

Approach to criticality in sandpiles

Anne Fey,¹ Lionel Levine,² and David B. Wilson³¹*Delft Institute of Applied Mathematics, Delft University of Technology, Delft, The Netherlands*²*Department of Mathematics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*³*Microsoft Research, Redmond, Washington 98052, USA*

(Received 16 March 2010; revised manuscript received 30 March 2010; published 15 September 2010)

A popular theory of self-organized criticality predicts that the stationary density of the Abelian sandpile model equals the threshold density of the corresponding fixed-energy sandpile. We recently announced that this “density conjecture” is *false* when the underlying graph is any of \mathbb{Z}^2 , the complete graph K_n , the Cayley tree, the ladder graph, the bracelet graph, or the flower graph. In this paper, we substantiate this claim by rigorous proof and extensive simulations. We show that driven-dissipative sandpiles continue to evolve even after a constant fraction of the sand has been lost at the sink. Nevertheless, we do find (and prove) a relationship between the two models: the threshold density of the fixed-energy sandpile is the point at which the driven-dissipative sandpile begins to lose a macroscopic amount of sand to the sink.

DOI: [10.1103/PhysRevE.82.031121](https://doi.org/10.1103/PhysRevE.82.031121)

PACS number(s): 64.60.av, 45.70.Cc

I. INTRODUCTION

In this paper, we expand on our results announced in [1] critiquing the theory of self-organized criticality developed by Dickman, Muñoz, Vespignani, and Zapperi (DMVZ) in a series of widely cited papers [2–6]. The DMVZ theory predicts a certain relationship between systems that are driven from the outside and closed systems with an absorbing state. We refute this prediction for the Abelian sandpile model of Bak, Tang, and Wiesenfeld [7] and its fixed-energy counterpart. In particular, we focus on the prediction that the stationary density ζ_s of the driven-dissipative sandpile model equals the threshold density ζ_c of the fixed-energy sandpile model (FES).

For several families of graphs, we have found precise values for both these densities, which are clearly not equal. We presented these values [1]; for completeness, we reproduce the table here (Table I). In this paper, we present our evidence for these values, which consists either of rigorous proof or extensive simulations. Our rigorous results (see Theorems 4 and 14) point to a somewhat different relationship than that posited in the DMVZ series of papers: the driven system exhibits a second-order phase transition at the threshold density of the closed system.

One hope of the DMVZ paradigm was that critical features of the driven-dissipative model, such as the exponents governing the distribution of avalanche sizes and decay of correlations, might be more easily studied in FES by examining the scaling behavior of these observables as $\zeta \uparrow \zeta_c$. However, several findings including ours suggest that these two models may not share the same critical features. Among these we note the discrepancies reported by Grassberger and Manna [8]; the discovery by De Menech, Stella, and Tebaldi that many observables of the driven-dissipative model do not show simple power-law scaling [9]; the finding of Bagnoli *et al.* [10] of nonergodicity in the FES; and the work of Peters and Pruessner [11,12], who numerically find different critical properties for driven and fixed-energy versions of the Ising model and the Oslo model. Our main findings—the inequality of ζ_s and ζ_c , and the continued change in density of

driven-dissipative sandpiles beyond ζ_c —constitute further evidence that driven-dissipative and fixed-energy sandpile models may not share the same critical features.

This paper is organized as follows: in Sec. II we define the two sandpile models, with the square grid graph \mathbb{Z}^2 as example. We present simulation results supplementing those in [1]. In the remaining sections, we discuss the other graph families. In Secs. III (bracelet graph), IV (complete graph), and V (flower graph), we give rigorous proofs for the exact values of the two densities. Moreover, in Sec. III we give the proof of Theorem 1 of [1], and in Sec. V of a similar theorem, both illustrated graphically in Fig. 2 of [1]. In Secs. VI (regular trees) and VII (ladder graph), we find the threshold densities by simulation. For the exact values of the stationary densities, we refer to the work of Jeng, Piroux, and Ruelle [13], Dhar and Majumdar [14], and Járai and Lyons [15].

II. SANDPILES ON THE SQUARE GRID \mathbb{Z}^2

In this section we give precise definitions of the stationary and threshold densities, and present the results of large-scale simulations on \mathbb{Z}^2 . The definitions in this section apply to general graphs, but we defer the discussion of results about other graphs to subsequent sections.

TABLE I. Stationary and threshold densities for different graphs. Exact values are in bold.

Graph	ζ_s	ζ_c
\mathbb{Z}	1	1
\mathbb{Z}^2	17/8 =2.125	2.125288...
Bracelet	5/2 =2.5	2.496608 ...
Flower graph	5/3 =1.666667...	1.668898 ...
Ladder graph	$\frac{7}{4} - \frac{\sqrt{3}}{12}$ =1.605662...	1.6082...
Complete graph	$n/2 + O(\sqrt{n})$	$n - O(\sqrt{n} \log n)$
3-regular tree	3/2	1.50000...
4-regular tree	2	2.00041...
5-regular tree	5/2	2.51167...

A. Driven-dissipative sandpiles and the stationary density ζ_s

Let $\hat{G}=(V, E)$ be a finite graph, which may have loops and multiple edges. Let $S \subset V$ be a nonempty set of vertices, which we will call *sinks*. The presence of sinks distinguishes the driven-dissipative sandpile from its fixed-energy counterpart. To highlight this distinction, throughout the paper, graphs denoted with a “hat” as in \hat{G} have sinks, and those without a hat as in G do not.

For vertices $v, w \in V$, write $a_{v,w}=a_{w,v}$ for the number of edges connecting v and w , and

$$d_v = \sum_{w \in V} a_{v,w}.$$

for the number of edges incident to v . A *sandpile* (or “configuration”) η on \hat{G} is a map

$$\eta: V \rightarrow \mathbb{Z}_{\geq 0}.$$

We interpret $\eta(v)$ as the number of sand particles at the vertex v ; we will sometimes call this number the *height* of v in η .

A vertex $v \notin S$ is called *unstable* if $\eta(v) \geq d_v$. An unstable vertex can *topple* by sending one particle along each edge incident to v . Thus, toppling v results in a new sandpile η' given by

$$\eta' = \eta + \Delta_v,$$

where

$$\Delta_v(w) = \begin{cases} a_{v,w}, & v \neq w \\ a_{v,v} - d_v, & v = w. \end{cases}$$

Sinks by definition are always stable, and never topple. If all vertices are stable, we say that η is stable.

Note that toppling a vertex may cause some of its neighbors to become unstable. The *stabilization* η° of η is a sandpile resulting from toppling unstable vertices in sequence, until all vertices are stable. By the *Abelian property* [16], the stabilization is unique: it does not depend on the toppling sequence. Moreover, the number of times a given vertex topples does not depend on the toppling sequence.

The most commonly studied example is the $n \times n$ square grid graph, with the boundary sites serving as sinks (Fig. 1). The driven-dissipative sandpile model is a continuous time Markov chain $(\eta_t)_{t \geq 0}$ whose state space is the set of stable sandpiles on \hat{G} . Let $V' = V \setminus S$ be the set of vertices that are not sinks. At each site $v \in V'$, particles are added at rate 1. When a particle is added, topplings occur instantaneously to stabilize the sandpile. Writing $\sigma_t(v)$ for the total number of particles added at v before time t , we have by the Abelian property

$$\eta_t = (\sigma_t)^\circ.$$

Note that for fixed t , the random variables $\sigma_t(v)$ for $v \in V'$ are independent and have the Poisson distribution with mean t .

The model just described is most commonly known as the Abelian sandpile model (ASM), but we prefer the term “driven dissipative” to distinguish it from the fixed-energy




FIG. 1. The square grid Z^2 .

sandpile described below, which is also a form of ASM. “Driven” refers to the addition of particles, and “dissipative” to the loss of particles absorbed by the sinks.

Dhar [16] developed the *burning algorithm* to characterize the recurrent sandpile states, that is, those sandpiles η for which, regardless of the initial state,

$$\text{Prob}(\eta_t = \eta \text{ for some } t) = 1.$$

Lemma 1 (Burning Algorithm [16]). A sandpile η is recurrent if and only if every nonsink vertex topples exactly once during the stabilization of $\eta + \sum_{s \in S} \Delta_s$, where the sum is over sink vertices S .

The recurrent states form an Abelian group under the operation of addition followed by stabilization. In particular, the stationary distribution of the Markov chain η_t is uniform on the set of recurrent states.

The combination of driving and dissipation organizes the system into a critical state. To measure the density of particles in this state, we define the *stationary density* $\zeta_s(\hat{G})$ as

$$\zeta_s(\hat{G}) = \mathbb{E}_\mu \left[\frac{1}{|V'|} \sum_{v \in V'} \eta(v) \right],$$

where $V' = V \setminus S$, and μ is the uniform measure on recurrent sandpiles on \hat{G} . The stationary density has another expression in terms of the Tutte polynomial of the graph obtained from \hat{G} by collapsing the set S of sinks to a single vertex; see Sec. IV.

Most of the graphs we will study arise naturally as finite subsets of infinite graphs. Let Γ be a countably infinite graph in which every vertex has finite degree. Let \hat{G}_n for $n \geq 1$ be a nested family of finite induced subgraphs with $\cup \hat{G}_n = \Gamma$. As sinks in \hat{G}_n we take the set of boundary vertices

$$S_n = \hat{G}_n - \hat{G}_{n-1}.$$

In cases where the free and wired limits are different, such as on regular trees, we will choose a sequence \hat{G}_n corresponding to the wired limit. We denote by μ_n the uniform measure on recurrent configurations on \hat{G}_n .

We are interested in the *stationary density*

$$\zeta_s(\Gamma) := \lim_{n \rightarrow \infty} \zeta_s(\hat{G}_n).$$

When $\Gamma = \mathbb{Z}^d$, it is known that the infinite-volume limit of measures $\mu = \lim_{n \rightarrow \infty} \mu_n$ exists and is translation-invariant [17]. In this case it follows that the limit defining $\zeta_s(\Gamma)$ exists and equals

$$\zeta_s = \mathbb{E}_\mu[\eta(0)],$$

where $0 \in \mathbb{Z}^d$ is the origin. For other families of graphs we consider, we will show that the limit defining $\zeta_s(\Gamma)$ exists.

Much is known about the limiting measure μ in the case $\Gamma = \mathbb{Z}^2$. The following expressions have been obtained for ζ_s and the single site height probabilities. The symbol $\stackrel{?}{=}$ denotes expressions that are rigorous up to a conjecture [13] that a certain integral, numerically evaluated as 0.5 ± 10^{-12} , is exactly $1/2$.

$$\zeta_s(\mathbb{Z}^2) \stackrel{?}{=} 17/8 \quad (\text{Ref. [13]}),$$

$$\mu\{\eta(x) = 0\} = \frac{2}{\pi^2} - \frac{4}{\pi^3} \quad (\text{Ref. [18]}),$$

$$\mu\{\eta(x) = 1\} \stackrel{?}{=} \frac{1}{4} - \frac{1}{2\pi} - \frac{2}{\pi^2} + \frac{12}{\pi^3} \quad (\text{Refs. [13,19]}),$$

$$\mu\{\eta(x) = 2\} \stackrel{?}{=} \frac{3}{8} + \frac{1}{\pi} - \frac{12}{\pi^3} \quad (\text{Ref. [13]}), \quad \text{and}$$

$$\mu\{\eta(x) = 3\} \stackrel{?}{=} \frac{3}{8} - \frac{1}{2\pi} + \frac{1}{\pi^2} + \frac{4}{\pi^3} \quad (\text{Ref. [13]}).$$

The equality $\zeta_s(\mathbb{Z}^2) \stackrel{?}{=} 17/8$ was first conjectured by Grassberger.

B. Fixed-energy sandpiles and the threshold density ζ_c

Next we describe the fixed-energy sandpile model, in which the driving and dissipation are absent, and the total number of particles is conserved. In the mathematical and computer science literature, this model goes by the name *parallel chip-firing* [20–22]. As before, let G be a finite graph, possibly with loops and multiple edges. Unlike the driven-dissipative model, we do not single out any vertices as sinks. The fixed-energy sandpile evolves in discrete time: at each time step, all unstable vertices topple once in parallel. Thus the configuration η_{j+1} at time $j+1$ is given by

$$\eta_{j+1} = \eta_j + \sum_{v \in U_j} \Delta_v,$$

where

$$U_j = \{v \in V: \eta_j(v) \geq d_v\}$$

is the set of vertices that are unstable at time j . We say that η_0 *stabilizes* if toppling eventually stops, i.e., $U_j = \emptyset$ for all sufficiently large j .

If η_0 stabilizes, then there is some site that never topples [23] (see also [[24], Theorem 2.8, item 4] and [[25], Lemma 2.2] for the case when G is infinite). Otherwise, for each site x , let $j(x)$ be the last time x topples. Choose a site x minimizing $j(x)$. Then each neighbor y of x has $j(y) \geq j(x)$, so y topples at least once at or after time $j(x)$. Thus x receives at least d_x additional particles and must topple again after time $j(x)$, a contradiction. Note that this argument uses in a crucial way the deterministic nature of the toppling rule. It gives a criterion that is very useful in simulations: as soon as every site has toppled at least once, we know that the system will not stabilize.

Let $[\sigma_\lambda(v)]_{\lambda \geq 0}$ be a collection of independent Poisson point processes of intensity 1, indexed by the vertices of G . So each $\sigma_\lambda(v)$ has the Poisson distribution with mean λ . We define the *threshold density* of G as

$$\zeta_c(G) = \mathbb{E}\Lambda_c,$$

where

$$\Lambda_c = \sup\{\lambda: \sigma_\lambda \text{ stabilizes}\}.$$

We expect that Λ_c is tightly concentrated around its mean when G is large. Indeed, if Γ is an infinite vertex-transitive graph, then the event that σ_λ stabilizes on Γ is translation-invariant. By the ergodicity of the Poisson product measure, this event has probability 0 or 1. Since this probability is monotone in λ , there is a (deterministic) threshold density $\zeta_c(\Gamma)$, such that

$$\text{Prob}[\sigma_\lambda \text{ stabilizes on } \Gamma] = \begin{cases} 1, & \lambda < \zeta_c(\Gamma) \\ 0, & \lambda > \zeta_c(\Gamma). \end{cases}$$

We expect the threshold densities on natural families of finite graphs to satisfy a law of large numbers such as the following.

Conjecture 2. With probability 1,

$$\Lambda_c(\mathbb{Z}_n^2) \rightarrow \zeta_c(\mathbb{Z}^2) \quad \text{as } n \rightarrow \infty.$$

DMVZ believed that the combination of driving and dissipation in the classical Abelian sandpile model should push it toward the critical density ζ_c of the fixed-energy sandpile. This leads to a specific testable prediction, which we call the *Density Conjecture*.

Conjecture 3 (Density Conjecture).

$$\zeta_c = \zeta_s.$$

In the case of the square grid, the conjecture $\zeta_c = 17/8$ can be found in [5]. Likewise, in [6], it is asserted that ‘‘FES are found to be critical only for a particular value $\zeta = \zeta_c$ (which as we will show turns out to be identical to the stationary energy density of its driven-dissipative counterpart).’’

Previous simulations ($n = 160$ [2]; $n = 1280$ [3]) to estimate the threshold density $\zeta_c(\mathbb{Z}^2)$ found a value of 2.125, in agreement with the stationary density $\zeta_s(\mathbb{Z}^2) \stackrel{?}{=} 17/8$. By performing larger-scale simulations, however, we find that ζ_c exceeds ζ_s .

Table II summarizes the results of our simulations refuting the density conjecture on \mathbb{Z}^2 . We find that $\zeta_c(\mathbb{Z}^2)$ equals 2.125 288 to six decimal places, whereas $\zeta_s(\mathbb{Z}^2)$ is known to be 2.125 000 000 000 to 12 decimal places. In each random

TABLE II. Fixed-energy sandpile simulations on $n \times n$ tori \mathbb{Z}_n^2 . The third column gives our empirical estimate of the threshold density $\zeta_c(\mathbb{Z}_n^2)$. The next four columns give the empirical distribution of the height of a fixed vertex in the stabilization $(\sigma_\lambda)^\circ$, for λ just below Λ_c . Each estimate of the expectation $\zeta_c(\mathbb{Z}_n^2)$ and of the marginals $\text{Prob}[h=i]$ has standard deviation less than 4×10^{-7} . The total number of topplings needed to stabilize σ_λ appears to scale as n^3 .

Grid size (n^2)	No. samples	$\zeta_c(\mathbb{Z}_n^2)$	Distribution of height h of sand				(No. topplings)/ n^3
			Prob[$h=0$]	Prob[$h=1$]	Prob[$h=2$]	Prob[$h=3$]	
64^2	268435456	2.124956	0.073555	0.173966	0.306447	0.446032	0.197110
128^2	67108864	2.125185	0.073505	0.173866	0.306567	0.446062	0.197808
256^2	16777216	2.125257	0.073488	0.173835	0.306609	0.446068	0.198789
512^2	4194304	2.125279	0.073481	0.173826	0.306626	0.446067	0.200162
1024^2	1048576	2.125285	0.073479	0.173822	0.306633	0.446066	0.201745
2048^2	262144	2.125288	0.073478	0.173821	0.306635	0.446065	0.203378
4096^2	65536	2.125288	0.073477	0.173821	0.306637	0.446064	0.205323
8192^2	16384	2.125288	0.073477	0.173821	0.306638	0.446064	0.206475
16384^2	4096	2.125288	0.073478	0.173821	0.306638	0.446064	0.208079
\mathbb{Z}^2 (stationary)		2.125000	0.073636	0.173900	0.306291	0.446172	

trial, we add particles one at a time at uniformly random sites of the $n \times n$ torus. After each addition, we perform topplings until either all sites are stable, or every site has toppled at least once since the last addition. In the latter case, the sandpile does not stabilize. We record m/n^2 as an empirical estimate of the threshold density, where m is the maximum number of particles for which the configuration stabilizes. We then average these empirical estimates over many independent trials. The one-site marginals we report are obtained from the stable configuration just before the $(m+1)^{\text{st}}$ particle was added, and the number of topplings reported is the total number of topplings required to stabilize the first m particles.

We used a random number generator based on the Advanced Encryption Standard (AES-256), which has been found to exhibit good statistical properties. Our simulations were conducted on a high performance computing (HPC) cluster of computers.

III. SANDPILES ON THE BRACELET

Next we examine a family of graphs for which we can determine ζ_c and ζ_s exactly and prove that they are not equal. Despite this inequality, we show that an interesting connection remains between the driven-dissipative and conservative dynamics: the threshold density of the conservative model is the point at which the driven-dissipative model begins to lose a macroscopic amount of sand to the sink.

The *bracelet graph* B_n (Fig. 2) is a multigraph with vertex set \mathbb{Z}_n (the n -cycle) with the usual edge set $\{(i, i+1 \bmod n) : 0 \leq i < n\}$ doubled. Thus all vertices have degree 4. The graph \hat{B}_n is the same, except that vertex 0 is distinguished as a sink from which particles disappear from the system. We denote by B_∞ the infinite path \mathbb{Z} with doubled edges.

For $\lambda > 0$, let σ_λ be the configuration with $\text{Poisson}(\lambda)$ particles independently on each site of \hat{B}_n . Let $\eta_\lambda = (\sigma_\lambda)^\circ$ be the stabilization of σ_λ , and let

$$\rho_n(\lambda) = \frac{1}{n-1} \sum_{x=1}^{n-1} \eta_\lambda(x)$$

be the final density. The following theorem gives the threshold and stationary densities of the infinite bracelet graph B_∞ , and identifies the $n \rightarrow \infty$ limit of the final density $\rho_n(\lambda)$ as a function of the initial density λ .

Theorem 4. For the bracelet graph,

- (1) The threshold density $\zeta_c(B_\infty)$ is the unique positive root of $\zeta = \frac{5}{2} - \frac{1}{2}e^{-2\zeta}$ (numerically, $\zeta_c = 2.496\ 608$).
- (2) The stationary density $\zeta_s(B_\infty)$ is $5/2$.
- (3) $\rho_n(\lambda) \rightarrow \rho(\lambda)$ in probability as $n \rightarrow \infty$, where

$$\rho(\lambda) = \min\left(\lambda, \frac{5 - e^{-2\lambda}}{2}\right) = \begin{cases} \lambda, & \lambda \leq \zeta_c \\ \frac{5 - e^{-2\lambda}}{2}, & \lambda > \zeta_c. \end{cases}$$

Part 3 of this theorem shows that the final density undergoes a second-order phase transition at ζ_c : the derivative of




FIG. 2. The bracelet graph B_{20} .




FIG. 3. Density $\rho(\lambda)$ of the final stable configuration as a function of initial density λ , for the driven sandpile on the bracelet graph \hat{B}_n as $n \rightarrow \infty$. A second-order phase transition occurs at $\lambda = \zeta_c$. Beyond this transition, the density of the driven sandpile continues to increase, approaching the stationary density ζ_s from below.

$\rho(\lambda)$ is discontinuous at $\lambda = \zeta_c$ (Fig. 3). Thus in spite of the fact that $\zeta_s \neq \zeta_c$, there remains a connection between the conservative dynamics used to define ζ_c and the driven-dissipative dynamics used to define ζ_s . For $\lambda < \zeta_c$, very little dissipation takes place, so the final density equals the initial density λ ; for $\lambda > \zeta_c$ a substantial amount of dissipation takes place, many particles are lost to the sink, and the final density is strictly less than the initial density. The sandpile continues to evolve as λ increases beyond ζ_c ; in particular its density keeps changing.

We believe that this phenomenon is widespread. As evidence, in Sec. V we introduce the “flower graph,” which looks very different from the bracelet, and prove (in Theorem 14) that a similar phase transition takes place there.

For the proof of Theorem 4, we compare the dynamics of pairs of particles on the bracelet graph to single particles on \mathbb{Z} . At each vertex x of the bracelet, we group the particles starting at x into pairs, with one “passive” particle left over if $\sigma_\lambda(x)$ is odd. Since all edges in the bracelet are doubled, we can ensure that in each toppling the two particles comprising a pair always move to the same neighbor, and that the passive particles never move. The toppling dynamics of the pairs are equivalent to the usual Abelian sandpile dynamics on \mathbb{Z} .

We recall the relevant facts about one-dimensional sandpile dynamics:

(i) In any recurrent configuration on a finite interval of \mathbb{Z} , every site has height 1, except for at most one site of height 0. Therefore, $\zeta_s = 1$ [26].

(ii) On \mathbb{Z} , an initial configuration distributed according to a nontrivial product measure with mean λ stabilizes almost surely (every site topples only finitely many times) if $\lambda < 1$, while it almost surely does not stabilize (every site topples infinitely often) if $\lambda \geq 1$ [24]. Thus, $\zeta_c = 1$.

Proof of Theorem 4 parts 1 and 2. Given $\lambda > 0$, let λ^* be the pair density $\mathbb{E}[\sigma_\lambda(x)/2]$, and let

$$p_{\text{odd}}(\lambda) = e^{-\lambda} \sum_{m \geq 0} \frac{\lambda^{2m+1}}{(2m+1)!} = \frac{1}{2}(1 - e^{-2\lambda}).$$

be the probability that a $\text{Poisson}(\lambda)$ random variable is odd. Then λ and λ^* are related by

$$\lambda = 2\lambda^* + p_{\text{odd}}(\lambda). \tag{1}$$

The configuration σ_λ stabilizes on B_∞ if and only if the pair configuration σ_λ^* stabilizes on \mathbb{Z} . Thus $\zeta_c(B_\infty)^* = \zeta_c(\mathbb{Z})$. Setting $\lambda = \zeta_c(B_\infty)$ in Eq. (1), using the fact that $\zeta_c(\mathbb{Z}) = 1$, and that λ^* is an increasing function of $\lambda > 0$, we conclude that $\zeta_c(B_\infty)$ is the unique positive root of

$$\zeta = 2 + p_{\text{odd}}(\zeta),$$

or $\zeta = \frac{5}{2} - \frac{1}{2}e^{-2\zeta}$. This proves part 1.

For part 2, by the burning algorithm, a configuration σ on \hat{B}_n is recurrent if and only if it has at most one site with fewer than two particles. Thus, in the uniform measure on recurrent configurations on \hat{B}_n ,

$$\text{Prob}(\sigma(x) = 2) = \text{Prob}(\sigma(x) = 3) = \frac{1}{2} - \frac{1}{2n},$$

$$\text{Prob}(\sigma(x) = 0) = \text{Prob}(\sigma(x) = 1) = \frac{1}{2n}.$$

We conclude that $\zeta_s(\hat{B}_n) = \mathbb{E}\sigma(x) = \frac{5}{2} - \frac{2}{n} \rightarrow \frac{5}{2}$ as $n \rightarrow \infty$.

To prove part 3 of Theorem 4, we use the following lemma, whose proof is deferred to the end of this section. Let \hat{Z}_n be the n -cycle with vertex 0 distinguished as a sink. Let σ'_λ be a sandpile on \hat{Z}_n distributed according to a product measure (not necessarily Poisson) of mean λ . Let η'_λ be the stabilization of σ'_λ , and let $\rho'_n(\lambda) = \frac{1}{n-1} \sum_{x=1}^{n-1} \eta'_\lambda(x)$ be the final density after stabilization.

Lemma 5. On \hat{Z}_n , we have $\rho'_n(\lambda) \rightarrow \min(\lambda, 1)$ in probability.

Proof of Theorem 4, part 3. Let η_λ be the stabilization of σ_λ on \hat{B}_n , and let η_λ^* be the stabilization of $\sigma_\lambda^* = \lfloor \sigma_\lambda / 2 \rfloor$ on \hat{Z}_n . Then

$$\eta_\lambda(x) = 2\eta_\lambda^*(x) + \omega_\lambda(x), \tag{2}$$

where $\omega_\lambda(x) = \sigma_\lambda(x) - 2\sigma_\lambda^*(x)$ is 1 or 0 accordingly as $\sigma_\lambda(x)$ is odd or even. Let

$$\rho_n^*(\lambda) = \frac{1}{n-1} \sum_{x=1}^{n-1} \eta_\lambda^*(x)$$

be the final density after stabilization of σ_λ^* on \hat{Z}_n . Then

$$\rho_n(\lambda) = 2\rho_n^*(\lambda) + \frac{1}{n-1} \sum_{x=1}^{n-1} \omega_\lambda(x).$$

By the weak law of large numbers, $\frac{1}{n-1} \sum_{x=1}^{n-1} \omega_\lambda(x) \rightarrow p_{\text{odd}}(\lambda)$ in probability as $n \rightarrow \infty$. If $\lambda < \zeta_c$, then $\lambda^* < 1$, so by Lemma 5, $\rho_n^*(\lambda) \rightarrow \lambda^*$ in probability, and hence

$$\rho_n(\lambda) \rightarrow 2\lambda^* + p_{\text{odd}}(\lambda) = \lambda$$

in probability. If $\lambda \geq \zeta_c$, then $\lambda^* \geq 1$, so by Lemma 5, $\rho_n^*(\lambda) \rightarrow 1$ in probability, hence

$$\rho_n(\lambda) \rightarrow 2 + p_{\text{odd}}(\lambda) = \frac{5 - e^{-2\lambda}}{2}$$

in probability. This proves part 3. ■

Proof of Lemma 5. We may view \hat{Z}_n as the path in \mathbb{Z} from $a_n = -\lfloor n/2 \rfloor$ to $b_n = \lfloor n/2 \rfloor$, with both endpoints serving as sinks. For $x \in \hat{Z}_n$, let $u_n(x)$ be the number of times that x topples during stabilization of the configuration σ'_λ on \hat{Z}_n . Let $u_\infty(x)$ be the number of times x topples during stabilization of σ'_λ on \mathbb{Z} . The procedure of ‘‘toppling in nested volumes’’ [24] shows that $u_n(x) \uparrow u_\infty(x)$ as $n \rightarrow \infty$.

We consider first $\lambda < 1$. In this case $u_\infty(x)$ is finite almost surely (a.s.). The total number of particles lost to the sinks on \hat{Z}_n is $u_n(a_n + 1) + u_n(b_n - 1)$, so the final density is given by

$$\rho'_n(\lambda) = \frac{1}{n-1} \left[\sum_{x=a_n+1}^{b_n-1} \sigma_\lambda(x) - u_n(a_n + 1) - u_n(b_n - 1) \right].$$

By the law of large numbers, $\frac{1}{n-1} \sum \sigma_\lambda(x) \rightarrow \lambda$ in probability as $n \rightarrow \infty$. Since $u_\infty(x)$ is a.s. finite, we have $\frac{u_n(a_n+1)+u_n(b_n-1)}{n-1} \rightarrow 0$ in probability, so $\rho'_n(\lambda) \rightarrow \lambda$ in probability.

Next we consider $\lambda \geq 1$. In this case we have $u_n(x) \uparrow u_\infty(x) = \infty$, a.s. Let $p(n, x) = \text{Prob}[u_n(x) = 0]$ be the probability that $x \in \hat{Z}_n$ does not topple. By the Abelian property, adding sinks cannot increase the number of topplings, so

$$p(n, x) \leq p(m, 0),$$

where $m = \min(x - a_n, b_n - x)$. Let

$$Y_n = \sum_{x=a_n+1}^{b_n-1} 1_{\{u_n(x)=0\}}$$

be the number of sites in \hat{Z}_n that do not topple. Since $u_n(0) \uparrow \infty$ a.s., we have $p(n, 0) \downarrow 0$, hence

$$\mathbb{E} \frac{Y_n}{n} = \frac{1}{n} \sum_{x=a_n+1}^{b_n-1} p(n, x) \leq \frac{2}{n} \sum_{m=1}^{n/2} p(m, 0) \rightarrow 0$$

as $n \rightarrow \infty$. Since $Y_n \geq 0$ it follows that $Y_n/n \rightarrow 0$ in probability.

In an interval where every site toppled, there can be at most one empty site. We have $Y_n + 1$ such intervals. Therefore, the number of empty sites is at most $2Y_n + 1$. Hence

$$\frac{n - 2Y_n - 2}{n - 1} \leq \rho'_n(\lambda) \leq 1.$$

The left side tends to 1 in probability, which completes the proof.

IV. SANDPILES ON THE COMPLETE GRAPH

Let K_n be the complete graph on n vertices: every pair of distinct vertices is connected by an edge. In \hat{K}_n , one vertex is distinguished as the sink. The maximal stable configuration on \hat{K}_n has density $n - 2$, while the minimal recurrent configurations have exactly one vertex of each height $0, 1, \dots, n - 2$, hence density $\frac{n-2}{2}$. The following result shows that the stationary and threshold densities are quite far apart: ζ_s is close to the minimal recurrent density, while ζ_c is close to the maximal stable density.

Theorem 6.

$$\zeta_s(\hat{K}_n) = \frac{n}{2} + O(\sqrt{n}),$$

$$\zeta_c(K_n) \geq n - O(\sqrt{n \log n}).$$

The proof uses an expression for the stationary density ζ_s in terms of the Tutte polynomial, due to Merino L3pez [27]. Our application will be to the complete graph, but we state Merino L3pez’ theorem in full generality. Let $G = (V, E)$ be a connected undirected graph with n vertices and m edges. Let v be any vertex of G , and write \hat{G} for the graph G with v distinguished as a sink. Let d be the degree of v .

Recall that the Tutte polynomial $T_G(x, y)$ is defined by

$$T_G(x, y) = \sum_{A \subseteq E} (x - 1)^{c(A) - c(E)} (y - 1)^{c(A) + |A| - |V|},$$

where $c(A)$ denotes the number of connected components of the spanning subgraph (V, A) .

Theorem 7 ([27]). The Tutte polynomial $T_G(x, y)$ evaluated at $x = 1$ is given by

$$T_G(1, y) = y^{d-m} \sum_{\sigma} y^{|\sigma|}$$

where the sum is over all recurrent sandpile configurations σ on \hat{G} , and $|\sigma|$ denotes the number of particles in σ .

Differentiating and evaluating at $y = 1$, we obtain

$$\left. \frac{d}{dy} T_G(1, y) \right|_{y=1} = \sum_{\sigma} (d - m + |\sigma|). \tag{3}$$

Referring to the definition of the Tutte polynomial, we see that $T_G(1, 1)$ is the number of spanning trees of G , and that the left side of Eq. (3) is the number of spanning unicyclic subgraphs of G . (In evaluating T_G at $x = y = 1$, we interpret 0^0 as 1.) The number of recurrent configurations equals the number of spanning trees of G , so the stationary density ζ_s may be expressed as

$$\zeta_s(\hat{G}) = \frac{1}{n T_G(1, 1)} \sum_{\sigma} |\sigma|.$$

Combining these expressions yields the following:

Corollary 8.

$$\zeta_s(\hat{G}) = \frac{1}{n} \left(m - d + \frac{u(G)}{\kappa(G)} \right),$$

where $\kappa(G)$ is the number of spanning trees of G , and $u(G)$ is the number of spanning unicyclic subgraphs of G .

Note that $m - d$ is the minimum number of particles in a recurrent configuration, so the ratio $u(G)/\kappa(G)$ can be interpreted as the average number of excess particles in a recurrent configuration.

Everything so far applies to general connected graphs G . The following is specific for the complete graph.

Theorem 9 (Wright [28]). The number of spanning unicyclic subgraphs of K_n is

$$u(K_n) = \left(\sqrt{\frac{\pi}{8}} + o(1) \right) n^{n-1/2}.$$

Proof of Theorem 6. For \hat{K}_n we have

$$m - d = \frac{n(n-1)}{2} - (n-1) = \frac{(n-2)(n-1)}{2}.$$

From Corollary 8, Theorem 9, and Cayley’s formula $\kappa(K_n) = n^{n-2}$, we obtain

$$\zeta_s(\hat{K}_n) = \frac{1}{n} \left(\frac{(n-2)(n-1)}{2} + \frac{u(K_n)}{\kappa(K_n)} \right) = \frac{n}{2} + \left(\sqrt{\frac{\pi}{8}} + o(1) \right) \sqrt{n}.$$

On the other hand, if we let

$$\lambda = n - 2\sqrt{n \log n}$$

and start with $\sigma(v) \sim \text{Poisson}(\lambda)$ particles at each vertex v of K_n , then for all v

$$\text{Prob}[\sigma(v) \geq n] < \frac{1}{n^2}.$$

So

$$\text{Prob}[\sigma(v) \geq n \text{ for some } v] < \frac{1}{n};$$

in other words, with high probability no topplings occur at all. Thus

$$\text{Prob}(\Lambda_c(K_n) \geq n - 2\sqrt{n \log n}) > 1 - \frac{1}{n}$$

which completes the proof.

One might guess that the large gap between ζ_s and ζ_c is related to the small diameter of \hat{K}_n : since the sink is adjacent to every vertex, its effect is felt with each and every toppling. This intuition is misleading, however, as shown by the lollipop graph \hat{L}_n consisting of K_n connected to a path of length n , with the sink at the far end of the path. Since L_n has the same number of spanning trees and unicyclic subgraphs as K_n , we have by Corollary 8

$$\zeta_s(\hat{L}_n) = \frac{1}{2n} \left(\frac{n(n-1)}{2} + n - 1 + \frac{u(L_n)}{\kappa(L_n)} \right) = \frac{n}{4} + O(\sqrt{n}).$$

On the other hand, by first stabilizing the vertices on the path, close to half of which end up in the sink without reaching the K_n , it is easy to see that with high probability

$$\Lambda_c(L_n) \geq \frac{2n}{3} - O(\sqrt{n \log n}).$$

V. SANDPILES ON THE FLOWER GRAPH

An interesting feature of parallel chip-firing is that further phase transitions appear above the threshold density ζ_c . On a finite graph $G=(V,E)$, since the time evolution is deterministic, the system will eventually reach a periodic orbit: for




FIG. 4. The flower graph F_{20} .

some positive integer m , we have $\eta_{t+m} = \eta_t$ for all sufficiently large t . The *activity density*, ρ_a , measures the proportion of vertices that topple in an average time step,

$$\rho_a(\lambda) = \mathbb{E}_\lambda \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} \frac{1}{|V|} \sum_{x \in V} \mathbf{1}_{\{\eta_s(x) \geq d_x\}}.$$

The expectation \mathbb{E}_λ refers to the initial state η_0 , which we take to be distributed according to the Poisson product measure with mean λ . Note that the limit in the definition of ρ_a can also be expressed as a finite average, due to the eventual periodicity of the dynamics.

Bagnoli *et al.* [10] observed that ρ_a tends to increase with λ in a sequence of flat steps punctuated by sudden jumps. This “devil’s staircase” phenomenon is so far explained only on the complete graph [20]: the number of flat stairs increases with n , and in the $n \rightarrow \infty$ limit there is a stair at each rational number height $\rho_a = p/q$.

On the cycle \mathbb{Z}_n [29] there are just two jumps: at $\lambda=1$, the activity density jumps from 0 to 1/2, and at $\lambda=2$, from 1/2 to 1. For the $n \times n$ torus, simulations [10] indicate a devil’s staircase, which is still not completely understood despite much effort [30].

In this section we study the “flower” graph, which was designed with parallel chip-firing in mind: the idea is that a graph with only short cycles should give rise to short period orbits under the parallel chip-firing dynamics. We find that there are four activity density jumps (Theorem 13). In addition, we determine the stationary and threshold densities of the flower graph, and find a second-order phase transition at ζ_c (Theorem 14).

The flower graph F_n consists of a central site together with $n \geq 1$ petals (Fig. 4). Each petal consists of two sites connected by an edge, each connected to the central site by an edge. Thus the central site has degree $2n$, and all other sites have degree 2. The number of sites is $2n+1$. The graph \hat{F}_n is the same, except one petal serves as sink.

Recall that we defined the density of a configuration as the total number of particles, divided by the total number of sites. Since the flower graph is not regular, the central site has a different expected number of particles than the petal sites.

Proposition 10. For parallel chip-firing on the flower

graph F_n , every configuration has eventual period at most 3.

The proof uses the following two lemmas. First recall that a directed graph is called Eulerian if each vertex has its in-degree equal to its out-degree. In particular, undirected graphs are Eulerian.

Lemma 11 ([31] Lemma 2.5). For parallel chip firing on an Eulerian directed graph, if the eventual period is not 1, then after some finite time, every site x has height at most $2d_x - 1$.

We also use an observation from [10] (see also [31] Proposition 6.2); it was stated and proved there for \mathbb{Z}^2 , but the same proof works for general Eulerian graphs.

Lemma 12 ([10,31]). For an Eulerian directed graph, let two height configurations η and ξ be “mirror images” of each other, that is, $\eta(x) = 2d_x - 1 - \xi(x)$ for all x . Then after performing a parallel chip-firing update on each, the two configurations are again mirror images of each other.

Proof of Proposition 10. Suppose at time t the model has settled into periodic orbit, and the period is not 1. Then by Lemma 11, at time t every petal site has height at most 3. Say that a petal is in state ij if it has i particles at one site and j particles at the other site (we do not distinguish between the states ij and ji). *A priori* there are ten possible petal states, listed in Table III, where each one has two possible successor states, depending on whether or not the central site is stable. If a petal is in state ij , then by $S(ij)$ we denote the state that it is in after one time step in which the central site does not topple, and likewise by $U(ij)$ after one time step in which the central site topples.

From this we see that a petal will be in state 00 only if the central site is always stable, and consequently each site is always stable, in which case the period is 1. Similarly, petal state 33 only occurs if the central site is unstable each step, in which case each site must be unstable each step, and the period is again 1. State 03 is not a successor of any state of these states, so it will not be a periodic petal state either. Thus the set of allowed periodic petal states is $\{01, 02, 11, 12, 13, 22, 23\}$.

If the central site is stable every other time step, then the possible petal states are $12 \rightarrow 02 \rightarrow 12$, $22 \rightarrow 11 \rightarrow 22$, and $13 \rightarrow 12 \rightarrow 13$, each of which has period 2. Then the period of

TABLE III. There are ten possible petal states, and each has two possible successor states depending on whether the central site was stable (S) or unstable (U).

State	S(state)	U(state)
00	00	11
01	01	12
02	01	12
03	11	22
11	11	22
12	02	13
13	12	23
22	11	22
23	12	23
33	22	33

the entire configuration is 2. Thus if the period is larger than 2, the central site must be stable for at least two consecutive time steps, or else unstable for at least two consecutive time steps. We will label a time step S if the central site is stable in that time step, otherwise we label it U. So, if the period is larger than 2 we will see SS or UU in the time evolution. In the latter case, we can study the mirror image, which will have the same period, and for which we will see SS.

Eventually the central site must be unstable again, since otherwise the period would be 1. Therefore, we can examine three time steps labeled SSU. Examining the evolution of the central site together with the petals, we see

S	S	U	
01, 02	01	01	12
12	02	01	12
11, 22	11	11	22
13, 23	12	02	12

Whenever we have SSU, during the second and third time steps each petal contributes at most two particles to the central site, while the central site topples, so the central site must again be stable. Thus SSUU cannot occur, and we see SSUS.

There are two cases for what the central site does next. Let us first consider SSUSU.

S	S	U	S	U	
01, 02	01	01	12	02	12
12	02	01	12	02	12
11, 22	11	11	22	11	22
13, 23	12	02	12	02	12

During the last two time steps, each petal contributes exactly 2 particles to the central site, and the central site topples once. Thus after two time steps not only the petals, but also the central site is in the same state. Therefore, the period becomes 2.

Next we consider SSUSS. At this stage each petal is in state 01 or 11, so if there were yet another S, the sandpile would be periodic with period 1. So we see SSUSSU, and because SSUU is forbidden, we conclude that we see SSUSSUS.

S	S	U	S	S	U	S
01, 02	01	01	12	02	01	12
12	02	01	12	02	01	12
11, 22	11	11	22	11	11	22
13, 23	12	02	12	02	01	12

At the time of the third S, each petal is in state 12 or 22. Between the third S and the fifth S, each petal contributes exactly two particles to the central site and returns to the same state, while the central site topples once. Thus the configuration is periodic with period 3.

We conclude from the above case analysis that the activity ρ_a is always one of 0, 1/3, 1/2, 2/3, or 1. Table IV summarizes the behavior of the periodic sandpile states for different values of ρ_a .

The following theorem shows that parallel chip-firing on the flower graph exhibits four distinct phase transitions where the activity ρ_a jumps in value: for each α

TABLE IV. Behavior of the central site and petals as a function of the activity ρ_a .

		Periodic sandpile states									
Activity ρ_a	0	1/3		1/2		2/3			1		
Central site	S	S	S	U	S	U	S	U	U	U	U
petals	01	12	02	01	12	02	23	12	13	≥ 22	
	11	22	11	11	22	11	22	11	22		
	00			13	12						

$\in \{0, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}, 1\}$, there is a nonvanishing interval of initial densities λ where $\rho_a = \alpha$ asymptotically almost surely.

Theorem 13. Let ζ_c be the unique root of $\frac{5}{3} + \frac{1}{3}e^{-3\zeta} = \zeta$, and let ζ'_c be the unique positive root of $\frac{10}{3} - \frac{1}{3}e^{-3\zeta} = \zeta$. (Numerically, $\zeta_c = 1.668\ 897\ 6\dots$ and $\zeta'_c = 3.333\ 318\ 2\dots$) With probability tending to 1 as $n \rightarrow \infty$, the activity density ρ_a of parallel chip-firing on the flower graph F_n is given by

$$\rho_a = \begin{cases} 0, & \text{if } 0 \leq \lambda < \zeta_c \\ 1/3, & \text{if } \zeta_c < \lambda < 2 \\ 1/2, & \text{if } 2 < \lambda < 3 \\ 2/3, & \text{if } 3 < \lambda < \zeta'_c \\ 1 & \text{if } \zeta'_c < \lambda. \end{cases}$$

Proof. In a given petal, let X denote the difference modulo 3 of the number of particles on the two sites of the petal. Observe that X is unaffected by toppling. Let Z denote the number of petals for which $X=0$, and R denote the total number of particles, in a given initial configuration. Using Table IV, we can relate Z , R , and the activity ρ_a .

When $\rho_a=0$, we have less than $2n$ particles at the central site, at most two particles for the Z petals of type $X=0$, and exactly one particle for the other $n-Z$ petals, so $0 \leq R < 2n + 2Z + (n-Z) = 3n + Z$.

When $\rho_a=1/3$, by considering the U time step, we have $R \geq 2n + 2Z + (n-Z) = 3n + Z$. By considering the preceding S time step, we have $R < 2n + 2Z + 2(n-Z) = 4n$. When $\rho_a=1/2$, by considering the U time step, we have $R \geq 4n$, and by considering the S time step, we get $R < 2n + 4Z + 4(n-Z) = 6n$.

When $\rho_a=2/3$, by considering the second U step, we have $R \geq 2n + 4Z + 4(n-Z) = 6n$. By considering the S time step, we have $R < 2n + 4Z + 5(n-Z) = 7n - Z$.

When $\rho_a=1$, we have $R \geq 2n + 4Z + 5(n-Z) = 7n - Z$. Since for given n and Z , these intervals on the values of R are disjoint, we see that the converse statements hold as well: the values of R and Z determine the activity ρ_a . We summarize these bounds:

$$\rho_a = \begin{cases} 0 & \text{if and only if } 0 \leq R < 3n + Z \\ 1/3 & \text{if and only if } 3n + Z \leq R < 4n \\ 1/2 & \text{if and only if } 4n \leq R < 6n \\ 2/3 & \text{if and only if } 6n \leq R < 7n - Z \\ 1 & \text{if and only if } 7n - Z \leq R. \end{cases}$$




FIG. 5. Density $\rho(\lambda)$ of the final stable configuration as a function of initial density λ on the flower graph \hat{F}_n for large n . A second-order phase transition occurs at $\lambda = \zeta_c$. Beyond this transition, the density of the driven sandpile decreases with λ . (In [1], the curve in this figure was correctly graphed, but mislabeled as $(5 + e^{-\lambda})/3$.)

Everything so far holds deterministically; next we use probability to estimate R and Z . By the weak law of large numbers, $R/n \rightarrow 2\lambda$ and $Z/n \rightarrow \text{Prob}(X=0)$ in probability. Thus, to complete the proof it suffices to show

$$\text{Prob}(X=0) = \frac{1}{3}(1 + 2e^{-3\lambda}). \quad (4)$$

We can think of building the initial configuration σ_λ by starting with the empty configuration and adding particles in continuous time. Then the value of X for a single petal as particles are added is a continuous time Markov chain on the state space $\{0, \pm 1\}$ with transitions $0 \rightarrow \pm 1$ at rate 2, and $\pm 1 \rightarrow 0$ and $\pm 1 \rightarrow \pm 1$ each at rate 1. Starting in state 0, after running this chain for time λ we obtain

$$[\text{Prob}(X=0), \text{Prob}(X \neq 0)] = [1, 0] \exp \left\{ \lambda \begin{bmatrix} -2 & 2 \\ 1 & -1 \end{bmatrix} \right\}.$$

The eigenvalues of the above matrix are 0 and -3 , with corresponding left eigenvectors $v_1 = [1, 2]$ and $v_2 = [1, -1]$. Since $[1, 0] = \frac{1}{3}v_1 + \frac{2}{3}v_2$, we obtain Eq. (4).

The following theorem describes a phase transition in the driven sandpile dynamics on the flower graph analogous to Theorem 4 for the bracelet graph. We remark on one interesting difference between the two transitions: for $\lambda > \zeta_c$, the final density $\rho(\lambda)$ is increasing in λ for the bracelet, and decreasing in λ for the flower graph (Fig. 5).

For $\lambda > 0$, let σ_λ be the configuration with $\text{Poisson}(\lambda)$ particles independently on each site of \hat{F}_n . Let $\eta_\lambda = (\sigma_\lambda)^\circ$ be the stabilization of σ_λ , and let

$$\rho_n(\lambda) = \frac{1}{n-1} \sum_{x=1}^{n-1} \eta_\lambda(x)$$

be the final density.

Theorem 14. For the flower graph with n petals, in the limit $n \rightarrow \infty$ we have

- (1) The threshold density ζ_c is the unique positive root of $\zeta = \frac{5}{3} + \frac{1}{3}e^{-3\zeta}$.
- (2) The stationary density ζ_s is $5/3$.
- (3) $\rho_n(\lambda) \rightarrow \rho(\lambda)$ in probability, where




FIG. 6. The Cayley trees (Bethe lattices) of degree $d=3,4,5$.

$$\rho(\lambda) = \min\left(\lambda, \frac{5}{3} + \frac{1}{3}e^{-3\lambda}\right) = \begin{cases} \lambda, & \lambda \leq \zeta_c \\ \frac{5}{3} + \frac{1}{3}e^{-3\lambda}, & \lambda > \zeta_c. \end{cases}$$

Proof. Part 1 follows from Theorem 13.

For Part 2, we use the burning algorithm. In all recurrent configurations on \hat{F}_n , the central site has either $2n-1$ or $2n-2$ particles. All other sites have at most one particle, and in each petal (except the sink) there is at least one particle. For each petal that is not the sink, there are two possible configurations with 1 particle, and one with 2 particles. Each of these occurs with equal probability in the stationary state, so the expected number of particles in the petals is $(n-1)\left(\frac{2}{3} \cdot 1 + \frac{1}{3} \cdot 2\right) = \frac{4n}{3} + O(1)$ as $n \rightarrow \infty$. Therefore, the total density is $\zeta_s = \lim_{n \rightarrow \infty} \frac{2n+4n/3}{2n-1} + o(1) = 5/3$.

For part 3, for the driven-dissipative sandpile on \hat{F}_n , we first stabilize all the petals, then topple the center site if it is unstable, then stabilize all the petals, and so on. For each toppling of the center site, the sandpile loses $O(1)$ particles to the sink. If the center topples at least once, then each petal will be in one of the states 11, 01, or 10, after which the number of particles at the center site is $R-n-Z+O(1)$. Recall from the proof of Theorem 13 that $R/n \rightarrow 2\lambda$ and $Z/n \rightarrow \frac{1}{3}(1+2e^{-3\lambda})$ in probability. Thus if $\lambda \leq \zeta_c$, then $\frac{R-n+Z}{2n} \rightarrow \lambda - \frac{2}{3} + \frac{1}{3}e^{-3\lambda} \leq 1$ in probability, so the sandpile does not lose a macroscopic amount of sand, and $\rho_n(\lambda) \rightarrow \lambda$ in probability.

If $\lambda > \zeta_c$, then the number of particles that remain after stabilization is $2n+n+Z+O(1)$. In this case, we have $\rho_n(\lambda) = \frac{3n+Z}{2n+1} + o(1) \rightarrow \frac{5}{3} + \frac{1}{3}e^{-3\lambda}$ in probability.

VI. SANDPILES ON THE CAYLEY TREE

Dhar and Majumdar [14] studied the Abelian sandpile model on the Cayley tree (also called the Bethe lattice) with

branching factor q , which has degree $q+1$ (Fig. 6). Implicit in their formulation is that they used wired boundary conditions, i.e., where all the vertices of the tree at a certain large distance from a central vertex are glued together and become the sink. (The other common boundary condition is free boundary conditions, where all the vertices at a certain distance from the central vertex become leaves, and one of them becomes the sink. The issue of boundary conditions becomes important for trees, because in any finite subgraph, a constant fraction of vertices are on the boundary. This is in contrast to \mathbb{Z}^2 , where free and wired boundary conditions lead to the same infinite-volume limit. See [32].)

The finite regular wired tree $T_{q,n}$ is the ball of radius n in the infinite $(q+1)$ -regular tree, with all leaves collapsed to a single vertex s . In $\hat{T}_{q,n}$ the vertex s serves as the sink. Maes, Redig, and Saada [33] show that the stationary measure on recurrent sandpiles on $\hat{T}_{q,n}$ has an infinite-volume limit, which is a measure on sandpiles on the infinite tree. Denoting this measure by Prob_q , if h denotes the number of particles at a single site far from the boundary, then we have [14]

$$\text{Prob}_q[h=i] = \frac{1}{(q^2-1)q^q} \sum_{m=0}^i \binom{q+1}{m} (q-1)^{q+1-m}.$$

From this formula we see that the stationary density is

$$\zeta_s = \mathbb{E}_q[h] = \frac{q+1}{2}.$$

For 3-regular, 4-regular, and 5-regular trees, these values are summarized in Table V.

TABLE V. Stationary distribution of the sandpile height at a vertex of the Bethe lattice (Cayley tree), from Dhar and Majumdar’s formula [14].

Tree		Distribution of height h of sand					
q	Degree	$\mathbb{E}_q[h]$	$\text{Prob}_q[h=0]$	$\text{Prob}_q[h=1]$	$\text{Prob}_q[h=2]$	$\text{Prob}_q[h=3]$	$\text{Prob}_q[h=4]$
2	3	3/2	1/12	4/12	7/12		
3	4	2	2/27	2/9	1/3	10/27	
4	5	5/2	81/1280	27/160	153/640	21/80	341/1280

TABLE VI. Data for the fixed-energy sandpile on a pseudorandom 3-regular graph on n nodes. Each estimate of $\mathbb{E}[h]$ has standard deviation less than 7×10^{-8} , and each estimate of the marginals $\text{Prob}[h=i]$ has standard deviation less than 3×10^{-7} . The data for $\mathbb{E}[h]$ appears to fit $3/2 + \text{const}/\sqrt{n}$ very well, and extrapolating to $n \rightarrow \infty$ it appears that $\mathbb{E}[h] \rightarrow 1.500\,000$ to six decimal places. However, apparently $\text{Prob}[h=0] \rightarrow 0.083\,331 < 1/12$.

n	No. samples	$\mathbb{E}[h]$	Distribution of height h of sand			(No. topplings)/($n \log^{1/2} n$)
			$\text{Prob}[h=0]$	$\text{Prob}[h=1]$	$\text{Prob}[h=2]$	
1048576	2097152	1.5004315	0.0833326	0.332903	0.583764	1.263145
2097152	1048576	1.5003054	0.0833321	0.333031	0.583637	1.258046
4194304	524288	1.5002161	0.0833314	0.333121	0.583548	1.253092
8388608	262144	1.5001528	0.0833311	0.333185	0.583484	1.247642
16777216	131072	1.5001081	0.0833311	0.333230	0.583439	1.242359
33554432	65536	1.5000765	0.0833307	0.333262	0.583407	1.237317
67108864	32768	1.5000540	0.0833307	0.333285	0.583385	1.232398
134217728	16384	1.5000382	0.0833307	0.333300	0.583369	1.227548
268435456	8192	1.5000269	0.0833308	0.333311	0.583358	1.222371
536870912	4096	1.5000191	0.0833308	0.333319	0.583350	1.214431
1073741824	2048	1.5000136	0.0833307	0.333325	0.583344	1.212751
∞ (stationary)		1.5	0.0833333	0.333333	0.583333	

Large-scale simulations on $T_{q,n}$ are rather impractical because the vast majority of vertices are near the boundary. Consequently, each simulation run produces only a small amount of usable data from vertices near the center.

To experimentally measure ζ_c for the Cayley trees, we generated large random regular graphs $G_{q,n}$, and used these as finite approximations of the infinite Cayley tree. We used the following procedure to generate random connected bipartite multigraphs of degree $q+1$ on n vertices (n even). Let M_0 be the set of edges $(i, i+1)$ for $i=1, 3, 5, \dots, n-1$. Then take the union of M_0 with q additional i.i.d. perfect matchings M_1, \dots, M_q between odd and even vertices. Each M_j is chosen uniformly among all odd-even perfect matchings whose union with M_0 is an n -cycle.

Most vertices of G_n will not be contained in any cycle smaller than $\log_q n + O(1)$ (see e.g., [34]), so these graphs are locally treelike. For this reason, we believe that as $n \rightarrow \infty$ the threshold density $\Lambda_c(G_n)$ will be concentrated at the threshold density of the infinite tree.

Since the choice of multigraph affects the estimate of ζ_c , we generated a new independent random multigraph for each trial. The results for random regular graphs of degree 3, 4 and 5 are summarized in Tables VI–VIII. We find that for the 5-regular tree, the threshold density is about 2.511 rather than 2.5, for the 4-regular tree the threshold density is very close to but decidedly larger than 2, while for the 3-regular tree the threshold density is extremely close to 1.5, with a discrepancy that we were unable to measure. However, for

TABLE VII. Data for the fixed-energy sandpile on a pseudorandom 4-regular graph on n nodes. Each estimate of $\mathbb{E}[h]$ and of the marginals $\text{Prob}[h=i]$ has standard deviation less than 3×10^{-7} .

n	No. samples	$\mathbb{E}[h]$	Distribution of height h of sand				(No. topplings)/($n \log^{1/2} n$)
			$\text{Prob}[h=0]$	$\text{Prob}[h=1]$	$\text{Prob}[h=2]$	$\text{Prob}[h=3]$	
1048576	2097152	2.001109	0.073884	0.221887	0.333466	0.370763	0.623322
2097152	1048576	2.000853	0.073881	0.221978	0.333547	0.370593	0.618848
4194304	524288	2.000688	0.073880	0.222037	0.333599	0.370484	0.620894
8388608	262144	2.000584	0.073878	0.222075	0.333631	0.370416	0.631324
16777216	131072	2.000518	0.073877	0.222100	0.333651	0.370372	0.649328
33554432	65536	2.000477	0.073877	0.222114	0.333664	0.370345	0.670838
67108864	32768	2.000451	0.073877	0.222123	0.333673	0.370328	0.691040
134217728	16384	2.000434	0.073876	0.222130	0.333678	0.370316	0.699706
268435456	8192	2.000424	0.073876	0.222134	0.333681	0.370310	0.695065
536870912	4096	2.000417	0.073876	0.222136	0.333683	0.370305	0.684507
1073741824	2048	2.000413	0.073876	0.222138	0.333684	0.370303	0.673061
∞ (stationary)		2	0.074074	0.222222	0.333333	0.370370	

TABLE VIII. Data for the fixed-energy sandpile on a pseudorandom 5-regular graph on n nodes. Each estimate of $\mathbb{E}[h]$ and of the marginals $\text{Prob}[h=i]$ has standard deviation less than 2×10^{-6} .

n	No. samples	$\mathbb{E}[h]$	Distribution of height h of sand					(No. topplings)/ n
			$\text{Prob}[h=0]$	$\text{Prob}[h=1]$	$\text{Prob}[h=2]$	$\text{Prob}[h=3]$	$\text{Prob}[h=4]$	
1048576	1048576	2.512106	0.062271	0.166547	0.237230	0.264711	0.269242	1.666086
2097152	524288	2.511947	0.062269	0.166579	0.237256	0.264727	0.269169	1.666244
4194304	262144	2.511847	0.062268	0.166599	0.237272	0.264737	0.269123	1.666404
8388608	131072	2.511781	0.062267	0.166613	0.237283	0.264743	0.269093	1.666589
16777216	65536	2.511743	0.062267	0.166621	0.237289	0.264748	0.269075	1.667322
33554432	65536	2.511716	0.062267	0.166627	0.237293	0.264750	0.269063	1.668196
67108864	32768	2.511700	0.062267	0.166630	0.237296	0.264752	0.269056	1.669392
134217728	16384	2.511689	0.062267	0.166632	0.237297	0.264755	0.269050	1.671613
268435456	8192	2.511683	0.062266	0.166634	0.237299	0.264753	0.269048	1.675479
536870912	4096	2.511680	0.062267	0.166633	0.237300	0.264755	0.269045	1.677092
1073741824	2048	2.511677	0.062266	0.166634	0.237300	0.264755	0.269044	1.688093
∞ (stationary)		2.5	0.063281	0.168750	0.239063	0.262500	0.266406	

the 3-regular tree there is a measurable discrepancy (about 2×10^{-6}) in the probability that a site has no particles.

VII. SANDPILES ON THE LADDER GRAPH

The examples in previous sections suggest that the density conjecture can fail for (at least) two distinct reasons: local toppling invariants, and boundary effects. A *toppling invariant* for a graph G is a function f defined on sandpile configurations on G which is unchanged by performing topplings; that is

$$f(\sigma) = f(\sigma + \Delta_x)$$

for any sandpile σ and any column vector Δ_x of the Laplacian of G . Examples we have seen are

$$f(\sigma) = \sigma(x) \bmod 2$$

where x is any vertex of the bracelet graph B_n ; and

$$f(\sigma) = \sigma(x_1) - \sigma(x_2) \bmod 3,$$

where x_1, x_2 are the two vertices comprising any petal on the flower graph F_n . Both of these toppling invariants are *local* in the sense that they depend only on a bounded number of vertices as $n \rightarrow \infty$.

The Cayley tree has no local toppling invariants, but the large number of sinks, comparable to the total number of vertices, produce a large boundary effect. The density conjecture fails even more dramatically on the complete graph (Theorem 6). One might guess that this is due to the high degree of interconnectedness, which causes boundary effects from the sink to persist as $n \rightarrow \infty$. A good candidate for a graph G satisfying the density conjecture, then, should have

- (i) no local toppling invariants,
- (ii) most vertices far from the sink.

The best candidate graphs G should be essentially one-dimensional, so that the sink is well insulated from the bulk of the graph, keeping boundary effects to a minimum. In-

deed, the only graph known to satisfy the density conjecture is the infinite path \mathbb{Z} .

Járai and Lyons [15] study sandpiles on graphs of the form $G \times P_n$, where G is a finite connected graph and P_n is the path of length n , with the end points serving as sinks. The simplest such graphs that are not paths are obtained when $G = P_1$ has two vertices and one edge (Fig. 7). These graphs are a good candidate for $\zeta_c = \zeta_s$, for the reasons described above. Nevertheless, we find that while ζ_c and ζ_s are very close, they appear to be different.

First we calculate ζ_s . Járai and Lyons [[15], Sec. 5] define recurrent configurations as Markov chains on the state space

$$X = \{(3,3), (3,2), (2,3), (3,1), (1,3), \overline{(3,2)}, \overline{(2,3)}\}$$

describing the possible transitions from one rung of the ladder to the next. States (i,j) and $\overline{(i,j)}$ both represent rungs whose left vertex has $i-1$ particles and whose right vertex has $j-1$ particles. The distinction between states $(3,2)$ and $\overline{(3,2)}$ lies only in which transitions are allowed. The adjacency matrix describing the allowable transitions is given by

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

Its largest eigenvalue is $2 + \sqrt{3}$, and the corresponding left and right eigenvectors are

$$u = (1 + \sqrt{3}, 1 + \sqrt{3}, 1 + \sqrt{3}, 1, 1, 1, 1),$$



FIG. 7. The ladder graph.

TABLE IX. Data for the fixed-energy sandpile on $2 \times n$ ladder graphs. Each estimate of $\mathbb{E}[h]$ and of the marginals $\text{Prob}[h=i]$ has a standard deviation smaller than 10^{-5} . To four decimal places, the threshold density ζ_c equals 1.6082, which exceeds the stationary density $\zeta_s = 7/4 - \sqrt{3}/12 = 1.6057$. The total number of topplings appears to scale as $n^{5/2}$.

n	No. samples	$\mathbb{E}[h]$	Distribution of height h of sand			(No. topplings)/ $n^{5/2}$
			$\text{Prob}[h=0]$	$\text{Prob}[h=1]$	$\text{Prob}[h=2]$	
256	4194304	1.60567	0.07695	0.24043	0.68262	0.094773
512	2097152	1.60693	0.07656	0.23996	0.68349	0.095366
1024	1048576	1.60757	0.07636	0.23970	0.68393	0.095864
2048	524288	1.60788	0.07626	0.23960	0.68414	0.096316
4096	262144	1.60805	0.07621	0.23952	0.68426	0.096545
8192	131072	1.60814	0.07618	0.23950	0.68432	0.096753
16384	65536	1.60816	0.07618	0.23949	0.68434	0.097113
32768	32768	1.60818	0.07617	0.23948	0.68435	0.096944
65536	16384	1.60820	0.07616	0.23948	0.68436	0.097342
131072	8192	1.60820	0.07617	0.23946	0.68437	0.097648
262144	4096	1.60821	0.07615	0.23949	0.68436	0.096158
∞ (stationary)		1.60566	0.07735	0.23964	0.68301	

$$v = (3 + \sqrt{3}, 1 + \sqrt{3}, 1 + \sqrt{3}, 1 + \sqrt{3}, 1 + \sqrt{3}, 1, 1)^T.$$

By the Parry formula [35], the stationary probabilities are given by $p(i) = u_i v_i / Z$, where Z is a normalizing constant. So

$$p(3,3) = (1 + \sqrt{3})(3 + \sqrt{3})/Z,$$

$$p(2,3) = p(3,2) = (1 + \sqrt{3})^2/Z,$$

$$p(1,3) = p(3,1) = (1 + \sqrt{3})/Z,$$

$$p(\overline{2},3) = p(3,\overline{2}) = 1/Z,$$

where

$$Z = (1 + \sqrt{3})(3 + \sqrt{3}) + 2(1 + \sqrt{3})^2 + 2(1 + \sqrt{3}) + 2.$$

Thus we find that for the ladder graph in stationarity, the number h of particles at a site satisfies

$$\text{Prob}[h = 0] = -\frac{1}{2} + \frac{\sqrt{3}}{3} = 0.0773503 \dots,$$

$$\text{Prob}[h = 1] = \frac{5}{4} - \frac{7\sqrt{3}}{12} = 0.2396370 \dots,$$

$$\text{Prob}[h = 2] = \frac{1}{4} + \frac{\sqrt{3}}{4} = 0.6830127 \dots,$$

$$\zeta_s = \mathbb{E}[h] = \frac{7}{4} - \frac{\sqrt{3}}{12} = 1.60566243 \dots$$

In contrast, the threshold density for ladders appears to be about 1.6082. Table IX summarizes simulation data on finite $2 \times n$ ladders.

VIII. CONCLUSIONS

We have rigorously demonstrated that Conjecture 3 does not hold for the Abelian sandpile model, so that the conclu-

sions of [6], ‘‘FES are shown to exhibit an absorbing state transition with critical properties coinciding with those of the corresponding sandpile model,’’ deserve to be re-evaluated.

In some recent papers such as [36], the DMVZ paradigm is explicitly restricted to stochastic models. In other recent papers [37,38] it is claimed to apply both to stochastic and deterministic sandpiles, although these papers focus on stochastic sandpiles, for the reason that deterministic sandpiles are said to belong to a different universality class. While our results refute the density conjecture for deterministic sandpiles, the validity of the density conjecture for stochastic sandpiles remains an intriguing open question.

An interesting possibility for further research is to examine initial conditions other than a Poisson(ζ) number of particles independently at each site. As Grassberger and Manna observed [8], the value of the FES threshold density depends on the choice of initial condition. One might consider a more general version of FES, namely adding independent Poisson($(\zeta - \zeta_0)$) numbers of particles to a ‘‘background’’ configuration τ of density ζ_0 already present. For example, taking τ to be the deterministic configuration on \mathbb{Z}^d of $2d - 2$ particles everywhere, by [[25], Prop 1.4] we obtain a threshold density of $\zeta_c = 2d - 2$. Many interesting questions present themselves: for instance, for which background does ζ_c take the smallest value, and for which backgrounds do we obtain $\zeta_c = \zeta_s$? It would also be interesting to replicate the phase transition for driven sandpiles (see Theorems 4 and 14) for different backgrounds.

ACKNOWLEDGMENT

L. L. was partially supported by the National Science Foundation.

- [1] A. Fey, L. Levine, and D. B. Wilson, *Phys. Rev. Lett.* **104**, 145703 (2010).
- [2] R. Dickman, A. Vespignani, and S. Zapperi, *Phys. Rev. E* **57**, 5095 (1998).
- [3] A. Vespignani, R. Dickman, M. A. Muñoz, and S. Zapperi, *Phys. Rev. Lett.* **81**, 5676 (1998).
- [4] R. Dickman, M. Muñoz, A. Vespignani, and S. Zapperi, *Braz. J. Phys.* **30**, 27 (2000).
- [5] A. Vespignani, R. Dickman, M. A. Muñoz, and S. Zapperi, *Phys. Rev. E* **62**, 4564 (2000).
- [6] M. A. Muñoz, R. Dickman, R. Pastor-Satorras, A. Vespignani, and S. Zapperi, in *Proc. 6th Granada Sem. on Comput. Physics* (American Institute of Physics, Melville, NY, 2001).
- [7] P. Bak, C. Tang, and K. Wiesenfeld, *Phys. Rev. Lett.* **59**, 381 (1987).
- [8] P. Grassberger and S. S. Manna, *J. Phys. France* **51**, 1077 (1990).
- [9] M. De Menech, A. L. Stella, and C. Tebaldi, *Phys. Rev. E* **58**, R2677 (1998).
- [10] F. Bagnoli, F. Cecconi, A. Flammini, and A. Vespignani, *EPL* **63**, 512 (2003).
- [11] G. Pruessner and O. Peters, *Phys. Rev. E* **73**, 025106R (2006).
- [12] O. Peters and G. Pruessner, e-print arXiv:0912.2305 (unpublished).
- [13] M. Jeng, G. Piroux, and P. Ruelle, *J. Stat. Mech.: Theory Exp.* (2006) P10015.
- [14] D. Dhar and S. N. Majumdar, *J. Phys. A* **23**, 4333 (1990).
- [15] A. A. Járai and R. Lyons, *Markov Processes Relat. Fields* **13**, 493 (2007).
- [16] D. Dhar, *Phys. Rev. Lett.* **64**, 1613 (1990).
- [17] S. R. Athreya and A. A. Járai, *Commun. Math. Phys.* **249**, 197 (2004).
- [18] S. N. Majumdar and D. Dhar, *J. Phys. A* **24**, L357 (1991).
- [19] V. B. Priezzhev, *J. Stat. Phys.* **74**, 955 (1994).
- [20] L. Levine, *Ergod. Theory Dyn. Syst.* (to be published).
- [21] A. Björner, L. Lovász, and P. Shor, *Eur. J. Comb.* **12**, 283 (1991).
- [22] J. Bitar and E. Goles, *Theor. Comput. Sci.* **92**, 291 (1992).
- [23] G. Tardos, *SIAM J. Discrete Math.* **1**, 397 (1988).
- [24] A. Fey-den Boer, R. Meester, and F. Redig, *Ann. Probab.* **37**, 654 (2009).
- [25] A. Fey, L. Levine, and Y. Peres, *J. Stat. Phys.* **138**, 143 (2010).
- [26] F. Redig, *Les Houches lecture notes* (2005).
- [27] C. Merino López, *Ann. Comb.* **1**, 253 (1997).
- [28] E. M. Wright, *J. Graph Theory* **1**, 317 (1977).
- [29] L. Dall'Asta, *Phys. Rev. Lett.* **96**, 058003 (2006).
- [30] M. Casartelli, L. Dall'Asta, A. Vezzani, and P. Vivo, *Eur. Phys. J. B* **52**, 91 (2006).
- [31] E. Prisner, *Complex Syst.* **8**, 367 (1994).
- [32] R. Lyons and Y. Peres, *Probability on Trees and Networks* (Cambridge University Press, Cambridge, England, 2010), in preparation. Available at <http://mypage.iu.edu/rdlyons/prbtree/prbtree.html>
- [33] C. Maes, F. Redig, and E. Saada, *Ann. Probab.* **30**, 2081 (2002).
- [34] B. Bollobás, *Extremal Graph Theory with Emphasis on Probabilistic Methods*, CBMS Regional Conference Series in Mathematics Vol. 62 (American Mathematical Society, Providence, Rhode Island, 1986).
- [35] W. Parry, *Trans. Am. Math. Soc.* **112**, 55 (1964).
- [36] J. A. Bonachela and M. A. Muñoz, *Phys. Rev. E* **78**, 041102 (2008).
- [37] R. Vidigal and R. Dickman, *J. Stat. Phys.* **118**, 1 (2005).
- [38] S. D. da Cunha, R. R. Vidigal, L. R. da Silva, and R. Dickman, *Eur. Phys. J. B* **72**, 441 (2009).