.2 Flooding

We shall use the following notation.

DEFINITION 4.4. *Fix an ambient space $\mathcal{X}$. For $S \subseteq V_0$, the* characteristic configuration *of S is defined as*

$$\chi_S(v) = \begin{cases} 1 & \text{if } v \in S \\ 0 & \text{otherwise} \end{cases} \qquad (4.5)$$

If $S = \{u\}$ for some $u \in V_0$, we shall write χ_u for χ_S.

PROPOSITION 4.1. *Let u and v be adjacent ordinary vertices of the ambient space $\mathcal{X}$. Let $\mathbf{c}$ be a configuration with $u \in \operatorname{supp}(\mathbf{c})$. Suppose $\mathbf{c} \vdash \mathbf{d}$. Then*

$$\operatorname{supp}(\mathbf{d}) \cap \{u, v\} \neq \emptyset. \qquad (4.6)$$

In other words, if u has been touched at some point in the course of an avalanche then at no later time in the avalanche can both u and v be empty.

DEFINITION 4.5. *Fix an ambient space $\mathcal{X}$. Let $S \subseteq V_0$. We say that configuration $\mathbf{c}$* floods *S if $S \subseteq T(\mathbf{c})$.*

We call an ordinary vertex v of the ambient space $\mathcal{X}$ isolated if v is not adjacent to any other ordinary vertex.

COROLLARY 4.2. *Assume $\mathcal{X}$ has no isolated vertices. Let $\mathbf{c}$ be a configuration on the ambient space $\mathcal{X}$ and let m be the maximum ordinary degree in $\mathcal{X}$. If $\mathbf{c}$ floods S then $(m+1)\mathbf{c} \succcurlyeq \chi_S$.*

Proof. We have $(m+1)\mathbf{c} \succcurlyeq (m+1)\chi_{\operatorname{supp}(\sigma(\mathbf{c}))}$. By Proposition 4.1, a single parallel toppling now puts a grain on every empty vertex in S without depleting any vertex. ■

COROLLARY 4.3. *Assume $\mathcal{X}$ has no isolated vertices. Let $\mathbf{c}$ be a configuration on an ambient space $\mathcal{X}$ and let m be the maximum ordinary degree in $\mathcal{X}$. If $T(\mathbf{c}) = V_0$ then the configuration $(m^2-1)\mathbf{c}$ is recurrent.*

Proof. By Corollary 4.2, $(m+1)\mathbf{c} \succcurlyeq \chi_{V_0}$ and therefore $(m-1)(m+1)\mathbf{c} \succcurlyeq (m-1)\chi_{V_0} \geq \mathbf{c}_{\max}$ and we are done by Proposition 2.4. ■

5 Avalanches on the infinite grid

5.1 Convexity: axis-monotone configurations

Essential to the proof of the main result is the invariance of a certain convexity property under stabilization.

The points of a grid will be called "lattice-points." We shall use the term *square* to mean the set of lattice-points inside an axis-aligned square. The term *diamond* will refer to the set of lattice points in a square of which the diagonals are axis-aligned. Note that a square is an ℓ_∞-ball on the grid, and a diamond is an ℓ_1-ball.

D_4 denotes the dihedral group of order 8 represented as the group of symmetries of a square of positive radius. Note that a diamond has the same group of symmetries.

DEFINITION 5.1. *Consider a function $f : \mathbb{Z}^2 \to \mathbb{N}$. We say that f has* D_4 symmetry *if there exists a lattice point $v = (i, j)$ such that f is symmetric about the lines $x = i, y = j, x - i = y - j$, and $x - i = -y + j$. We say that f is* centered at *v.*

We also apply this concept to functions f defined on a finite grid by extending the definition of f to the infinite grid by assigning the value 0 to all lattice points outside our finite grid.

All our functions will have nonempty finite supports, so the *center* defined above will be unique.

D_4 symmetry involves 4 axes of symmetry which divide the plane into eight *octants*, each bounded by two half-axes of symmetry at a 45-degree angle which meet at the center of the D_4 symmetry.

DEFINITION 5.2. *Consider a function $f : \mathbb{Z}^2 \to \mathbb{N}$ with nonempty finite support. Assume f has D_4 symmetry. We say that f is* axis monotone *if $f(r) \leq f(s)$ for each pair of lattice points $r, s \in \mathbb{Z}^2$ such that $s - r$ is perpendicular to some axis A of symmetry of f and s is closer to A than r.*

In other words, we require that the function be nondecreasing as we move perpendicularly toward any axis of the D_4 symmetry of f (see Figure 1).

DEFINITION 5.3. *We define the "diamond" and the "square" of radius r, centered at $(i, j) \in \mathbb{Z}^2$, as*

$$\diamond_r(i,j) = \{\, (a,b) \in \mathbb{Z}^2 : |i-a| + |j-b| \leq r \} \qquad (5.7)$$

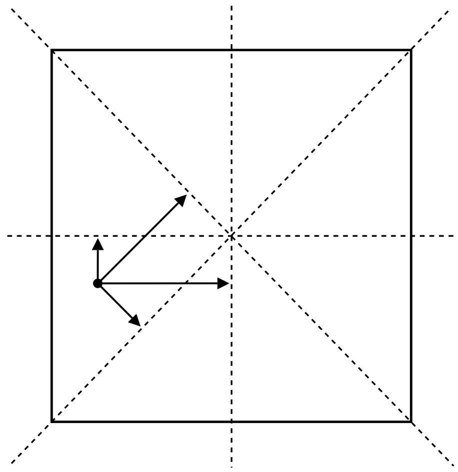

Figure 1: An illustration of axis monotonicity; the arrows indicate the directions in which the values of an axis monotone configuration would increase. Note that if a configuration increases in the two outermost directions (here, north and southeast), then it follows that it increases in all four directions.

and

(5.8) $\square_r(i,j) = \{ (a,b) \in \mathbb{Z}^2 : |i-a| \le r, \; |j-b| \le r \}$,

respectively.

We omit (i,j) from the square and diamond notation if the center is irrelevant or can be inferred from the context. We observe that

$$|\diamond_r| = 2r^2 + 2r + 1 \quad \text{and} \quad |\square_r| = (2r+1)^2. \tag{5.9}$$

PROPOSITION 5.1. *Assume the function $f : \mathbb{Z}^2 \to \mathbb{N}$ is axis-monotone, has D_4 symmetry, and has nonempty finite support. Then there exists $s \in \mathbb{N}$ such that*

$$\diamond_s \subseteq \mathrm{supp}(f) \subseteq \square_s. \tag{5.10}$$

We call the value s the radius *of* $\mathrm{supp}(f)$.

5.2 Preserving axis-monotonicity If $\mathcal{X}$ is the infinite grid, we can bound the number of vertices touched in terms of the weight of the configuration.

PROPOSITION 5.2. *Let $\mathbf{c}$ be a configuration of weight k on the infinite grid. Then*

$$k/3 \le |T(\mathbf{c})| \le 2k. \tag{5.11}$$

Proof. On the infinite grid no grains are lost, so the weight of $\sigma(\mathbf{c})$ is k. We have $|T(\mathbf{c})| \le 2k$ by Proposition 4.1. On the other hand, $k \le 3|T(\mathbf{c})|$; this is because $\sigma(\mathbf{c})$ is by definition stable, so no vertex may hold more than 3 grains. ■

Assumption. Henceforth throughout this section we assume that $\mathbf{c}$ *is a nonempty configuration on the infinite grid with D_4 symmetry and* $\mathbf{c}$ *is axis-monotone.*

The following lemma is inspired by Theorem 4 from Le Borgne and Rossin [LR]. While our methods are similar to those of [LR], we note that our choice of definitions strengthened both the method and the claim proven, even in the special case considered by [LR] (diffusion of a single stack of grains). In particular, the introduction of single-step parallel topplings made it possible to prove a stronger monotonicity for $\mathbf{z}_t$ than what was proven for the analogous quantity $c(t)$ in [LR] by avoiding an annoying parity problem, thus allowing for tighter and cleaner proofs in which fewer cases need to be considered.

LEMMA 5.3. $\mathbf{z}_t^{\mathbf{c}}$ *is axis-monotone for all t.*

Proof. We use the following notation: if v is a lattice point then $v = (v_1, v_2)$ and $v' = (v_1, v_2+1)$. Assume the center of $\mathbf{c}$ is $a = (a_1, a_2)$. Suppose we are approaching a horizontal axis from below (as in Figure 1). We need to show that for all v and t, if $v_2 < a_2$ then $\mathbf{z}_t(v) \le \mathbf{z}_t(v')$.

We proceed by induction on t. By the axis-monotonicity of $\mathbf{c}$, we have $\mathbf{c}(v) \le \mathbf{c}(v')$, settling the case $t = 0$. Assuming now the statement for t, let us prove $\mathbf{z}_{t+1}(v) \le \mathbf{z}_{t+1}(v')$. By the nature of single-step parallel toppling we know that $\mathbf{z}_{t+1}(v) \le 1 + \mathbf{z}_t(v)$, so we are done if $\mathbf{z}_t(v) < \mathbf{z}_t(v')$. We may therefore assume $\mathbf{z}_t(v) = \mathbf{z}_t(v')$.

For the inductive step, we use equation (4.4). Applying the $v \mapsto v'$ operator to all variables in this equation (so $u \mapsto u'$) we see that by the inductive hypothesis, each term of the expression for $\mathbf{z}_{t+1}(v')$ dominates the corresponding term for $\mathbf{z}_{t+1}(v)$, except possibly if $v_2 + 1 = a_2$ (v' is on the horizontal axis of symmetry). In this case, the required inequality between the "northern" neighbors of v and v' is not immediately evident; what is required is $\mathbf{z}_t(v') \le \mathbf{z}_t(v'')$. But v'' and v are now positioned symmetrically about the horizontal axis, so $\mathbf{z}_t(v'') = \mathbf{z}_t(v)$ and by our assumption this is equal to $\mathbf{z}_t(v')$, completing the case when we approach a horizontal axis.

The other case to consider is when $w - v$ is perpendicular to a diagonal axis. In this case the inductive step again works trivially unless v is on the axis of symmetry; but if it is then the symmetry yields the required termwise inequality. ■

It follows by Proposition 5.1 that for each t there exists a radius s such that

$$\diamond_s \subseteq \mathrm{supp}(\mathbf{z}_t^{\mathbf{c}}) \subseteq \square_s. \tag{5.12}$$

COROLLARY 5.4. *For each t there exists a radius r_t such that* $\Diamond_{r_t} \subseteq T_t(\mathbf{c}) \subseteq \Box_{r_t}$.

Proof. It suffices to recall that a vertex is touched by time t if it is in $\operatorname{supp}(\mathbf{c})$ or it has a neighbor in $\operatorname{supp}(\mathbf{z}_t^{\mathbf{c}})$. Let s satisfy equation (5.12) and let $\Diamond_q \subseteq \operatorname{supp}(\mathbf{c}) \subseteq \Box_q$ (Proposition 5.1). Set $r_t = \max\{s+1, q\}$. ■

It is clear that the function r_t is monotone and continuous:

$$r_t \le r_{t+1} \le r_t + 1. \tag{5.13}$$

COROLLARY 5.5. *For any R in the interval $r_0 \le R \le r_\infty$, there exists t such that $r_t = R$.*

Next, we estimate r_∞.

PROPOSITION 5.6. *If $\mathbf{c}$ has weight k then*

$$r_\infty \ge \sqrt{k/12} + (1/2). \tag{5.14}$$

Proof. Assume $R < \sqrt{k/12} + (1/2)$. Then $|T(\mathbf{c})| \ge k/3 > (2R-1)^2 = |\Box_{R-1}|$, therefore $r_\infty \ge R$. ■

Now we come to the main result of this section. The result estimates how tall a stack of grains we need to put on every vertex of a diamond in order to touch a larger diamond without spilling beyond the corresponding larger square.

THEOREM 5.7. *Assume the integers R, r, h satisfy $R > r \ge 1$ and let $h \ge 6(R/r)^2$. Then there exists t such that*

$$\Diamond_R \subseteq T_t(h\chi_{\Diamond_r}) \subseteq \Box_R. \tag{5.15}$$

Proof. We show that the weaker inequality

$$h > (3(2R-1)^2)/(2r^2 + 2r + 1) \tag{5.16}$$

suffices in lieu of $h \ge 6(R/r)^2$.

Let $\mathbf{c} = h\chi_{\Diamond_r}$ and $k = h(2r^2 + 2r + 1)$. Then k is the weight of $\mathbf{c}$. By Proposition 5.2, $|T(\mathbf{c})| \ge k/3 > (2R-1)^2 = |\Box_{R-1}|$. So $r_\infty \ge R$ and the result follows from Corollary 5.5. ■

6 The $n \times n$ grid: an upper bound on the transience class

For the $n \times n$ grid we use $V_0 = \{1, \ldots, n\} \times \{1, \ldots, n\}$.

6.1 Reduction to a single stack Note that in the grid, the characteristic configuration χ_S has D_4 symmetry and is axis-monotone if and only if the set S has D_4 symmetry and is axis-convex (convex in the directions of the axes of symmetry).

The following general observation allows us to reduce the question of a polynomial bound on the transience class to the case of a single stack of grains.

LEMMA 6.1. *For $v \in V_0$, let $N(v)$ denote the smallest value such that $N(v)\chi_v$ floods the entire grid. Let $N = \sum_{v \in V_0} N(v)$. Then every configuration of weight $\ge 15N$ is recurrent.*

Proof. Let the configuration $\mathbf{c}$ have weight $\ge 15N$. Then for some $v \in V_0$ we have $\mathbf{c}(v) \ge 15N(v)$. Now Corollary 4.3 implies that $15N\chi_v$ is recurrent; therefore so is $\mathbf{c} \ge 15N\chi_v$. ■

So it suffices to prove that there exists a polynomial P such that for any $v \in V_0$, the configuration $P(n)\chi_v$ floods the entire grid.

6.2 The traveling diamond lemma The crux of the proof is the observation that if a stack of a certain constant height is put on every node of a diamond, then it floods a larger diamond that is closer to the center; repetition of this process will grow the diamonds and move them toward the center in a geometric progression.

This idea is formalized in the following lemma.

DEFINITION 6.1. *We say that a subset of the $n \times n$ grid is* constrained *if it contains at least one point of the boundary of the grid.*

It is clear that $\Diamond_r(i,j)$ and $\Box_r(i,j)$ are constrained at the same time, namely, if and only if

$$r = \min\{i-1, j-1, n-i, n-j\}. \tag{6.17}$$

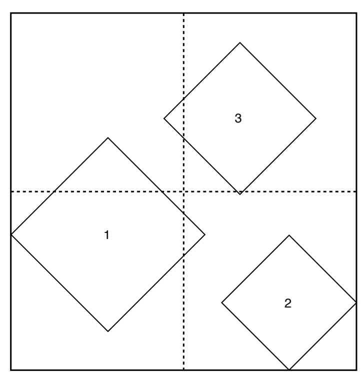

Figure 2: Diamonds on a finite grid. Diamonds 1 and 2 are constrained, while diamond 3 is not.

LEMMA 6.2. *Suppose $1 \le i \le j \le \lceil n/2 \rceil$. Let $\Diamond_r(i,j)$ be a constrained diamond and let $\mathbf{c}$ be a configuration that floods $\Diamond_r(i,j)$. For integers $p, q \ge 0$ such that $\max\{i+p, b+q\} \le \lceil n/2 \rceil$, define*

$$d = \min\{i+p-1, j+q-1\}.$$

Then the diamond $\diamond_d(i+p, j+q)$ is constrained. Moreover if $\max\{p,q\} \le \lfloor r/4 \rfloor$ then this diamond is flooded by the configuration $50\mathbf{c}$; and if $p \le \lfloor r/2 \rfloor$ and $q = 0$ then it is flooded by $70\mathbf{c}$.

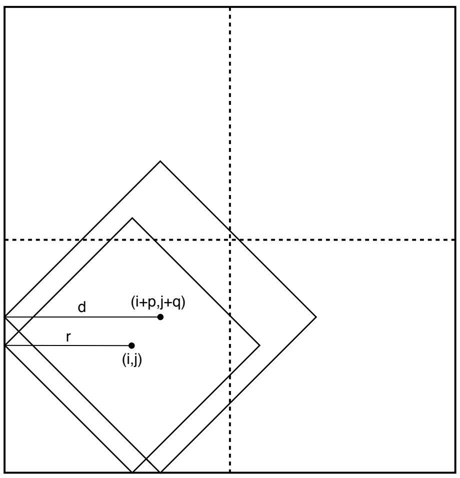

Figure 3: Lemma 6.2 states that if a configuration $\mathbf{c}$ floods the smaller constrained diamond of radius r and $p, q \le r/4$ then $50\mathbf{c}$ will flood the larger constrained diamond of radius d.

Proof. First consider the case $p, q \le \lfloor r/4 \rfloor$. Let $s = \lfloor r/2 \rfloor$. Clearly, $\diamond_s(i+p, j+q) \subset \diamond_r(i,j)$. It is easy to see that $d \le 5r/4 \le 5s/2$. Since $\mathbf{c}$ floods $\diamond_r(i,j)$, by Corollary 4.2 we have $5\mathbf{c} \succcurlyeq \chi_{\diamond_r(i,j)} \ge \chi_{\diamond_s(i+p,j+q)}$ and therefore $50\mathbf{c} \succcurlyeq 10\chi_{\diamond_s(i+p,j+q)}$. We claim that $10\chi_{\diamond_s(i+p,j+q)}$ floods $\diamond_d(i+p, j+q)$. This follows from Theorem 5.7. Indeed, when toppling $10\chi_{\diamond_s(i+p,j+q)}$, we can pretend that we are on the infinite grid as long as we don't spill beyond the boundaries of our $n \times n$ square.

Setting $R = d$ in Theorem 5.7, we see that $6(R/r)^2 \le 6(5/4)^2 < 10$. Therefore, for some t, after t steps of single-step parallel toppling of $10\chi_{\diamond_s(i+p,j+q)}$, the larger diamond $\diamond_R(i+p, j+q)$ is flooded while all the grains still stay within the square $\square_R(i+p, j+q)$ and therefore within the $n \times n$ grid.

The other case works analogously, noting that the expansion factor is $d/r \le 3/2$ and $6(3/2)^2 < 14$. ■

6.3 Upper bound strategy We exhibit a polynomial P such that given any $v, v' \in V_0$, the configuration $P(n)\chi_v$ touches v' (or equivalently, for any $v \in V_0$ the configuration $P(n)\chi_v$ floods the entire $n \times n$ grid). We do this by way of the following process (illustrated in the figures in the Appendix).

For simplicity, we assume n is odd.

First, by putting enough grains on v we flood a constrained diamond of radius ≥ 4. Call its center v_1.

Given $v_1, v' \in V_0$ we chart a path from v_1 (the "departure vertex," not to be confused with our "start vertex" v)) to v' (the "destination vertex") that runs through the center of the grid, consisting of segments that run either diagonally or parallel to a coordinate axis; this is illustrated in Figure 4.

In our *expansion phase* we move from v_1 to the center along the shortest path consisting of a diagonal segment followed by a segment along a median.

Having flooded our initial constrained diamond $\diamond$ of radius ≥ 4 centered at v_1, we add grains at the vertices of $\diamond$ so that toppling the vertices of $\diamond$ floods a larger constrained diamond whose center is now slightly farther along our path toward the center. We repeat this, adding grains and creating progressively larger constrained diamonds whose centers move toward the center of the grid, until finally we have flooded the (unique) constrained diamond centered at the center of grid. Each diamond is larger by a constant factor ($\approx 5/4$) than the previous one. This completes the expansion phase; this is illustrated in Figures 5 and 6.

Our *contraction phase* starts at the center and moves to v' along the shortest path consisting of a diagonal segment followed by a segment parallel to an axis.

During this phase we run the above process "backwards" along our path from the center to v', adding grains to create a sequence of "shrinking" constrained diamonds whose centers move along the diagonal segment of this path toward v', until finally one of the diamonds contains v'. (We do not need to make any moves along the axis-parallel segment of this second path.) This completes the contraction phase; it is illustrated in Figure 6.

Now instead of putting grains on our various diamonds, we are only permitted to put more grains on v. Each diamond to be covered by stacks of a constant height will require an increase of the initial stack at v by a constant factor. Since we traverse our path in a logarithmic number of steps (the diamonds grow, and then shrink, in geometric progressions), a polynomial bound on the height of the initial stack follows.

We give the details below.

6.4 Expansion phase Starting at vertex $v \in V_0$, we first need to create our *departure vertex* v_1, the center of a constrained diamond of radius $r_1 \ge 4$ we can flood at a modest cost. For $v \in V_0$, let $d(v)$ denote the distance of v from the boundary, i.e., if $v = (i,j)$ then $d(v) = \min\{i-1, j-1, n-i, n-j\}$.

LEMMA 6.3. *There exists a constant C such that for* *every $v \in V_0$ the configuration $Cd(v)^2\chi_v$ floods a constrained diamond of radius $r_1 = \max\{4, d(v)\}$.*

Proof. It follows from general principles that for some absolute constant C_1, the stack of height C_1 at v floods a neighborhood of radius 12 of v. If $d(v) < 4$ then this neighborhood contains a constrained diamond of radius 4. If $d(v) \geq 4$ then a stack of height $6d(v)^2$ will flood the diamond of radius $d(v)$ centered at v without spilling beyond the boundary of the $n \times n$ grid Theorem 5.7 (applied with $r = 1$). ∎

We call the center of this diamond $v_1 = (i_1, j_1)$, our "departure point." Assume without loss of generality that v_1 is in the lower left quadrant, i. e., $i_1, j_1 \leq \lceil n/2 \rceil$. Let $r_1 = \min\{i_1, j_1\} - 1$; this is the radius of our first diamond. Inductively, let

$$i_{k+1} = \min\{i_k + s_k, \lceil n/2 \rceil\}, \tag{6.18}$$

$$j_{k+1} = \min\{j_k + s_k, \lceil n/2 \rceil\}, \tag{6.19}$$

$$r_{k+1} = \min\{i_{r+1}, j_{r+1}\} - 1; \tag{6.20}$$

where $s_k = \lfloor r_k/4 \rfloor$ if $\max\{i_k, j_k\} < \lceil n/2 \rceil$, and $s_k = \lfloor r_k/2 \rfloor$ otherwise.

Let K be the smallest k such that v_k is the center ($i_K = j_K = \lceil n/2 \rceil$). Our sequence ends here.

By this definition, the vertices (i_k, j_k) move diagonally until one of the coordinates increases to $\lceil n/2 \rceil$; at that point they continue moving along the median toward the center of the grid. Note that i_k and j_k are never greater than $\lceil n/2 \rceil$.

Clearly, for each k, $\diamond_{r_k}(i_k, j_k)$ is a constrained diamond. Moreover, by Lemma 6.2, if the configuration $\mathbf{c}$ floods $\diamond_{r_1}(i_1, j_1)$, then $50^{k-1}\mathbf{c}$ floods $\diamond_{r_k}(i_k, j_k)$.

Consider the first, diagonal segment of movement. It is true by definition that $r_k = \min\{i_k, j_k\} - 1$, and since both i_k and j_k increase by $\lfloor r_k/4 \rfloor$ at each step of this stage, we have

$$r_k = r_{k-1} + \left\lfloor \frac{r_{k-1}}{4} \right\rfloor \tag{6.21}$$

and therefore

$$\frac{5}{4} r_{k-1} \geq r_k \geq \frac{5}{4} r_{k-1} - \frac{3}{4}. \tag{6.22}$$

If we take ℓ_1 steps along the diagonal segment, we reach a radius $R_1 = r_{\ell_1+1}$ satisfying

$$(5/4)^{\ell_1} r_1 \geq R_1 \geq (5/4)^{\ell_1}(r_1 - 3) + 3. \tag{6.23}$$

The movement is even faster along the second segment (a median): At this stage, one of the coordinates is always $\lceil n/2 \rceil$, and the other increases by $\lfloor r_k/2 \rfloor$ at each step. Since again we have $r_k = \min\{i_k, j_k\} - 1$, it follows that

$$r_k = r_{k-1} + \left\lfloor \frac{r_{k-1}}{2} \right\rfloor \tag{6.24}$$

and therefore

$$\frac{3}{2} r_{k-1} \geq r_k \geq \frac{3}{2} r_{k-1} - \frac{1}{2}. \tag{6.25}$$

If we make ℓ_2 steps along the median, the radius increases to $R_2 = r_{\ell_1+\ell_2+1}$, where

$$(3/2)^{\ell_2} R_1 \geq R_2 \geq (3/2)^{\ell_2}(R_1 - 1) + 1. \tag{6.26}$$

Noting that $r_{\ell_1+\ell_2} < n/2$ we obtain the inequality

$$(3/2)^{\ell_2 - 1}(5/4)^{\ell_1}(r_1 - 3) < n/2. \tag{6.27}$$

On the other hand, each diagonal step costs us a factor of 50 while each step on the median costs a factor of 70 in the height of the initial stack according to Lemma 6.2.

Combining these two inequalities we find that, if the configuration $\mathbf{c}$ floods the constrained diamond $\diamond_{r_1}$ then the configuration $(n/(r_1 - 3))^\alpha \mathbf{c}$ will flood the central diamond (of radius $\lceil n/2 \rceil$), where

$$\alpha = \log 50 / \log(5/4) \approx 17.5314. \tag{6.28}$$

(The highest cost estimate is obtained when we start from the corner and have to move diagonally all the way to the center).

Combined with Lemma 6.3, we obtain the following.

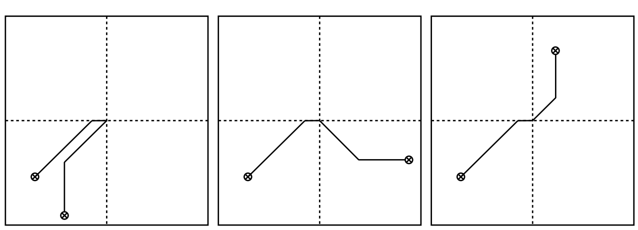

Figure 4: Three examples of a departure/destination vertex pair (the departure vertex is the same in all three). In the proof of the polynomial bound, given such a pair we chart a path from the departure vertex to the center of the grid consisting of a diagonal segment followed by a segment along a median. Subsequently we move from the center to the destination, first diagonally and then parallel to an axis.

LEMMA 6.4. *For any $v \in V_0$, the configuration $f(r_1)\chi_v$ will flood $\diamond_{\lfloor n/2 \rfloor}$, where $r_1 = \max\{d(v), 4\} \geq 4$, and*

$$f(r_1) = Cr_1^2 \left(\frac{n}{r_1 - 3} \right)^\alpha. \tag{6.29}$$

This concludes the analysis of the expansion phase.

The contraction phase is analogous but somewhat simpler; we omit the detailed analysis and instead summarize it in the following lemma.

LEMMA 6.5. *Assume that the configuration* $\mathbf{c}$ *floods the central diamond* $\diamond_{\lfloor n/2 \rfloor}$. *Then the configuration* $C'n^\beta \mathbf{c}$ *floods the entire* $n \times n$ *grid, where* C' *is an absolute constant and*

$$\beta = \log 70 / \log(3/2) \approx 10.4781. \tag{6.30}$$

6.5 Putting it all together We are almost done.

THEOREM 6.6. *Let* $\mathcal{X}_n$ *denote the ambient space associated with the* $n \times n$ *grid. Then* $\operatorname{tcl}(\mathcal{X}_n) = O(n^\gamma)$ *where*

$$\gamma = \alpha + \beta + 1 \approx 29.0095. \tag{6.31}$$

Proof. Fix an ordinary vertex v. Lemmas 6.4 and 6.5 together imply that the configuration $N(v)\chi_v$ floods the entire grid, where $N(v) = C'n^\beta f(r_1)$. Now we can apply Lemma 6.1. By equation (6.29) and in the light of the convergence of the series $\sum_{r=4}^{\infty} r^2/(r-3)^\alpha$, we have $\sum_{v \in V_0} N(v) = \Theta(n^{\alpha+\beta+1})$ and the result follows. ■

7 The rank of grid configurations

In this section we prove Proposition 3.2 and Corollary 3.3.

Proof. [Proof of Proposition 3.2] We use the concepts discussed in Remark 2.1. Dhar et al. [DRSV] proved that the rank of the sandpile group $\mathcal{G}_n$ associated with the grid $\mathcal{X}_n$ is n. The following is implicit in their proof. Let $v_1, \dots, v_n$ be the sites along one of the edges of the grid. Let $\mathbf{e}$ denote the identity configuration (identity element of $\mathcal{G}_n$). Then the group $\mathcal{G}_n$ is generated by the (recurrent) configurations $\chi_{v_i} \oplus \mathbf{e}$ $(i = 1, \dots, n)$.

Now we claim that all of $\mathcal{G}_n$ is generated by the χ_{v_i} alone. To see this, all we need to observe that χ_{v_1} generates $\mathbf{e}$. Indeed, for some N, the configuration $N\chi_{v_1}$ is recurrent, so it belongs to the group $\mathcal{G}_n$. Let its order in this group be N'; then $\sigma(N'N\chi_{v_1}) = \mathbf{e}$. ■

Proof. [Proof of Corollary 3.3] Let $\mathcal{T}_n(m)$ be the set of (not necessarily stable) transient configurations on $\mathcal{X}_n$ whose support consists of at most m vertices $(1 \le m \le n^2)$. Further, let $\mathcal{T}_n^s$ be the set of stable, transient configurations on $\mathcal{X}_n$. We shall show that for some constant $c > 0$, if $m < cn^2/\log n$ then $|\mathcal{T}_n(m)|/|\mathcal{T}_n^s|$ goes to zero exponentially in n^2. This implies that with extremely high probability, a random stable transient configuration has rank $\ge cn^2/\log n$.

Let $\mathbf{c} \in \mathcal{T}_n(m)$. Then the weight of $\mathbf{c}$ is $\le n^\gamma$ for some constant γ by our main result, Theorem 3.1.

It follows that

$$|\mathcal{T}_n(m)| \le \binom{n^2}{m} (n^\gamma)^m < n^{(\gamma+2)m}. \tag{7.32}$$

Moreover, it is easy to see that $|\mathcal{T}_n^s| = 4^{n^2}(1 - o(1))$. Consequently,

$$\frac{|\mathcal{T}_n(m)|}{|\mathcal{T}_n^s|} \lesssim \frac{n^{(\gamma+2)m}}{2^{2n^2}}. \tag{7.33}$$

So if $m \le n^2/(\gamma+2)\log_2 n$ then $|\mathcal{T}_n(m)|/|\mathcal{T}_n^s| \lesssim 2^{-n^2}$, proving the result. ■

8 Open questions

The obvious open question raised by our main result is to reduce the exponent of 30 in our $O(n^{30})$ estimate of the transience class of the $n \times n$ grid. Computer experiments suggest a much lower exponent, perhaps $O(n^3)$.

We believe that the longest sequence of transient configurations occurs when grains are added at a single site. Although this would only save a factor of n in our case, the question is of interest in its own right, so we state it as a conjecture, not only for the grid but for any (undirected) ambient space.

CONJECTURE 8.1. *Assume the ambient space* $\mathcal{X}$ *induces a connected graph on the set of ordinary vertices. Then the transience class of* $\mathcal{X}$ *(the largest weight of any transient configuration on* $\mathcal{X}$*) is the height of the tallest transient stack of grains placed on a single site.*

CONJECTURE 8.2. *In the case of the* $n \times n$ *grid, the site of the tallest transient stack of grains is a corner of the grid.*

We believe that the $\Omega(n^2/\log n)$ lower bound in Corollary 3.3 is not tight.

CONJECTURE 8.3. *In the* $n \times n$ *grid, almost all stable configurations have rank* $\Omega(n^2)$.

(Note that over the grid, almost all stable configurations are transient, so inserting the adjective "transient" would not change the meaning of this Conjecture.)

PROBLEM 8.4. Find general conditions on the ambient space $\mathcal{X}$ which guarantee that the rank of almost all stable configurations is close to the number of vertices. Decide if this situation is "typical" in some sense.

We know that the rank of a recurrent configuration on the $n \times n$ grid is at most n, but we don't have any lower bound.

Conjecture 8.5. *The rank of almost all recurrent configurations on the $n \times n$ grid is $\Omega(n)$.*

Finally we ask, to what sequences of ambient spaces does the polynomiality of the transience class extend. Our ambient spaces were all derived as certain finite parts of the infinite square grid, with the boundary linked to a sink. Similar constructions can be applied to other locally finite infinite graphs X with high symmetry. Is it the case that if X grows polynomially then the transience class of the corresponding family of finite models is polynomially bounded? Is it the case that for the tilings of the hyperbolic plane, the transience class is exponential?

References

[BT] L. Babai, E. Toumpakari: A structure theory of the sandpile monoid for digraphs. In preparation, see http://people.cs.uchicago.edu/~laci/sandpiles.html

[BGS] L. Babai, I. Gorodezky, A. Shapiro: It's a long way to recurrence: sandpile models with exponentially large transience class. In preparation, see http://people.cs.uchicago.edu/~laci/sandpiles.html

[BTW] P. Bak, C. Tang, K. Wiesenfeld: Self-organized criticality. *Phys. Rev. A* **3** (1988), 364–374.

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

[BLS] A. Björner, L. Lovász, P. W. Shor: Chip-firing games on graphs. *Europ. J. Comb.* **12** (1991), 283–291

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

[DRSV] D. Dhar, P. Ruelle, S. Sen, D. Verma: Algebraic aspects of abelian sandpile models. *Jour. Phys. A* **28** (1995), 805–831

[HLP] A. Holroyd, L. Levine, J. Propp: Quasirandom sampling. In preparation, see [Kl]

[Kl] M. Kleber: Goldbug Variations. *Math. Intelligencer* **27/1** (2005), 55–63

[Je] H. Jensen: *Self-Organized Criticality.* Cambridge Lecture Notes in Physics **10** (1998)

[LR] Y. Le Borgne, D. Rossin: On the identity of the sandpile group. *Discrete Math.* **256** (2002), 775–790

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

[MD] S. N. Majumdar, D. Dhar: Equivalence between the abelian sandpile model and the $q \to 0$ limit of the Potts model. *Physica A* **185** (1992), 129–145

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

Appendix: Figures

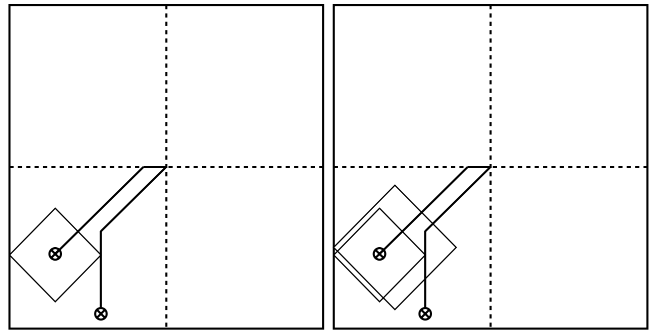

Figure 5: The left image shows that we begin the expansion phase by flooding a constrained diamond around the "departure vertex." In the right image, we continue to flood a sequence of progressively larger constrained diamonds whose centers are on the diagonal segment of the path toward the center of the grid.

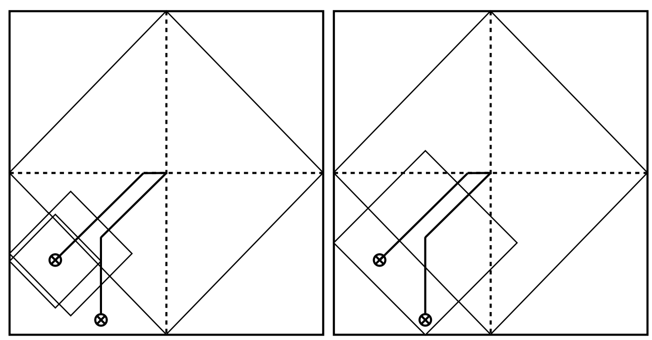

Figure 6: Now we flood a sequence of diamonds whose centers move along the median segment of the path, until finally the central diamond has been flooded (left image). The right image depicts the contraction phase; we flood progressively smaller constrained diamonds whose centers move diagonally towards the destination vertex. Once we have flooded a constrained diamond centered at the end of the diagonal segment of the path, the destination vertex has been touched and we stop.