

Bottom-up model of self-organized criticality on networksPierre-André Noël,^{1,2,*} Charles D. Brummitt,^{1,3} and Raissa M. D'Souza^{1,2,4,5}¹*Complexity Sciences Center, University of California, Davis, California 95616, USA*²*Department of Computer Science, University of California, Davis, California 95616, USA*³*Department of Mathematics, University of California, Davis, California 95616, USA*⁴*Department of Mechanical and Aerospace Engineering, University of California, Davis, California 95616, USA*⁵*Santa Fe Institute, Santa Fe, New Mexico 87501, USA*

(Received 20 October 2013; published 15 January 2014)

The Bak-Tang-Wiesenfeld (BTW) sandpile process is an archetypal, stylized model of complex systems with a critical point as an attractor of their dynamics. This phenomenon, called self-organized criticality, appears to occur ubiquitously in both nature and technology. Initially introduced on the two-dimensional lattice, the BTW process has been studied on network structures with great analytical successes in the estimation of macroscopic quantities, such as the exponents of asymptotically power-law distributions. In this article, we take a microscopic perspective and study the inner workings of the process through both numerical and rigorous analysis. Our simulations reveal fundamental flaws in the assumptions of past phenomenological models, the same models that allowed accurate macroscopic predictions; we mathematically justify why universality may explain these past successes. Next, starting from scratch, we obtain microscopic understanding that enables mechanistic models; such models can, for example, distinguish a cascade's area from its size. In the special case of a 3-regular network, we use self-consistency arguments to obtain a zero-parameter mechanistic (bottom-up) approximation that reproduces nontrivial correlations observed in simulations and that allows the study of the BTW process on networks in regimes otherwise prohibitively costly to investigate. We then generalize some of these results to configuration model networks and explain how one could continue the generalization. The numerous tools and methods presented herein are known to enable studying the effects of controlling the BTW process and other self-organizing systems. More broadly, our use of multitype branching processes to capture information bouncing back and forth in a network could inspire analogous models of systems in which consequences spread in a bidirectional fashion.

DOI: [10.1103/PhysRevE.89.012807](https://doi.org/10.1103/PhysRevE.89.012807)

PACS number(s): 89.75.Hc, 02.30.Yy, 05.65.+b, 45.70.Ht

I. INTRODUCTION

Many complex systems affecting modern life, from infrastructure systems such as power grids to the natural catastrophes that threaten them, appear to be poised near criticality. For instance, power-law distributions seem to characterize the sizes of electrical blackouts [1], financial fluctuations [2], neuronal avalanches [3–5], earthquakes [6], landslides [7], overspill in water reservoirs [8], forest fires [9,10], and solar flares [11,12]. Since its introduction in 1987 [13,14], the Bak-Tang-Wiesenfeld (BTW) sandpile has served as a useful paradigm for the self-organizing dynamics that may drive these systems toward critical points [also called self-organized criticality (SOC)]. The recent prevalence of cascading failures and overloads in networked infrastructures [1,15,16] is one motivation for studying the sandpile model on random graphs [17–25]. Here we attain deeper understanding of the BTW sandpile model on networks, which provides important lessons and techniques for studies of self-organized critical processes in natural and engineered systems in general.

We begin in Sec. II with a brief background on complex networks and on the BTW sandpile process on networks. We then provide a macroscopic understanding of how these systems self-organize, including a few past results of interest.

Though past work has significantly advanced our macroscopic understanding of the sandpile process on networks, we demonstrate in Sec. III some fundamental gaps between this macroscopic view and the microscopic reality. Sections IV–VI

seek to reconcile these macroscopic and microscopic perspectives. Before doing so, we emphasize in Sec. III B the need for our models of the sandpile process to be self-consistent, a feature that plays a central role in the rest of the paper.

In Sec. IV, we start from scratch and develop a microscopic analytical understanding of the BTW sandpile process on networks. We prove rigorous results concerning the inner workings of a cascade, which are summarized in Theorems 1–3 and proved in the Appendixes.

Section V combines these rigorous results with self-consistency arguments to obtain a zero-parameter model for random 3-regular networks. This model allows the study of the BTW sandpile process in a regime prohibitively costly to investigate in simulations. As reported in [26], this approach is an analytical model that can separately calculate cascade size and cascade area.

Section VI generalizes to configuration model networks some of the results for random 3-regular graphs in Sec. V, including the independent calculations of cascade size and cascade area. The same method as in Sec. V could be used to obtain a zero-parameter model. However, doing so appears more amenable to a case-by-case study (e.g., a mix of nodes of degrees 3 and 4).

Section VII discusses the impact of our work in the general context of SOC on networks. We stress that self-consistency may be more important to the success of an analytical approximation—as it captures the self-organization mechanism—than direct consistency with the original system. We also argue that the success of our mechanistic self-consistent modeling approach is an important proof of concept

*noel.pierre.andre@gmail.com

and that this method's applicability should extend beyond the BTW model.

Proofs and additional justifications are presented in Appendixes.

II. STUDIED PROCESS

The BTW sandpile model considers a large collection of nodes that shed load to their neighbors whenever they reach their capacity. We drop a discrete unit of load (called a grain of sand) randomly onto the system after each cascade finishes. A single grain can cause a large cascade (avalanche) of sand to move around the system as the system restabilizes. Such cascades typically occur in sizes distributed according to a power law because the system self-organizes to a critical point, as some have argued occurs to some extent in electrical grids [1], financial markets [27], neuronal avalanches [3,5], and some natural catastrophes [6,7,9–12]. Inspired by the excess load and stress that can cascade among critical infrastructure systems—such as blackouts in power grids, patients in overwhelmed hospitals, and excess travelers in transportation—here we study these cascades occurring on a network.

This section presents background material needed for the rest of the article. Section II A covers important concepts concerning complex networks and Sec. II B quickly describes the BTW process on networks. Finally, Sec. II C explains how the BTW process self-organizes and it presents the resulting observables for various network structures.

A. Networks with quenched or annealed structure

Networks (graphs) consist of *nodes* (vertices), representing the elements of a system, connected by *links* (edges), representing interactions among those elements. Two nodes are *neighbors* (adjacent) if they are joined by a link, and the *degree* of a node is its number of neighbors. A network's *degree distribution* is the sequence $\{p_k : k \geq 0\}$ such that p_k is the fraction of nodes with degree k . The *configuration model* samples random graphs from all graphs on N nodes with a specified degree distribution. Our numerical implementations of the configuration model uses the following algorithm [28,29]. (i) Create N isolated nodes. (ii) Assign to each node a number of half links sampled from the degree distribution $\{p_k : k \geq 0\}$. The total number of half links should be even (otherwise, discard one node's degree and resample until the total number of half links is even). (iii) Select two half links uniformly at random and pair them to form a link. Repeat until no half links remain. Although this process may create parallel edges and self-loops, they occur so rarely that they can be neglected.

This article studies networks with *quenched structure*: Once created, the structure of the network does not change. Thus neighbors remain neighbors throughout the process. At the other extreme, the network could have *annealed structure*, in which the structure of the network changes at a rate arbitrarily faster than the considered process. In this case, the network “forgets” the identities of past neighbors and at each time step the network is an independent realization from the ensemble of networks. This distinction between quenched and annealed is important: Quenched structure allows intricate correlations to appear among the internal states of nearby nodes, whereas such correlations among nodes' states are impossible in a network with annealed structure.

B. The BTW sandpile model on networks

Originally introduced on the plane, the BTW sandpile process [13,14] can be generalized to networks in a few natural ways that differ only in specifics [17–25,30]. Throughout this paper, we consider the following natural formulation [21,24,25].

Each node holds grains of sand. We call a node *i-sand* if it holds i grains of sand. The *capacity* of a node is the maximal amount of sand that it can hold. In this article, we set the capacity of every node to one less than its degree [19,21–25]. Hence, a $(k - 1)$ -sand node of degree k is *at capacity*, which means that it holds as much sand as it can withstand. Adding one or more grains to this node would bring it *over capacity*. A node brought over capacity *topples*, which means that it sheds its load by sending one grain to each of its neighbors.

The BTW sandpile process consists of a sequence of *cascades* (avalanches), defined as follows. Drop a grain of sand on a node chosen uniformly at random, called the *root* of the cascade. If this addition does not bring the root over capacity, then that cascade is finished. However, if the root is over capacity, then the root topples and sheds one grain to each of its neighbors. Any node that exceeds its capacity topples in the same way, until all nodes hold a number of grains less than or equal to its capacity. The *size* of a cascade is the number of toppling events; the *area* of a cascade is the number of nodes that topple. Subsequently, we begin a new cascade by dropping a grain on a root node chosen uniformly at random. Details of the simulation algorithm are provided in [31].

Some mechanism is required to avoid inundating the system with sand. In this article, we choose *annealed dissipation*: Whenever a grain of sand moves from one node to another, independently and with probability ϵ , this grain disappears from the system. Another viable choice would be *quenched dissipation*: A fraction ϵ of nodes are *sinks* and grains of sand sent to sinks disappear rather than pile up. (Note that the difference between these two alternatives is conceptually much less important than the one between networks with quenched structure and with annealed structure.)

C. Macroscopic understanding of the BTW sandpile process on networks

For a sufficiently large network with quenched structure ($N \rightarrow \infty$) and for a sufficiently small probability of annealed dissipation ($\epsilon \rightarrow 0$), the BTW sandpile process on networks self-organizes to a stationary state in which the cascade sizes follow a power-law distribution of exponent τ (i.e., the probability for a cascade to be of size s is proportional to $s^{-\tau}$ for large s). Figure 1 demonstrates this well-known result using Monte Carlo simulations. Historically, only cascades in which no sand dissipates were analyzed [21–23]; the size distribution of these so-called bulk cascades are indicated by plus symbols in Fig. 1. In this article, we focus on the distribution for all cascades, illustrated by circles in Fig. 1. Note that the size distribution of bulk cascades can be approximated using that of all cascades by applying the exponentially decaying weight $\{(1 - \epsilon)^{\langle k \rangle s} : s \geq 0\}$, where s is the size of the cascade and $\langle k \rangle$ is the average node degree. This estimation of the chance that no grains dissipate uses the approximation that $\langle k \rangle s$ grains are shed during the cascade. Note that the cascades

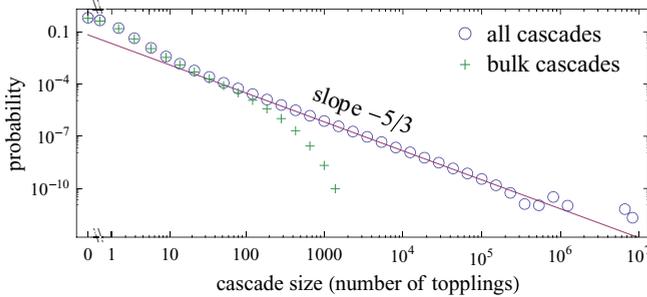


FIG. 1. (Color online) Cascade size distribution for a scale-free graph generated by the static model [32] with $N = 10^7$ nodes, power-law degree exponent $\gamma = 2.5$, average degree 6.3, and dissipation rate $\epsilon = 10^{-3}$. We use only the largest component (containing 9 503 592 nodes), and we collect statistics on the last 10^8 cascades out of 1.2×10^8 cascades. We logarithmically bin the data for all cascades (open circles) and for the bulk cascades in which no sand dissipates (pluses). Past work based on the $1/k$ assumption [21] successfully predicts that the cascade size distribution should follow a power law with exponent $\tau = \gamma/(\gamma - 1) = 5/3$ (line).

of size greater than 2×10^6 are an artifact of the combined effect of the finite size of the system and the heavy-tailed degree distribution.¹ The power-law behavior of the cascade size distribution indicates that the system is in a critical state: When a node topples, it directly causes on average $R_0 \approx 1$ other nodes to topple. An intuitive explanation reveals why the branching factor R_0 should approach unity. The crucial observation is that a cascade of size s destroys on average $\kappa \epsilon s$ grains of sand, where $\kappa \approx \langle k \rangle$ is the expected number of grains shed by a node that topples. On the one hand, if $R_0 < 1$, then the distribution of cascade size s falls exponentially, so there exists an ϵ sufficiently small so that $\kappa \epsilon \langle s \rangle < 1$. Because one grain of sand is added at each cascade and because each cascade destroys less than one grain on average, the amount of sand slowly builds up in the network, hence increasing R_0 toward unity. On the other hand, if $R_0 > 1$, then giant cascades with $s \propto N$ become possible. For fixed ϵ , there exists an N large enough so that $\kappa \epsilon \langle s \rangle > 1$, so more sand is destroyed than added and R_0 thus decreases toward unity. Hence, the critical value $R_0 = 1$ is an attractor of the dynamics (SOC).

The value of the exponent τ of the power-law cascade size distribution depends on the network structure. In the case of Fig. 1, a scale-free random graph with degree exponent $\gamma = 2.5$ (i.e., the probability that a node selected uniformly at random has degree k is proportional to $k^{-\gamma}$ over many orders of

¹In Fig. 1, the presence of a gap of width 5×10^6 in the cascade size distribution followed by a cluster of larger cascades appears to imply the existence of giant cascades. However, this gap is a combined effect of the finite size of the system (here $N = 10^7$ nodes) and the heavy-tailed degree distribution (here scale-free with exponent $\gamma = 2.5$). For fixed ϵ , the expected fractional size of giant cascades and the probability for such giant cascades both appear to converge to zero as $N \rightarrow \infty$, so these appear to not be real giants. For a network with a light-tailed degree distribution (e.g., a random 3-regular graph, in which every node has degree 3), no such giant cascades are observed, even for finite N .

magnitude) was generated using the static model [32], which results in an exponent $\tau = 5/3$ for the power-law cascade size distribution. Using the assumption that at equilibrium a degree- k node receiving a grain of sand has probability $1/k$ to topple—which we hereafter refer to as the $1/k$ assumption—past work [21] predicts that for $2 < \gamma < 3$ one should obtain a cascade size distribution of exponent $\tau = \gamma/(\gamma - 1)$. Empirical observations confirm this prediction. For $\gamma > 3$ and for other networks with light-tailed degree distributions, the mean-field value $\tau = 3/2$ is observed [19]. For instance, for a random 3-regular graph, the degree distribution is light tailed, so the slope $-\tau = -3/2$; see Fig. 4.

For networks with annealed structure, it has been shown that a node of degree k selected uniformly at random is i -sand with probability $1/k$ for $0 \leq i < k$ [33], which implies that the $1/k$ assumption holds in this case. However, the proof hinges on the assumption of annealed structure, so it does not apply for networks with quenched structure. Past works report that the $1/k$ assumption approximately holds in simulations on networks with quenched structure [21–25]. Moreover, the $1/k$ assumption enables the aforementioned successful analytical prediction that $\tau = \gamma/(\gamma - 1)$ for $2 < \gamma < 3$ [21–23]. Notwithstanding this success, we show next in Sec. III that the full picture is more intricate.

III. MOTIVATION FOR FURTHER ANALYSIS

Past works on the BTW sandpile process on networks focus on the asymptotic behavior at the critical state. Within this paradigm, a great analytical success is the prediction of the power-law exponent of the cascade size distribution for random networks with light-tailed [19] and scale-free [21–23] degree distributions. However, there are still fundamental gaps in our understanding of the BTW process on networks.

Section III A exposes some of these gaps, which motivate the in-depth analysis presented in the rest of this paper. Section III B justifies why our approach for this in-depth analysis places so much emphasis on self-consistency.

A. Caveats in the $1/k$ assumption: Going beyond universality

Although we confirm that the specific statement “at equilibrium a degree- k node receiving a grain of sand has probability $1/k$ to topple” is approximately true in simulations, we identify three subtle points that are not appropriately reflected by this $1/k$ assumption.

First, denoting by p_k the probability that a node selected uniformly at random has degree k , the assumption that a degree- k node is i -sand with probability $1/k$ for $0 \leq i < k$ predicts for the scale-free graph with degree exponent $\gamma = 2.5$ an average amount of sand per node of

$$\sum_k \sum_{i=0}^{k-1} i \mathbb{P}(\text{degree } k, i\text{-sand}) = \sum_k p_k \sum_{i=0}^{k-1} i \frac{1}{k} \approx 2.66. \quad (1)$$

However, the observed value in our simulations is 21% higher (i.e., 3.22 grains per node). For a random 3-regular network, the difference is 50% (namely, 1.50 grains per nodes instead of the $\frac{0}{3} + \frac{1}{3} + \frac{2}{3} = 1$ grain per node predicted by the $1/k$ assumption). In all the networks that we studied, instead of being flatly distributed among the possible values $0 \leq i < k$,

the probability that a degree- k node selected uniformly at random is i -sand is skewed toward larger values of i . In particular, we observed that the probability for this node to be at capacity [i.e., $(k-1)$ -sand] is typically much larger than $1/k$.

This higher probability for a node to be at capacity is somehow counterbalanced by a second observation: Nodes are unlikely to topple more than once during the same cascade, especially nodes that are many hops away from the root of this cascade. The reason is intuitive: A node that recently toppled had its amount of sand reset to zero at the moment of toppling, so it is unlikely to topple again during the same cascade. Hence, when a nonroot node u of degree k topples, one of the k grains shed by u is unlikely to cause further topplings because that grain is sent to the node v that caused u to topple in the first place.

A third subtlety is that intricate correlations exist between the amounts of sand on neighboring nodes. For simplicity, we consider the example of a random 3-regular graph. We define ψ_i to be the probability that a uniformly random node is i -sand, and we define ϕ_{ij} to be the probability that a uniformly random neighbor of a uniformly random i -sand node is j -sand. Numerical simulations for a network of $N = 10^7$ nodes using a dissipation of $\epsilon = 10^{-3}$ provide the values

$$(\psi_0 \quad \psi_1 \quad \psi_2) \approx (0.083 \ 97 \quad 0.334 \ 03 \quad 0.581 \ 99), \quad (2a)$$

$$\begin{pmatrix} \phi_{00} & \phi_{01} & \phi_{02} \\ \phi_{10} & \phi_{11} & \phi_{12} \\ \phi_{20} & \phi_{21} & \phi_{22} \end{pmatrix} \approx \begin{pmatrix} 0.000 \ 50 & 0.251 \ 48 & 0.748 \ 02 \\ 0.063 \ 21 & 0.313 \ 49 & 0.623 \ 30 \\ 0.107 \ 95 & 0.357 \ 73 & 0.534 \ 32 \end{pmatrix}. \quad (2b)$$

If the amount of sand on neighboring nodes were uncorrelated, then we would expect ϕ_{ij} to be identically equal to ψ_j , independently of the value of i on which we condition the probability. However, ϕ_{ij} deviates from ψ_j by amounts ranging from 7% to 33%, except for ϕ_{00} , which deviates from ψ_0 by a factor of about 170. The reason why ϕ_{00} is so low is intuitive: Sand dissipates with small probability ϵ (e.g., $\epsilon = 10^{-3}$ in the above simulation), and 00-sand links can appear only when a grain of sand dissipates as it is sent from a node to a 0-sand neighbor.

There are also correlations of higher order than the pairwise correlations shown in Eq. (2b), one of the simplest being three-star correlations (i.e., correlations between the amounts of sand on a node and on its three neighbors). Consider, for instance, $\theta_{\tilde{k}}$, the probability for a 2-sand node to have \tilde{k} many 2-sand neighbors out of its three neighbors. If there were no three-star correlations, then one would expect the number of 2-sand neighbors to be binomially distributed, i.e., $\theta_{\tilde{k}} = \binom{3}{\tilde{k}}(\phi_{22})^{\tilde{k}}(1 - \phi_{22})^{3-\tilde{k}}$. Using $\phi_{22} \approx 0.534 \ 32$ obtained in Eq. (2b), this hypothesis predicts

$$(\theta_0 \ \theta_1 \ \theta_2 \ \theta_3) \approx (0.100 \ 99 \quad 0.347 \ 61 \quad 0.398 \ 85 \quad 0.152 \ 55). \quad (3a)$$

However, simulations on a network of $N = 10^7$ nodes using a dissipation of $\epsilon = 10^{-3}$ provide the values

$$(\theta_0 \ \theta_1 \ \theta_2 \ \theta_3) \approx (0.083 \ 59 \quad 0.369 \ 41 \quad 0.407 \ 44 \quad 0.139 \ 56), \quad (3b)$$

which are significantly different from the predicted ones, revealing the presence of nontrivial three-star correlations. Nevertheless, ignoring the effects of three-star correlations by assuming independence beyond pairwise correlations is more acceptable than assuming no pairwise correlations (i.e., $\phi_{ij} \equiv \psi_j$).

In view of these subtleties, it may seem surprising that the $1/k$ assumption performs so well at predicting the cascade size distribution's power-law exponent $\tau = \gamma/(\gamma-1)$ for scale-free graphs with degree exponent $2 < \gamma < 3$. We conjecture that this success occurs because the branching process based on the $1/k$ assumption [21–23] belongs to the same universality class as the BTW process on a network. See Appendix A for mathematical justifications of this conclusion.

Although the exponent τ appears to be unaffected by the aforementioned caveats of the $1/k$ assumption, there are measures affected by these caveats, such as the total amount of sand in the system, the cascade area distribution, and the details of the cascade size distribution (besides its tail behavior). Moreover, the $1/k$ assumption does not allow one to study the system outside of the critical regime ($\epsilon > 0$), nor how the system would behave if one were to control it [26]. Sections IV–VI take a completely different, bottom-up perspective that is not fundamentally subject to these limitations like the $1/k$ assumption is. Before this endeavor, Sec. III B covers one last point important for modeling self-organizing (and especially SOC) systems.

B. Importance of self-consistency in models of self-organized dynamics

Given a dynamical system, a dynamical model of this system that uses some form of approximation may be understood as a different dynamical system. This apparently trivial statement has important implications in the presence of self-organization, especially for self-organized criticality. Although here we consider the special case where the original dynamical system is the BTW sandpile process, this general observation has wider relevance [26].

Consider the system given by the BTW sandpile process on a quenched graph \mathcal{G} sampled from the ensemble of random 3-regular graphs on N nodes. Denote by \mathcal{S} the rules for updating the system (i.e., dropping grains of sand, toppling nodes, and dissipating sand). Next, consider the system as a Markov chain over the set $\{0, 1, 2\}^N$ of all possible numbers of grains on each of the N many nodes; we denote by $\widehat{\mathcal{S}}$ the vector of probabilities for the system to be in each possible state after infinitely many cascades have occurred. The triplet $(\mathcal{G}, \mathcal{S}, \widehat{\mathcal{S}}_n)$ specifies the full system after n cascades have occurred. In the limit of many cascades, the state of the system approaches a stationary state (i.e., stationary distribution), denoted by $\widehat{\mathcal{S}}$.

As mentioned previously in Sec. II C, the stationary state $\widehat{\mathcal{S}}$ approaches a critical point (i.e., the branching factor R_0 approaches 1) in the limit of infinite system size $N \rightarrow \infty$ and vanishing dissipation $\epsilon \rightarrow 0$. Because of this criticality, the cascade size distribution for the system $(\mathcal{G}, \mathcal{S}, \widehat{\mathcal{S}})$ has a power-law tail. Moreover, as mentioned in Sec. III A, the stationary state $\widehat{\mathcal{S}}$ contains correlations among the amounts of sand on nearby nodes of the network.

Consider the Bethe lattice \mathcal{B} with coordination number 3 (i.e., an infinite 3-regular graph without boundaries). In the limit $N \rightarrow \infty$, assuming that all cascades are finite, the finite system $(\mathcal{G}, \mathcal{S}, \widehat{\mathcal{S}})$ becomes indistinguishable from the infinite system $(\mathcal{B}, \mathcal{S}, \widehat{\mathcal{S}})$. In Sec. V, we define the model $(\mathcal{B}, \mathcal{M}, \widehat{\mathcal{M}})$ that approximates the infinite system $(\mathcal{B}, \mathcal{S}, \widehat{\mathcal{S}})$ by neglecting all higher-order correlations beyond pairwise correlations. The equilibrium state $\widehat{\mathcal{M}}$ is fully specified by ψ_i and ϕ_{ij} because, by definition, the model \mathcal{M} cannot contain higher-order correlations. One could define another state $\widehat{\mathcal{M}}_{\widehat{\mathcal{S}}}$ by setting the model's values of ψ_i and ϕ_{ij} to be equal to those measured in the stationary state $\widehat{\mathcal{S}}$. However, it is unlikely that this empirical state $\widehat{\mathcal{M}}_{\widehat{\mathcal{S}}}$ would be critical under the rules \mathcal{M} : Among the numerous possible equilibrium states $\widehat{\mathcal{M}}$, rare are the states critical under the rules \mathcal{M} . Hence, the model $(\mathcal{B}, \mathcal{M}, \widehat{\mathcal{M}}_{\widehat{\mathcal{S}}})$ with empirically measured parameters would typically predict a cascade size distribution *without* a power-law tail.

Since the rules \mathcal{M} approximate the rules \mathcal{S} , it is very possible that they have a similar SOC mechanism. If this is the case (which we show to be true in Sec. V), then the equilibrium state $\widehat{\mathcal{M}}$ approaches a critical point as $\epsilon \rightarrow 0$ (note that the limit $N \rightarrow \infty$ is already accounted for in \mathcal{B}). Thus, like in the infinite-size equilibrated system $(\mathcal{B}, \mathcal{S}, \widehat{\mathcal{S}})$, the model $(\mathcal{B}, \mathcal{M}, \widehat{\mathcal{M}})$ predicts a power-law cascade size distribution.

Hence, although the model with empirical parameters $\widehat{\mathcal{M}}_{\widehat{\mathcal{S}}}$ is closer to the equilibrated system $\widehat{\mathcal{S}}$ than the self-consistent equilibrated model $\widehat{\mathcal{M}}$ is in terms of the ψ_i and ϕ_{ij} , the self-consistent model $(\mathcal{B}, \mathcal{M}, \widehat{\mathcal{M}})$ performs better at reproducing the behavior of the system $(\mathcal{B}, \mathcal{S}, \widehat{\mathcal{S}})$ than does the model with empirical parameters $(\mathcal{B}, \mathcal{M}, \widehat{\mathcal{M}}_{\widehat{\mathcal{S}}})$. By using a model that is self-consistent (i.e., in its own equilibrium state), we harness the power of the SOC mechanism encoded in the rules of that model. We expect that a similar statement (that letting the model reach equilibrium is probably more effective than forcing it to have the parameters observed in the real system) holds for network structures other than random 3-regular graphs. Moreover, the self-organized system of interest need not be critical for the self-consistency of a model to be important, although the presence of a critical point induces a strong dependence in the model's parameters that amplifies the importance of self-consistency. Sections VC and VD present methods resulting in such a self-consistent model.

IV. MICROSCOPIC UNDERSTANDING OF THE BTW SANDPILE PROCESS ON NETWORKS

This section provides rigorous results concerning the inner workings of a sandpile cascade. We begin by characterizing in Sec. IVA the constraints affecting the number of times that nodes can topple in a cascade. Our main results are summarized in Theorem 1 (for the root of any cascade) and Theorem 2 (for nonroot nodes in a treelike cascade). In Corollary 1, we unpack the results of these theorems to make them useful for approximating cascades using branching processes in Secs. V–VI. Using Theorems 1 and 2, we prove in Sec. IVB sufficient conditions for a cascade to form a tree, a case in which we obtain strong results. The main result, Theorem 3, specifies sufficient conditions for using the full toolbox developed in Sec. IVA. An example use of Theorem 3 is Corollary 2, which establishes a correspondence between

cascade area and bond percolation on the subgraph induced by nodes at capacity. Other examples of using Theorem 3 are given in Secs. V–VI. Proofs are deferred to the Appendixes.

A. Constraints applicable to treelike cascades

In treelike sandpile cascades on networks that are not necessarily treelike, causality constrains the shape of cascades, the number of times that each node topples, and how many grains neighboring nodes exchange. The purpose of this section is to identify such constraints because they are essential to our later calculations of cascade area and size. Understanding Corollary 1 and the associated Fig. 2 suffices to understand the sections that follow, whereas Theorems 1 and 2 are the main mathematical results that concisely summarize our characterization of the sandpile model on networks. Proofs are in Appendix B.

Our strong statements (Theorems 1 and 2) require a notion of time and causality. To make time considerations well defined and independent of the numerical implementation of

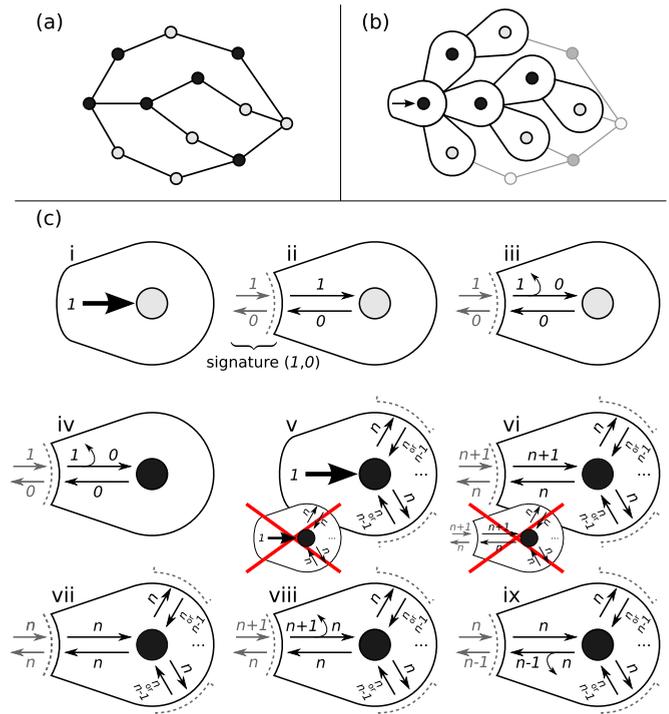


FIG. 2. (Color online) Finite treelike cascades as an assemblage of patterns. (a) The graph before the cascade. Dark circles represent nodes initially at capacity, while light circles represent nodes initially not at capacity. (b) Starting at the root (indicated by an arrow), the cascade is assembled from patterns that encode how the grains of sand are shed during the process. Note that this assemblage may form a tree even though the original graph contained cycles. (c) The only possible patterns and rules for assembling them. The roman numerals correspond to the cases in Corollary 1. The convex shapes i and v represent root patterns, while the concave shapes represent nonroot patterns. The large arrow represents the first grain dropped on the root, a curved arrow indicates dissipation, and the other arrows indicate exchanges of sand between neighboring nodes. The signature of a nonroot pattern may be inferred by the arrows on its left. A large \times marks a forbidden special case for patterns v and vi.

the sandpile process, we assume that each grain sent from one node to another takes a positive and possibly random amount of time to reach its target or to dissipate and we assume that nodes topple as soon as they exceed their capacity.

We are now ready to formulate strong constraints on the root (and on its surroundings) in any cascade on any network \mathcal{G} . Note that the phrase “a node v receives n grains from a single neighbor” has the intended meaning “for at least one of v 's neighbors, v received n grains from that neighbor.”

Theorem 1: Strong constraints on the root in any cascade. In any cascade, the root topples zero times if and only if the root was not initially at capacity. Moreover, for any positive integer n , the root topples n times by time $t \geq 0$ if and only if (a) the root was initially at capacity and (b) by time t the root received from each of its neighbors either n or $n - 1$ grains (except not n grains from each of its neighbors).

Whereas Theorem 1 provides strong constraints for the root in any cascade, Theorem 2 provides strong constraints for nonroot nodes in cascades that form a finite tree (which is the case, or nearly the case, for most cascades on treelike networks, including configuration model networks). To precisely define this notion of a cascade forming a tree, for a cascade on the graph \mathcal{G} , we define the graph \mathcal{G}^\dagger to have all the nodes of \mathcal{G} that have sand sent toward them in the cascade and all the edges of \mathcal{G} along which sand is sent in the cascade. That is, \mathcal{G}^\dagger is the subgraph induced by the root, the nodes that topple, and the neighbors of the nodes that topple, from which we remove the links between pairs of nodes that both do not topple. We say that a cascade forms a finite tree \mathcal{G}^\dagger if \mathcal{G}^\dagger is a finite tree.

We introduce the following nomenclature for a cascade that forms a tree. When a node sends a grain to each of its neighbors, those neighbors are the node's *children* in the cascade and the toppled node is the *parent* of these children. Only the first time at which a node receives sand during a cascade matters for this nomenclature. Hence, if a child later sends a grain to its parent, it does not acquire a new title (the child remains a child, and its parent remains its parent). We say that the root belongs to generation 0, and the children of a node in generation g belong to generation $g + 1$. The *ancestors* (resp. *descendants*) of a node is the set of nodes composed of its parent (resp. children), if any, and of the ancestors of this parent (resp. the descendants of these children). Hence the root is the ancestor of every other node receiving sand during a cascade and these nodes are all descendants of the root.

Theorem 2: Strong constraints on nonroot nodes in cascades forming a finite tree. For a cascade that forms a finite tree \mathcal{G}^\dagger , a nonroot node v topples zero times by time $t > 0$ if and only if v was not initially at capacity or v receives zero grains from its parent by time t . Moreover, under the same conditions and for any positive integer n , a nonroot node v topples n times by time $t > 0$ if and only if all of the following conditions hold: (a) v was initially at capacity; (b) v received n or $n + 1$ grains from its parent by time t ; and (c) v received n or $n - 1$ grains from each of its children by time t (except not n grains from every child if v received $n + 1$ grains from its parent).

Together, Theorems 1 and 2 provide strong constraints for any node in a cascade forming a finite tree, and this result holds at any time during such a cascade. Because the ultimate outcome of a cascade is of particular interest to our application, we specialize the mathematical apparatus of Theorems 1 and

2 to the form in Corollary 1, which turns out to be useful when elaborating a branching process that predicts the outcome of a cascade. Before doing so, some definitions are required.

For a cascade that forms a finite tree \mathcal{G}^\dagger , we associate a *pattern* with each node v in \mathcal{G}^\dagger . A pattern inherits all properties of its associated node, such as whether the node is the root, whether it is at capacity, and whether it is a parent or child with respect to another node.

Each nonroot pattern has a *signature* given by a pair of integers (n, n') , which characterizes grains exchanged between this pattern and its parent. Specifically, given a node v with parent u , node v is associated with a pattern of signature (n, n') if and only if the parent u sends n grains toward v and the parent u receives n' grains from v . Note that, due to dissipation, the child v may receive fewer than n grains from u and v may send more than n' grains toward u . The intuition to keep in mind is that we count grains *from the parent's perspective* (i.e., n and n' are the numbers of grains sent from and received by the parent u with respect to this particular child v). Though each pattern has a single well-defined signature, two different patterns may share the same signature. For simplicity, we say that a nonroot node v has signature (n, n') if v is associated with a pattern that has signature (n, n') .

We are now ready to provide the corollary of Theorems 1 and 2 enumerating the rules (illustrated in Fig. 2) that enable our approximations of cascade size and area in Secs. V–VI.

Corollary 1: Constraints for patterns. The following statements hold for a cascade that forms a finite tree \mathcal{G}^\dagger .

A node v in \mathcal{G}^\dagger topples zero times if and only if exactly one of the following holds.

- (i) v is the root and is not initially at capacity.
- (ii) v is nonroot; v is not initially at capacity; v has signature $(1, 0)$; and the grain sent by v 's parent toward v reaches v .
- (iii) v is nonroot; v is not initially at capacity; v has signature $(1, 0)$; and the grain sent by v 's parent toward v dissipates.
- (iv) v is nonroot; v is initially at capacity; v has signature $(1, 0)$; and the grain sent by v 's parent toward v dissipates.

For a positive integer n , a node v in \mathcal{G}^\dagger topples n times if and only if exactly one of the following holds.

- (v) v is the root, and each of v 's children has signature (n, n) or $(n, n - 1)$, except not all of its children may have signature (n, n) .
- (vi) v is nonroot; v has signature $(n + 1, n)$; and each of v 's children has signature (n, n) or $(n, n - 1)$, except not all of its children may have signature (n, n) .
- (vii) v is nonroot; v has signature (n, n) ; and each of v 's children has signature (n, n) or $(n, n - 1)$.
- (viii) v is nonroot; v has signature $(n + 1, n)$; each of v 's children has signature (n, n) or $(n, n - 1)$; and the last grain sent by v 's parent toward v dissipates.

- (ix) v is nonroot; v has signature $(n, n - 1)$; each of v 's children has signature (n, n) or $(n, n - 1)$; and the last grain sent by v toward its parent dissipates.

Figure 3 gives an explicit example of a cascade represented by the patterns (i)–(ix) enumerated in Corollary 1.

B. Sufficient condition for cascades to form a tree on locally treelike networks: A tool for calculating cascade area and size

Equipped with the results of Sec. IV A regarding the possible patterns that may appear in a cascade, we now

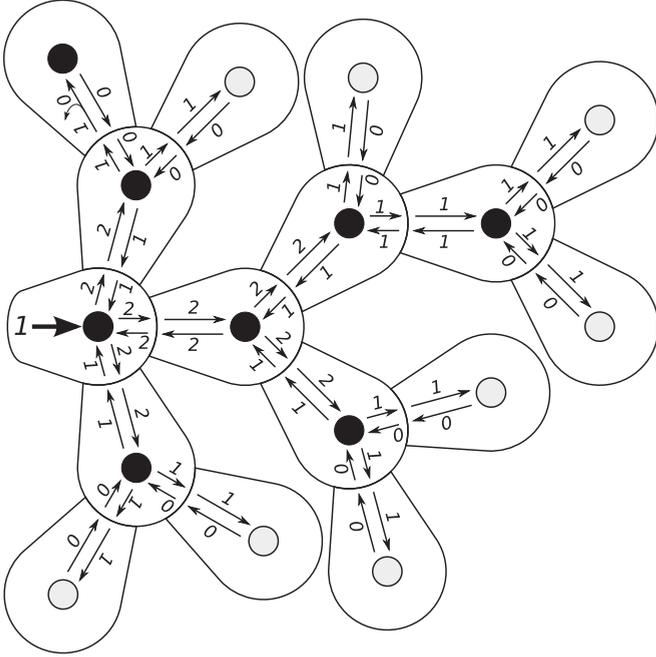


FIG. 3. Example of a cascade forming a finite tree \mathcal{G}^\dagger from the pattern perspective. Refer to the caption of Fig. 2 for notation. There are 16 nodes in \mathcal{G}^\dagger , so 16 patterns are required. All 8 nodes initially not at capacity do not topple [all pattern (ii)]. A single node at capacity does not topple (top left) because the grain sent towards it dissipates [pattern (iv)]. The 7 remaining nodes are all initially at capacity: 2 nodes each topple twice [the root, center left, has pattern (v) and the node immediately to the right of the root has pattern (vii)] and the 5 others topple once [4 have pattern (vi) and 1 has pattern (vii)]. The resulting cascade area is 7 and the cascade size is 9.

establish a sufficient condition for the cascade to form a tree, a case in which we can access the full power of Theorems 1 and 2. Readers interested only in the application of these rigorous results may skip this section with the message that, for locally treelike networks, a finite cascade forms a finite tree, so only the nodes that are initially at capacity can topple during that cascade (Theorem 3). A consequence of this result is that, under some reasonable conditions, obtaining the probability distribution for cascade area amounts to a bond percolation problem with bond occupation probability $1 - \epsilon$ on the subgraph of the nodes that were at capacity before the cascade (Corollary 2). Proofs are provided in Appendix C.

Now we make these messages precise. First, we make a definition for a node's surroundings to be locally treelike. Our definition is strict: We require all sufficiently small neighborhoods of a node to be a tree. Specifically, fix a positive integer M and a node v in a finite graph \mathcal{G} . We say that the triplet (\mathcal{G}, M, v) is *good* if, for every subgraph U of \mathcal{G} that contains node v and at most $M - 1$ other nodes, v belongs to a component of U that is a tree (i.e., that contains no cycles).

Denote by $\tilde{\mathcal{G}}$ the subgraph of \mathcal{G} induced by the nodes at capacity. Consider a cascade occurring on $\tilde{\mathcal{G}}$ as a cascade occurring on \mathcal{G} (starting with the same root) but with the nodes not at capacity treated as sinks that dissipate all sand sent to them. Theorem 3 specifies conditions guaranteeing that these

nodes not at capacity would not topple even if they were not treated as sinks.

Theorem 3: Sufficient condition for the cascade to form a tree and for ignoring nodes not at capacity. Let \mathcal{G}' be the subgraph induced by the nodes that topple in a cascade begun from the root v on a graph \mathcal{G} and let \mathcal{A} be the area of this cascade. Let $\tilde{\mathcal{G}}'$ be the subgraph induced by the nodes that topple in a cascade begun from the same root v on the graph $\tilde{\mathcal{G}}$ induced by the nodes at capacity in \mathcal{G} and let $\tilde{\mathcal{A}}$ be the area of this cascade. (If v is not in $\tilde{\mathcal{G}}$, define $\tilde{\mathcal{G}}'$ as empty and $\tilde{\mathcal{A}} = 0$.) If M is an integer such that (\mathcal{G}, M, v) is good and if $\tilde{\mathcal{A}} \in \{0, 1, \dots, M - 1\}$, then the following hold: The cascade forms a finite tree \mathcal{G}^\dagger , $\tilde{\mathcal{G}}'$ is a tree; $\mathcal{G}' = \tilde{\mathcal{G}}'$; and $\mathcal{A} = \tilde{\mathcal{A}}$.

We demonstrate the usefulness of Theorem 3 by the example of Corollary 2: Under the conditions of Theorem 3, bond percolation is the same as cascade area. Denote by $\tilde{\mathcal{G}}^{(1-\epsilon)}$ the graph that results from bond percolation on $\tilde{\mathcal{G}}$ with bond occupation probability $1 - \epsilon$. That is, $\tilde{\mathcal{G}}^{(1-\epsilon)}$ is the subgraph of $\tilde{\mathcal{G}}$ in which every link of $\tilde{\mathcal{G}}$ has probability $1 - \epsilon$ to appear.

Corollary 2: Correspondence between cascade area and bond percolation on nodes at capacity. Suppose that a cascade begins at a node v in a graph \mathcal{G} whose subgraph of nodes at capacity is $\tilde{\mathcal{G}}$. If M is an integer such that (\mathcal{G}, M, v) is good, then for any $x \in \{0, 1, \dots, M - 1\}$, the chance that the cascade area \mathcal{A} equals x is identical to the chance that v belongs to a component of size x in $\tilde{\mathcal{G}}^{(1-\epsilon)}$.

Corollary 2 enables a shortcut to estimating the cascade area distribution. Given a random root node v in a graph \mathcal{G} drawn from some ensemble of random graphs, if the neighborhood of v is a tree [specifically, if (\mathcal{G}, M, v) is good], then Corollary 2 guarantees that the cascade area distribution is identical to the component size distribution in bond percolation on the subgraph of nodes at capacity, at least for cascade area smaller than M .

More generally, if (\mathcal{G}, M, v) is good, then by Theorem 3 a cascade on \mathcal{G} starting from root v and of area smaller than M forms a finite tree \mathcal{G}^\dagger , so Theorems 1 and 2 and Corollary 1 hold. This conclusion allows the calculation of not only the cascade area but also of the inner workings of the cascade, including the cascade size.

We leave for future work the task of characterizing the probability that (\mathcal{G}, M, v) is good for various random graphs. Large configuration model random graphs (including random 3-regular graphs) are locally treelike, so for a random node v it is likely that (\mathcal{G}, M, v) is good for reasonably large M . We explore such random graphs next.

V. ZERO-PARAMETER APPROXIMATION OF THE BTW PROCESS ON RANDOM 3-REGULAR GRAPHS

Our goal in this section and the next is to calculate the probability distributions of cascade area and size for the BTW process on networks using the rigorous results of Sec. IV whenever possible. In this section, we focus on random 3-regular networks and use self-consistency arguments to obtain a zero-parameter model that poises itself at a stationary state.

Section V A uses standard percolation methods (justified by Corollary 2) to estimate the cascade area distribution using a single-type branching process. This step facilitates introducing

the multitype branching process in Sec. VB that harnesses the powerful results in Sec. IV to obtain both the cascade size and area distributions. Up to this point, the model has two parameters that need to be fixed (i.e., ψ_2 and ϕ_{22}). In Sec. VC, we further improve the model so that it predicts the expected changes of its own parameters. Finally, Sec. VD uses the assumption that these parameters reached an equilibrium, which results in our self-consistent zero-parameter model.

A. Standard percolation: Cascade area

In order to facilitate the introduction of alternative techniques in the following sections, here we outline a rather standard percolation approach to calculate cascade area for the BTW sandpile process on a random 3-regular network \mathcal{G} . The intuition behind the calculation is that if the neighborhood of the root is a tree and if the cascade area is small enough to fit inside this tree neighborhood, then Corollary 2 shows that cascade area is given by bond percolation on the subgraph $\tilde{\mathcal{G}}$ of nodes at capacity.

Suppose, however, that $\tilde{\mathcal{G}}$ is unknown to us and that we only know three assumptions: (i) \mathcal{G} is a large random 3-regular network; (ii) a uniformly random node is at capacity with probability ψ_2 ; and (iii) a node adjacent to a node at capacity is itself at capacity with probability ϕ_{22} . We thus *approximate* $\tilde{\mathcal{G}}$ from this information and we obtain an estimate of the cascade area distribution from the component size distribution of the graph $\tilde{\mathcal{G}}^{(1-\epsilon)}$, the result of bond percolation on $\tilde{\mathcal{G}}$ with bond occupancy $1 - \epsilon$.

To this end, we use a standard single-type branching process approach [29,34] based on probability generating functions (PGFs). We define the cascade area PGF $H(x) \equiv \sum_{a=0}^{\infty} \mathbb{P}(\text{area is } a)x^a$; the dummy variable x is called the generator for the cascade area. The function $H(x)$ is obtained by solving the system

$$F(x) = [1 - (1 - \epsilon)\phi_{22}] + (1 - \epsilon)\phi_{22}x[F(x)]^2, \quad (4a)$$

$$H(x) = (1 - \psi_2) + \psi_2x[F(x)]^3. \quad (4b)$$

The PGF $F(x)$ [Eq. (4a)] gives the contribution to the cascade area from a node sending a grain to a neighbor v that has not yet toppled: The grain reaches node v with probability $1 - \epsilon$ and v is at capacity with probability ϕ_{22} . If both these events occur, then v topples (factor x) and sends grains toward its three neighbors, two of which have not yet toppled (factor $[F(x)]^2$). In $H(x)$ [Eq. (4b)], the root is initially at capacity with probability ψ_2 , in which case it topples (factor x) and sends a grain toward its three neighbors (factor $[F(x)]^3$).

B. Inner workings: Cascade size

We now improve Eq. (4) to explicitly obtain the cascade size distribution in addition to the cascade area. This time, instead of Corollary 2 we base our approach on the more fundamental Corollary 1 and on a pattern representation of the cascade.

Keeping the generator x for the cascade area, we introduce the generator w for the cascade size such that $H(w, x) \equiv \sum_{a=0}^{\infty} \sum_{s=0}^{\infty} \mathbb{P}(\text{area is } a, \text{ size is } s)x^a w^s$. Instead of $F(x)$ in Eq. (4a) we need two families of functions ${}_n\mathcal{A}(w, x)$ and ${}_n\mathcal{B}(w, x)$, which track the contributions of patterns of

signature $(n, n - 1)$ and (n, n) , respectively. Dropping explicit dependences in w and x for simplicity, the system of equations becomes

$${}_1\mathcal{A} = \epsilon + \left(\frac{\epsilon}{1 - \epsilon}\right) {}_1\mathcal{B} + (1 - \epsilon)(1 - \phi_{22}), \quad (5a)$$

$${}_{n>1}\mathcal{A} = \epsilon_{n-1} {}_n\mathcal{B} + \left(\frac{\epsilon}{1 - \epsilon}\right) {}_n\mathcal{B} + (1 - \epsilon)^{2n-1} \phi_{22} x w^{n-1} \\ \times [({}_{n-1}\mathcal{A} + {}_{n-1}\mathcal{B})^2 - ({}_{n-1}\mathcal{B})^2], \quad (5b)$$

$${}_n\mathcal{B} = (1 - \epsilon)^{2n} \phi_{22} x w^n ({}_n\mathcal{A} + {}_n\mathcal{B})^2, \quad (5c)$$

$$H = (1 - \psi_2) + \psi_2 x \sum_{n=1}^{\infty} w^n [({}_n\mathcal{A} + {}_n\mathcal{B})^3 - ({}_n\mathcal{B})^3]. \quad (5d)$$

Next we explain how to obtain Eq. (5).

The factor ${}_n\mathcal{B}(w, x)$ [Eq. (5c)] corresponds to the contribution of a node v with pattern of signature (n, n) [i.e., pattern (vii) in Corollary 2 and Fig. 2]. The factor $(1 - \epsilon)^{2n}$ accounts for the probability that all the grains traveling from and to the parent of v do not dissipate, while ϕ_{22} accounts for the probability that v is at capacity at the beginning of the cascade (knowing that its parent was at capacity at the beginning of the cascade). The generators xw^n count v 's contribution of 1 to the cascade area and n to the cascade size. Finally, the factor $({}_n\mathcal{A} + {}_n\mathcal{B})^2$ accounts for the contribution of the two children of v , each of which may have signature $(n, n - 1)$ or (n, n) .

The three terms composing ${}_1\mathcal{A}(w, x)$ [Eq. (5a)] correspond to the contribution of nodes v that have different patterns of signature $(1, 0)$. The first term ϵ accounts for the probability ϵ that the grain sent toward v dissipates before reaching it [i.e., pattern (iii) or (iv) with $n = 1$], in which case there is no contribution to the cascade area or size. The second term considers the possibility that v topples, but the grain sent to its parent dissipates before reaching it [i.e., pattern (ix)]: Expanding that second term of Eq. (5a) gives

$$\left(\frac{\epsilon}{1 - \epsilon}\right) {}_1\mathcal{B} = (1 - \epsilon)\epsilon\phi_{22}xw({}_1\mathcal{A} + {}_1\mathcal{B})^2,$$

where the factor $(1 - \epsilon)\epsilon$ accounts for the probability that v receives the grain sent by its parent, but not the converse; and the rest of the expression is obtained in the same way as in ${}_1\mathcal{B}$. Finally, the third term of ${}_1\mathcal{A}$ corresponds to pattern (ii): v receives the grain sent by its parent [factor $(1 - \epsilon)$], but v does not topple because it is initially not at capacity (factor $1 - \phi_{22}$).

Similarly, the three terms composing ${}_n\mathcal{A}(w, x)$ with $n > 1$ [Eq. (5b)] correspond to the contribution of the different possible patterns of signature $(n, n - 1)$. Here the first term corresponds to the case in which the last grain sent by the parent dissipates [i.e., pattern (viii)], which amounts to ${}_{n-1}\mathcal{B}$ with a factor ϵ for the extra dissipation. Likewise, the second term corresponds to the case in which the last grain sent by v to its parent dissipates [i.e., pattern (ix)], which amounts to ${}_n\mathcal{B}$ with a factor $\epsilon/(1 - \epsilon)$ rectifying the probability for one extra sand dissipation and one fewer successful sand transfer. The last term of Eq. (5b) corresponds to pattern (vi): None of the grains exchanged between v and its parent dissipate [factor $(1 - \epsilon)^{2n-1}$]; v is initially at capacity (factor ϕ_{22}); and v topples

$n - 1$ times (factor xw^{n-1}). The remaining factor $[(_{n-1}\underline{A} + _{n-1}\underline{B})^2 - (_{n-1}\underline{B})^2]$ accounts for the contribution of the two children of this pattern, each of which may have signature $(n - 1, n - 2)$ or $(n - 1, n - 1)$, except not both may simultaneously have signature $(n - 1, n - 1)$.

Finally, $H(w, x)$ [Eq. (5d)] tracks the contribution of root patterns. With probability $1 - \psi_2$, the root is not initially at capacity, so it does not topple [i.e., pattern (i)]. Conversely, with probability ψ_2 , the root is initially at capacity and topples $n \geq 1$ times (factor xw^n). The three children of this pattern (v) may have signature $(n, n - 1)$ or (n, n) , but not all of them may have signature (n, n) (factor $[(_{n-1}\underline{A} + _n\underline{B})^3 - (_n\underline{B})^3]$).

One may check that setting $w = 1$ in Eq. (5) recovers Eq. (4): First verify that the ansatz $_n\underline{B}(1, x) = _{n+1}\underline{A}(1, x) + _{n+1}\underline{B}(1, x)$ satisfies Eq. (5), and then observe that $F(x) = _1\underline{A}(1, x) + _1\underline{B}(1, x)$ and that $H(x) = H(1, x)$ through a telescoping series.

Except for the structure of the network \mathcal{G} and for the dissipation ϵ (with the case $\epsilon \rightarrow 0$ being of particular interest), the BTW sandpile process on the network defined in Sec. II B has no free parameter. However, both Eq. (4) (based on a standard percolation approach) and the more refined Eq. (5) (accounting for the inner workings of a cascade) depend on two unknown quantities: ψ_2 and ϕ_{22} . These two quantities are not real parameters of the BTW process; instead they reflect the state of the system once it has reached equilibrium.

As discussed in Sec. III B, one could measure ψ_2 and ϕ_{22} from the steady state of numerical simulations (i.e., $\widehat{\mathcal{S}}$) and then use these values in Eq. (4) or (5) to estimate the probability distributions for cascade area and/or size [which amounts to the model $(\mathcal{B}, \mathcal{M}, \widehat{\mathcal{M}}_{\widehat{\mathcal{S}}})$]. Proceeding in this way suffers the major disadvantage that results cannot be obtained outside the regime in which simulations were performed. Moreover, due to the reasons discussed in Sec. III B, the resulting model would be very inaccurate: Using values from Eq. (2), both Eqs. (4) and (5) predict a branching factor of $R_0 = 2(1 - \epsilon)\phi_{22} \approx 1.07$, which is supercritical. Hence, feeding the model the ψ_2 and ϕ_{22} observed in simulations would predict the existence of giant cascades spanning a considerable fraction of the network in the limit $N \rightarrow \infty$, in contradiction with numerical simulations. (Note that the nongiant cascades would be affected by an exponential cutoff.) The following sections show how a self-consistency argument provides a zero-parameter model not subject to these disadvantages [which amounts to the model $(\mathcal{B}, \mathcal{M}, \widehat{\mathcal{M}})$].

C. Changes in the number of i -sand nodes and ii -sand links

Here we augment the multitype branching process in Eq. (5) to obtain the effect of a cascade on the numbers of i -sand nodes and of ii -sand links in the network. This augmentation is a crucial step toward obtaining the zero-parameter model presented next: We will poise the model at an equilibrium by requiring that the quantities ψ_i and ϕ_{ij} are on average unaffected by a cascade.

As before, w and x are generators for the cascade size and area, respectively. In addition, we define vectors of generators $\mathbf{y} = (y_0, y_1, y_2)$ and $\mathbf{z} = (z_0, z_1, z_2)$ such that y_i and z_i generate the changes in the numbers of i -sand nodes and of ii -sand links, respectively. Positive powers of y_i or z_i correspond to

an increase in the respective counts, whereas negative powers correspond to decreases. For convenience, we define

$$Y_i^{i'}(\mathbf{y}) \equiv \frac{y_{i'}}{y_i}, \quad Z_{ij}^{i'j'}(\mathbf{z}) \equiv \frac{1 + \delta_{i'j'}(z_{i'} - 1)}{1 + \delta_{ij}(z_i - 1)},$$

$$L_i^{i'}(\mathbf{z}) = \sum_{j=0}^2 \phi_{ij} Z_{ij}^{i'j'}, \quad (6)$$

where δ_{ij} is Kronecker's delta. A factor of $Y_i^{i'}(\mathbf{y})$ should be included whenever an i -sand node becomes i' -sand because $Y_i^{i'}(\mathbf{y})$ accounts for the change in the amount of sand on the node associated with this pattern (i.e., one fewer i -sand node and one more i' -sand node). Similarly, a factor of $Z_{ij}^{i'j'}(\mathbf{z})$ should be included whenever an ij -sand link becomes $i'j'$ -sand. Nonroot patterns (whether they topple or not) are responsible for tracking the factor $Z_{ij}^{i'j'}(\mathbf{z})$ due to the link joining them to their parent. We call a leaf a node v that does not receive sand but has a neighbor u that does receive sand (i.e., v is a leaf of the branching process). If u was initially i -sand and ends up being i' -sand after the cascade, then v has probability ϕ_{ij} to initially be (and to remain) j -sand, so a factor $L_i^{i'}(\mathbf{z})$ should be included for the link joining u to v .

In order to account for the factor $Z_{ij}^{i'j'}(\mathbf{z})$, a nonroot pattern must know how many grains of sand end up being on its parent at the end of the cascade. Hence we define the two families of functions $_n^i\underline{A}(w, x, \mathbf{y}, \mathbf{z})$ and $_n^i\underline{B}(w, x, \mathbf{y}, \mathbf{z})$ that track the contributions of nodes that have patterns of signature $(n, n - 1)$ and (n, n) , respectively, and that have parents who end up being i' -sand after the cascade. The new PGF $H(w, x, \mathbf{y}, \mathbf{z})$ is now obtained by solving a system of equations of structure very similar to Eq. (5),

$$_1^i\underline{A} = \epsilon L_2^{i'} + \left(\frac{\epsilon}{1 - \epsilon} \right) _1^i\underline{B} + (1 - \epsilon) \sum_{j=0}^1 \phi_{2j} Y_j^{j+1} Z_{2j}^{i'(j+1)} (L_j^{j+1})^2, \quad (7a)$$

$$_{n>1}^i\underline{A} = \epsilon _{n-1}^i\underline{B} + \left(\frac{\epsilon}{1 - \epsilon} \right) _n^i\underline{B} + (1 - \epsilon)^{2n-1} \phi_{22} x w^{n-1} \times \sum_{j'=1}^2 \binom{2}{j'-1} Y_2^{j'} Z_{22}^{i'j'} (_{n-1}^i\underline{A})^{2-(j'-1)} (_{n-1}^i\underline{B})^{j'-1}, \quad (7b)$$

$$_n^i\underline{B} = (1 - \epsilon)^{2n} \phi_{22} x w^n \sum_{j'=0}^2 \binom{2}{j'} Y_2^{j'} Z_{22}^{i'j'} (_n^i\underline{A})^{2-j'} (_n^i\underline{B})^{j'}, \quad (7c)$$

$$H = \sum_{i=0}^1 \psi_i Y_i^{i+1} (L_i^{i+1})^3 + \psi_2 x \sum_{n=1}^{\infty} w^n \sum_{i'=0}^2 \binom{3}{i'} Y_2^{i'} (_n^i\underline{A})^{3-i'} (_n^i\underline{B})^{i'}. \quad (7d)$$

Next we explain how to obtain Eq. (7).

All the terms present in the equation for $_n^i\underline{B}$ [Eq. (5c)] are still present in the equation for $_n^i\underline{B}$ [Eq. (7c)]. This time, we explicitly consider the signature of the two children of the node

v with pattern (vii). Suppose $j' \in \{0,1,2\}$ of these children have signature (n,n) and $2-j'$ have signature $(n,n-1)$. Summing the two grains initially on the node v , the n grains received from v 's parent, and the $nj' + (n-1)(2-j')$ received from v 's children, we obtain $3n + j'$ grains, enough for node v to topple n times and to end up j' -sand. A factor $Y_2^{j'}$ is required because v was 2-sand and becomes j' -sand. A factor $Z_{22}^{i'j'}$ is required because the link from v to its parent was 22-sand and becomes $i'j'$ -sand. The factor $\binom{j'}{n} \binom{2-j'}{n} \binom{j'}{j'}$ accounting for v 's children passes on the information that v ends up being j' -sand. The combinatorial factor $\binom{2}{j'}$ accounts for the number of ways to choose the signatures of the children.

Similarly, the terms present in the equation for $1_{\underline{A}}$ [Eq. (5a)] are also present in the equation for $1'_{\underline{A}}$ [Eq. (7a)]. The first term of Eq. (7a) corresponds to a leaf node [i.e., a node that does not receive any sand, pattern (iii) or (iv)], which requires a factor $L_2^{i'}$. The second term is obtained similarly to $1'_{\underline{B}}$ in Eq. (7c). The third term corresponds to the pattern (ii): For $j \in \{0,1\}$, the considered node v is initially j -sand with probability ϕ_{2j} , so v receives a single grain (factor Y_j^{j+1}) and does not topple. The link from v to its parent mandates a factor $Z_{2j}^{i'(j+1)}$, and the two other neighbors of v qualify as leaves, mandating a factor $(L_j^{j+1})^2$.

Again, the terms present in the equation for $n_{>1}\underline{A}$ [Eq. (5b)] are also present in the equation for $n_{>1}'\underline{A}$ [Eq. (7b)]. The first two terms are obtained similarly as in Eq. (7c). The last term corresponds to a node v with pattern (vi): $j' - 1 \in \{0,1\}$ of the two children of v have signature $(n-1, n-1)$, while the remaining $2 - (j' - 1)$ children have signature $(n-1, n-2)$. Summing the two grains initially on v , the n grains received from v 's parent, and the $(n-1)(j' - 1) + (n-2)[2 - (j' - 1)]$ grains received from v 's children, we obtain $3(n-1) + j'$ grains, enough for v to topple $n-1$ times and end up j' -sand. Contributions similar to those in Eq. (7c) are thus required.

Finally, the first term of the equation for H [Eq. (7d)] corresponds to the root not being at capacity [pattern (i)]. For $i \in \{0,1\}$, the root is i -sand with probability ψ_i and thus becomes $(i+1)$ -sand (factor Y_i^{i+1}). Moreover, the three neighbors of the root qualify as leaves [factor $(L_i^{i+1})^3$]. The second term corresponds to the root being at capacity [pattern (v)]: $i' \in \{0,1,2\}$ of the root's children have signature (n,n) and $3-i'$ have signature $(n,n-1)$. Summing the two grains initially on the root, the one grain dropped on the root to start the cascade, and the $ni' + (n-1)(3-i')$ grains received by the root from its children, we obtain $3n + i'$ grains, enough for the root to topple n times and to end up i' -sand. The other factors follow the same logic as in Eq. (5d).

One may check that setting $y_i = z_i = 1$ for all $i \in \{0,1,2\}$ in Eq. (7) recovers Eq. (5). Specifically, defining $\mathbf{1} = (111)$, we have $Y_i^{i'}(\mathbf{1}) = Z_{ij}^{i'j'}(\mathbf{1}) = L_i^{i'}(\mathbf{1}) = 1$, $1'_{\underline{A}}(w,x,\mathbf{1},\mathbf{1}) = 1_{\underline{A}}(w,x)$, and $1'_{\underline{B}}(w,x,\mathbf{1},\mathbf{1}) = 1_{\underline{B}}(w,x)$ for all the possible index values, which gives $H(w,x,\mathbf{1},\mathbf{1}) = H(w,x)$.

D. Bootstrapping ψ_i and ϕ_{ij}

Now we obtain the self-consistent zero-parameter model by enforcing equilibrium. The function $H(w,x,y,z)$ in Eq. (7d)

is a PGF of the multivariate probability distribution for the cascade size (generator w), the cascade area (generator x), the change in the number of i -sand nodes (generator y_i), and the change in the number of ii -sand links (generator z_i). By the standard techniques for PGFs, we obtain expectation values through differentiation with respect to the appropriate generator. In this case, we are interested in $h_{(i)}$, the average change in the number of i -sand nodes, and in $\eta_{(i)}$, the average change in the number of ii -sand links:

$$h_{(i)} = \left. \frac{\partial H(1,1,\mathbf{y},\mathbf{1})}{\partial y_i} \right|_{\mathbf{y}=\mathbf{1}}, \quad \eta_{(i)} = \left. \frac{\partial H(1,1,\mathbf{1},\mathbf{z})}{\partial z_i} \right|_{\mathbf{z}=\mathbf{1}}. \quad (8)$$

An explicit expression of these derivatives $h_{(i)}$ and $\eta_{(i)}$ is provided in [31].

By hypothesis, the system has reached a stationary state, so both these average changes should be zero. Hence, the relations

$$h_{(i)} = 0 \quad \forall i \in \{0,1,2\}, \quad \eta_{(i)} = 0 \quad \forall i \in \{0,1,2\} \quad (9a)$$

provide six constraints for the 12 unknowns $\psi_i \forall i \in \{0,1,2\}$ and $\phi_{ij} \forall i,j \in \{0,1,2\}$. Only five of these constraints are independent because each of the $h_{(i)} = 0$ relations can be obtained from the two others. Because the ψ_i are probabilities and the ϕ_{ij} are conditional probabilities, they must obey the seven additional, independent constraints

$$\sum_{i=0}^2 \psi_i = 1, \quad \sum_{j=0}^2 \phi_{ij} = 1 \quad \forall i, \quad \psi_i \phi_{ij} = \psi_j \phi_{ji} \quad \forall i,j. \quad (9b)$$

Thus, we have a total of 12 unknowns and 12 independent constraints on them. Because the constraints (9a) are nonlinear, *a priori* there is no guarantee that a valid solution exists or that there is a single solution.

However, starting from educated guesses, numerical solution of the system (9) does provide values of ψ_i and ϕ_{ij} that are consistent with those observed in the Monte Carlo simulations. For example, the solution of Eq. (9) for $\epsilon = 10^{-3}$ gives

$$(\psi_0 \quad \psi_1 \quad \psi_2) \approx (0.09574 \quad 0.35439 \quad 0.54987), \quad (10a)$$

$$\begin{pmatrix} \phi_{00} & \phi_{01} & \phi_{02} \\ \phi_{10} & \phi_{11} & \phi_{12} \\ \phi_{20} & \phi_{21} & \phi_{22} \end{pmatrix} \approx \begin{pmatrix} 0.00050 & 0.27802 & 0.72148 \\ 0.07511 & 0.34286 & 0.58203 \\ 0.12562 & 0.37512 & 0.49926 \end{pmatrix}. \quad (10b)$$

Comparison with the results from simulations in Eq. (2) reveals that, while the correspondence is not exact, a large fraction of the previously unexplained correlations is now accounted for. Part of the deviations is explained by the fact that Eq. (2) is obtained from simulations with $N = 10^7$ nodes, while Eq. (10) assumes $N \rightarrow \infty$. As discussed in Sec. III B, other deviations are likely due to Eq. (10) only considering pairwise correlations, while correlations between the amount of sand on nodes arbitrarily far from one another could appear in the network simulations, especially for small ϵ . Section VI B provides guidelines on how future work could consider such higher-order correlations by obtaining a more refined version of Eq. (7).

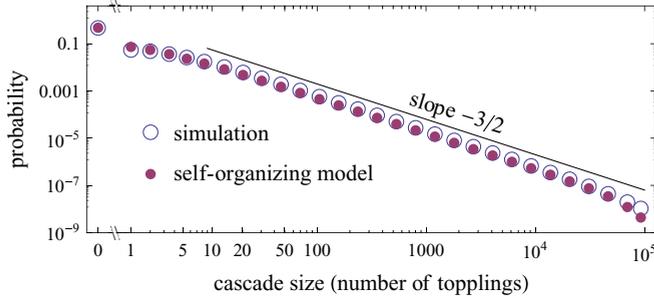


FIG. 4. (Color online) Cascade size distribution for a 3-regular network with $N = 10^7$ nodes and dissipation rate $\epsilon = 10^{-3}$. The data for both Monte Carlo simulations (open circles) and theory (closed circles) have been logarithmically binned. The theory is obtained for ψ_i and ϕ_{ij} shown in Eq. (10), i.e., those satisfying the constraints [Eq. (9)] of the self-organizing model.

Using $\psi_2 \approx 0.54987$ and $\phi_{22} \approx 0.49926$ [i.e., the values from Eq. (10)] in Eq. (5) provides estimates for the cascade area and size. Figure 4 confirms the accuracy of the resulting cascade size distribution by comparing it to numerical simulations. Because cascade size and area are often close to one another, we verify in Fig. 5 that Eq. (5) predicts the right distribution for the difference of a cascade's size and area. This approach allows the independent study of cascade size and area.

In the range $0 < \epsilon < 0.01$, the ψ_i and ϕ_{ij} show close-to-linear behavior (e.g., $\phi_{00} \approx \epsilon/2$) and the total variation of each of these probabilities is less than 0.01. Linear extrapolation provides the limiting behavior $\epsilon \rightarrow 0$,

$$\lim_{\epsilon \rightarrow 0} (\psi_0 \quad \psi_1 \quad \psi_2) \approx (0.09519 \quad 0.35406 \quad 0.55075), \quad (11a)$$

$$\lim_{\epsilon \rightarrow 0} \begin{pmatrix} \phi_{00} & \phi_{01} & \phi_{02} \\ \phi_{10} & \phi_{11} & \phi_{12} \\ \phi_{20} & \phi_{21} & \phi_{22} \end{pmatrix} \approx \begin{pmatrix} 0.00000 & 0.27676 & 0.72324 \\ 0.07441 & 0.34226 & 0.58333 \\ 0.12500 & 0.37500 & 0.50000 \end{pmatrix}. \quad (11b)$$

These results support our intuition that there should be no 00-links in the limit $\epsilon \rightarrow 0$ and that the system tunes itself to the onset of a giant component (which happens at $\phi_{22} = 1/2$

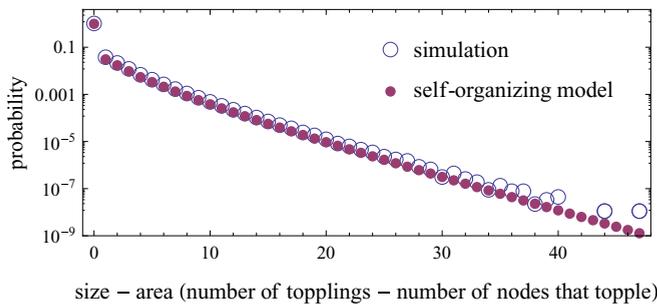


FIG. 5. (Color online) Distribution of the difference between cascade size and cascade area for a 3-regular network in the same conditions as in Fig. 4. We consider the difference between the two quantities to facilitate the comparison. Note that the size of a cascade is always greater than or equal to its area.

when $\epsilon \rightarrow 0$). Other quantities appear to approach ratios of small integers: $\phi_{20} \approx 1/8$, $\phi_{21} \approx 3/8$, and $\phi_{12} \approx 7/12$. Note that these rational numbers are indicative only: Excluding ϕ_{00} and ϕ_{22} , we have no reason to offer as to why the ϕ_{ij} should be close to the ratio of small integers (and this proximity to simple ratios may well be coincidental).

VI. GENERALIZATION TO CONFIGURATION MODEL NETWORKS

Here we show how the methods presented in Sec. V for 3-regular networks generalize to configuration model networks. Hence, Sec. VIA (resp. Sec. VIB) repeats the treatment of Sec. VA (resp. Sec. VB) to predict the cascade area distribution (resp. cascade size and area distributions) on configuration model networks. However, we do not perform the two final steps (Secs. VC and VD) that would provide a zero-parameter model: Doing so reveals to be tedious in the case of the fully general configuration model. Nonetheless, future work could apply the same approach on a case-by-case basis to other network structures of interest.

A. Standard percolation: Cascade area

We consider the same problem as in Sec. VA, that is, to estimate the cascade area distribution using standard percolation techniques, but this time we consider the general case of a configuration model random graph \mathcal{G} instead of a random 3-regular graph. Again, we base our model on Corollary 2: If the neighborhood of the root is a tree and if the cascade area is small enough to fit inside this tree neighborhood, then the cascade area is given by bond percolation on the subgraph $\tilde{\mathcal{G}}$ of nodes at capacity. This preliminary step facilitates the presentation of Sec. VIB, which considers the inner workings of the cascade to estimate both the cascade area and size distribution.

As described in Sec. IIA, the configuration model random graph \mathcal{G} is drawn from an ensemble specified by the degree distribution $\{p_k : k \geq 0\}$ and by N , the number of nodes in \mathcal{G} . We now make the assumption (approximation) that $\tilde{\mathcal{G}}$, which is the subgraph of \mathcal{G} induced by the nodes at capacity, is also a configuration model random graph. Specifically, we assume that $\tilde{\mathcal{G}}$ is a configuration model random graph drawn from an ensemble specified by the degree distribution $\{\tilde{p}_{\tilde{k}} : \tilde{k} \geq 0\}$ and by σN , the number of nodes in $\tilde{\mathcal{G}}$. Stated differently, σ is the fraction of nodes in \mathcal{G} that are at capacity at the beginning of the cascade and $\tilde{p}_{\tilde{k}}$ is the probability that a random node at capacity in \mathcal{G} has \tilde{k} neighbors at capacity in \mathcal{G} .

Following Newman *et al.* [29,34], we define $\tilde{q}_{\tilde{k}}$ as the probability that a node in $\tilde{\mathcal{G}}$ reached by following a random link of $\tilde{\mathcal{G}}$ has degree $\tilde{k} + 1$ (i.e., has \tilde{k} links in $\tilde{\mathcal{G}}$ other than the one along which we reached it). Using the standard calculation for configuration model random graphs, we obtain $\tilde{q}_{\tilde{k}}$ from $\tilde{p}_{\tilde{k}}$ as $\tilde{q}_{\tilde{k}} = (\tilde{k} + 1)\tilde{p}_{\tilde{k}+1} / \sum_{\tilde{k}'} \tilde{k}'\tilde{p}_{\tilde{k}'}$. For convenience, we define the PGFs

$$\tilde{P}(\xi) := \sum_{\tilde{k}=0}^{\infty} \tilde{p}_{\tilde{k}} \xi^{\tilde{k}}, \quad \tilde{Q}(\xi) := \sum_{\tilde{k}=0}^{\infty} \tilde{q}_{\tilde{k}} \xi^{\tilde{k}} = \frac{d\tilde{P}(\xi)}{d\xi} \Big|_{\xi=1}, \quad (12)$$

which encode the same information as the probability distributions $\tilde{p}_{\tilde{k}}$ and $\tilde{q}_{\tilde{k}}$ of degree and excess degree in $\tilde{\mathcal{G}}$. We use these functions in Eq. (12) to build a branching process providing

the component sizes in $\tilde{\mathcal{G}}^{(1-\epsilon)}$. By adapting [34], we obtain the PGFs

$$G(x) = \epsilon + (1 - \epsilon)x\tilde{Q}(G(x)), \quad (13a)$$

$$H(x) = (1 - \sigma) + \sigma x\tilde{P}(G(x)). \quad (13b)$$

The PGF $G(x)$ [Eq. (13a)] generates the probability distribution for the size of the component of $\tilde{\mathcal{G}}^{(1-\epsilon)}$ reached by following a random link of $\tilde{\mathcal{G}}$. Note that this link in $\tilde{\mathcal{G}}$ is also in $\tilde{\mathcal{G}}^{(1-\epsilon)}$ with probability $1 - \epsilon$, so this component size equals zero with probability ϵ . The full component size distribution of $\tilde{\mathcal{G}}^{(1-\epsilon)}$ is generated by $x\tilde{P}(G(x))$. A cascade of area zero occurs with probability $1 - \sigma$ (the fraction of nodes not at capacity). Hence, $H(x)$ [Eq. (13b)] is the PGF for our approximation of the distribution of the area of cascades in $\tilde{\mathcal{G}}$.

One may verify that setting $\sigma = \psi_2$, $\tilde{P}(\xi) = [(1 - \phi_{22}) + \phi_{22}\xi]^3$, and $\tilde{Q}(\xi) = [(1 - \phi_{22}) + \phi_{22}\xi]^2$ in Eq. (13) recovers the 3-regular special case (4), with the correspondence $F(x) = (1 - \phi_{22}) + \phi_{22}G(x)$.

B. Inner workings: Cascade size

This section is to Sec. VIA as Sec. VB is to Sec. VA: We improve Eq. (13) by accounting for the inner workings of a cascade to explicitly obtain both the cascade area and size distribution [analogously to our improvement to Eq. (4) in Sec. VB for random 3-regular graphs]. In addition to $\tilde{P}(\xi)$ and $\tilde{Q}(\xi)$ [Eq. (12)], we define the related PGFs

$$\tilde{P}^*(\xi) = \sum_{\tilde{k}=0}^{\infty} \tilde{p}_k^* \xi^{\tilde{k}}, \quad \tilde{Q}^*(\xi) = \sum_{\tilde{k}=0}^{\infty} \tilde{q}_k^* \xi^{\tilde{k}}. \quad (14)$$

Here \tilde{p}_k^* is the probability that a node at capacity has \tilde{k} neighbors in \mathcal{G} and that all \tilde{k} neighbors are at capacity. Similarly, \tilde{q}_k^* is the probability for a node at capacity to have $\tilde{k} + 1$ neighbors in \mathcal{G} that are all at capacity, given that this node was reached by following a link from a node that was at capacity. Note that, unlike $\tilde{P}(1)$ and $\tilde{Q}(1)$, neither $\tilde{P}^*(1)$ nor $\tilde{Q}^*(1)$ will typically equal one. The purpose of $\tilde{P}^*(\xi)$ and $\tilde{Q}^*(\xi)$ is to properly consider particular constraints imposed by Corollary 1, such as the fact that the children of patterns (v) and (vi) may not all simultaneously have signature (n, n) .

We again use the generator x to track the cascade area and the generator w to track the cascade size. We define the two families of PGFs ${}_n\mathcal{C}(w, x)$ and ${}_n\mathcal{D}(w, x)$ that are related to the families ${}_n\mathcal{A}(w, x)$ and ${}_n\mathcal{B}(w, x)$, respectively, except that we are now only considering nodes at capacity (i.e., nodes in $\tilde{\mathcal{G}}$). Specifically, ${}_n\mathcal{C}(w, x)$ [resp. ${}_n\mathcal{D}(w, x)$] accounts for the contribution of a nonroot pattern at capacity of signature $(n, n - 1)$ [resp. $(n, n - 1)$], i.e., patterns (iv), (vi), (viii), and (ix) [resp. pattern (vii)]. The relations defining $H(w, x)$ are

$${}_1\mathcal{C} = \epsilon \left(1 + \frac{{}_1\mathcal{D}}{1 - \epsilon} \right), \quad (15a)$$

$${}_2\mathcal{C} = \epsilon \left({}_1\mathcal{D} + \frac{{}_2\mathcal{D}}{1 - \epsilon} \right) + (1 - \epsilon)^3 x w [\tilde{Q}({}_1\mathcal{C} + {}_1\mathcal{D}) - \tilde{Q}^*({}_1\mathcal{D})], \quad (15b)$$

$${}_{n>2}\mathcal{C} = \epsilon \left({}_{n-1}\mathcal{D} + \frac{{}_n\mathcal{D}}{1 - \epsilon} \right) + (1 - \epsilon)^{2n-1} x w^{n-1} \times [\tilde{Q}^*({}_{n-1}\mathcal{C} + {}_{n-1}\mathcal{D}) - \tilde{Q}^*({}_{n-1}\mathcal{D})], \quad (15c)$$

$${}_1\mathcal{D} = (1 - \epsilon)^2 x w \tilde{Q}({}_1\mathcal{C} + {}_1\mathcal{D}), \quad (15d)$$

$${}_{n>1}\mathcal{D} = (1 - \epsilon)^{2n} x w^n \tilde{Q}^*({}_n\mathcal{C} + {}_n\mathcal{D}), \quad (15e)$$

$$H = (1 - \sigma) + \sigma x \left\{ w [\tilde{P}({}_1\mathcal{C} + {}_1\mathcal{D}) - \tilde{P}^*({}_1\mathcal{D})] + \sum_{n=2}^{\infty} w^n [\tilde{P}^*({}_n\mathcal{C} + {}_n\mathcal{D}) - \tilde{P}^*({}_n\mathcal{D})] \right\}. \quad (15f)$$

The ${}_n\mathcal{D}(w, x)$ [Eqs. (15d) and (15e), pattern (vii)] are obtained very similarly to the ${}_n\mathcal{B}$ [Eq. (5c)]. In the case of ${}_1\mathcal{D}(w, x)$, there is no particular constraint to be satisfied, so the children of the considered pattern contribute a total of $\tilde{Q}({}_1\mathcal{C} + {}_1\mathcal{D})$ (the children not at capacity, and hence not in $\tilde{\mathcal{G}}$, each contribute a factor 1). In the case of ${}_n\mathcal{D}(w, x)$ for $n > 1$ [Eq. (15e)], we know that the children of the considered node all topple at least once, so they were all at capacity (i.e., they are all in $\tilde{\mathcal{G}}$) and thus contribute a total of $\tilde{Q}^*({}_n\mathcal{C} + {}_n\mathcal{D})$.

Obtaining ${}_1\mathcal{C}(w, x)$ [Eq. (15a)] is very similar to obtaining ${}_1\mathcal{A}(w, x)$ [Eq. (5a)], except that the patterns (ii) and (iii) need not be considered because the considered node is known to be at capacity. Similarly, the case of ${}_2\mathcal{C}(w, x)$ and ${}_{n>2}\mathcal{C}(w, x)$ [Eqs. (15b) and (15c)] is similar to the one of ${}_{n>1}\mathcal{A}(w, x)$ [Eq. (5b)]. For ${}_2\mathcal{C}(w, x)$, only the contribution of the children in the last term [corresponding to pattern (vi)] requires further explanation. Using $\tilde{Q}({}_1\mathcal{C} + {}_1\mathcal{D})$ would consider that children may have signature (1, 0) or (1, 1), which includes a spurious term where all children have signature (1, 1); subtracting $\tilde{Q}^*({}_1\mathcal{D})$ cancels this spurious term. A similar cancellation is used in ${}_{n>2}\mathcal{C}(w, x)$, but this time all the children are known to be at capacity (since they topple at least once).

Finally, $H(w, x)$ [Eq. (15f)] is obtained very similarly to its counterpart for 3-regular networks [Eq. (5d)]. With probability $1 - \sigma$, the root is not at capacity [pattern (i)]. With probability σ , the root is at capacity [pattern (v)]: A cancellation similar to the one of ${}_2\mathcal{C}(w, x)$ [resp. ${}_{n>2}\mathcal{C}(w, x)$] is used if the root topples once (resp. more than once).

One may recover Eq. (13) by setting $w = 1$ [with the correspondences $G(x) = {}_1\mathcal{C}(1, x) + {}_1\mathcal{D}(1, x)$ and $H(x) = H(1, x)$, using the ansatz ${}_n\mathcal{D}(1, x) = {}_{n+1}\mathcal{C}(1, x) + {}_{n+1}\mathcal{D}(1, x)$]. Alternatively, one may recover Eq. (5) by setting $\sigma = \psi_2$, $\tilde{P}(\xi) = [(1 - \phi_{22}) + \phi_{22}\xi]^3$, $\tilde{Q}(\xi) = [(1 - \phi_{22}) + \phi_{22}\xi]^2$, $\tilde{P}^*(\xi) = (\phi_{22}\xi)^3$, and $\tilde{Q}^*(\xi) = (\phi_{22}\xi)^2$, with the correspondences ${}_1\mathcal{A}(w, x) = (1 - \phi_{22}) + {}_1\mathcal{C}(w, x)\phi_{22}$, ${}_{n>1}\mathcal{A}(w, x) = {}_{n>1}\mathcal{C}(w, x)\phi_{22}$, and ${}_n\mathcal{B}(w, x) = {}_n\mathcal{D}(w, x)\phi_{22}$; the PGF $H(w, x)$ remains $H(w, x)$.

It was mentioned earlier that one could improve the model presented in Sec. V by accounting for three-star correlations. A starting point for doing so could be the following special case of the configuration model considered in Eq. (15) obtained by fixing $\sigma = \psi_2$ and

$$\tilde{P}(\xi) = \theta_0 + \theta_1\xi + \theta_2\xi^2 + \theta_3\xi^3, \quad \tilde{P}^*(\xi) = \theta_3\xi^3, \quad (16a)$$

$$\tilde{Q}(\xi) = \frac{\theta_1 + 2\theta_2\xi + 3\theta_3\xi^2}{\theta_1 + 2\theta_2 + 3\theta_3}, \quad \tilde{Q}^*(\xi) = \frac{3\theta_3\xi^2}{\theta_1 + 2\theta_2 + 3\theta_3}. \quad (16b)$$

Directly using the values of Eq. (3b) would cause the same problems as described at the end of Sec. VB. In particular, the predicted branching factor $R_0 = (2\theta_2 + 6\theta_3)/(\theta_1 + 2\theta_2 + 3\theta_3) \approx 1.03$ is still supercritical. We expect that accounting for higher-order correlations should gradually decrease the predicted branching factor towards unity.² More generally for configuration model networks, one will face the same kind of problems when feeding values obtained from simulations into Eq. (15). To circumvent this problem, one could continue to generalize the method of Sec. V by implementing a system similar to Eq. (7), this time for configuration model networks, and then perform an analysis similar to Sec. VD to obtain the equilibrium state without the need for simulations. However, considering the general case may prove tedious and specifically considering special cases of interest [e.g., Eq. (16)] is likely more promising. We defer to future work the consideration of these questions.

VII. CONCLUSION

The BTW sandpile process is an archetypal example of self-organized criticality, a term blanketing any process that has a critical point as an attractor. We argue that, although this article focuses on the BTW sandpile process on networks, the understanding and tools developed herein are applicable to the modeling of a vast array of SOC processes.

To predict the macroscopic observables of an SOC process (e.g., the exponent of the power-law tail of the cascade size distribution in the BTW process), it may suffice to build a model that reproduces the symmetries of the SOC process [e.g., the branch distribution (the probability distribution for the number of events caused by an event) should have mean one and the right tail behavior]. This approach will succeed if the resulting model falls in the same universality class as the studied SOC process; other details may not matter. We conjecture in Sec. III A and Appendix A that sharing a universality class with the BTW process on networks explains the success of an earlier model [21] despite its flawed $1/k$ assumption.

In other contexts, a macroscopic understanding may not suffice because, for example, we seek quantities of a microscopic nature and/or because the process is not critical. Instead, a microscopic understanding of the process is likely required. In the context of the BTW process, we studied the internal workings of a cascade (Sec. IV), which enabled the possibility to distinguish cascade size and area (Secs. VB and VIB).

As discussed in Sec. III B, it is important to acknowledge that a model of an SOC process is a different dynamical system in itself: Forcing the parameters of a model to those measured in the SOC process may not be the best approach. The alternative that we recommend is to instead force the model to be self-consistent, with the hope that the model self-organizes at the right equilibrium state. Sections VC and VD perform this feat in the context of the BTW process on random 3-regular networks: The resulting model predicts with appreciable

accuracy nontrivial pairwise correlations in the equilibrium state, and it allows us to explore the BTW process in ranges prohibitively costly to simulate. Section VIB explains how one could generalize the approach to more complex networks.

Furthermore, our method can allow one to study the effects of controlling an SOC process [26]. Though our results were obtained in a specific context, this work serves as a proof of concept that paves the road for powerful, self-consistent, and microscopically accurate models of real world systems that self-organize to critical points.

Finally, in the broader context of modeling processes taking place on random graphs, we emphasize that our use of multitype branching processes effectively allows repeated back-and-forth exchanges of information between a parent vertex and its children. In fact, the signature of a pattern associated with a vertex v provide to v 's ancestors, to v 's siblings, and to the descendants of v 's siblings important information concerning v 's descendants and vice versa.³ We feel that this perspective is currently underused in the literature, despite its great potential to analytically model systems in which consequences spread in a bidirectional fashion.

ACKNOWLEDGMENTS

The authors thank Kwang-Il Goh for useful discussions. This work was supported in part by the Defense Threat Reduction Agency Basic Grant No. HDTRA1-10-1-0088, the Army Research Laboratory Cooperative Agreement W911NF-09-2-0053, the Army Research Laboratory MURI Grant No. W911NF-13-1-0340, the Department of Defense through the National Defense Science & Engineering Graduate Fellowship Program (C.D.B.), and the Fonds de recherche du Québec—Nature et technologies (P.-A.N.).

APPENDIX A: JUSTIFICATION OF THE UNIVERSALITY CONJECTURE

This appendix provides informal justifications to the conjecture made in Sec. III A, namely, that the reason the relation $\tau = \gamma/(\gamma - 1)$ for scale-free graphs with degree exponent $2 < \gamma < 3$ [21–23] is verified by numerical simulations *despite* the known flaws of the $1/k$ assumption is because the branching process leading to that relation belongs to the same universality class as the BTW process on networks.

We first consider a generic Galton-Watson branching process starting with a single particle. The process is completely defined by the sequence (u_0, u_1, u_2, \dots) such that u_k is the probability that a particle disintegrates into k new identical particles. We seek the sequence (v_0, v_1, v_2, \dots) such that v_s is the probability for the total number of particles ever existing to be s .

³To give an example, suppose that v has signature (3,2), i.e., v receives three grains from its parent, who receives two grains from v . Through that signature, v 's parent effectively learns that v topples at least two times, which implies that v 's children topple at least one time. Moreover, v learns that all of v 's ancestors topple at least three times, that v 's siblings topple at least two times, and that the children of v 's siblings all topple at least one time.

²Indeed, the sand in large clusters of nodes at capacity will likely disperse due to large cascades occurring on those nodes. Hence we expect the equilibrium state to contain smaller clusters of nodes at capacity than one would expect in a randomized surrogate.

We define the PGFs

$$U(\xi) = \sum_{k=0}^{\infty} u_k \xi^k, \quad V(w) = \sum_{s=0}^{\infty} v_s w^s. \quad (\text{A1})$$

The normalization of the probabilities u_k forces $U(1) = 1$. A well known result is that if $\sum_k k u_k = 1$ (i.e., $R_0 = 1$) then (i) the fixed point of $V(w) = w U(V(w))$ determines $V(w)$; (ii) the distribution v_s is normalized [i.e., $\sum_{s=0}^{\infty} v_s = V(1) = 1$]; and (iii) v_s asymptotically follows a power law $v_s \sim s^{-\tau}$ (possibly with some corrections). Note that the condition $\sum_k k u_k = 1$ may be understood as $\lim_{\xi \rightarrow 1^-} U'(\xi) = 1$, where the prime denotes differentiation and the limit is taken from the left along the real axis. The case where u_k is light tailed, i.e., $U(\xi)$ is analytic in the complex neighborhood of $\xi = 1$, leads to the mean-field regime $\tau = 3/2$.

Goh *et al.* [21] studied the case of the heavy-tailed distribution $u_k^{\text{Goh}} = \alpha k^{-\gamma}$ for $k > 0$ with α and u_0^{Goh} chosen such that $U_{\text{Goh}}(1) = 1$ and $\lim_{\xi \rightarrow 1^-} U'_{\text{Goh}}(\xi) = 1$, i.e.,

$$U_{\text{Goh}}(\xi) = 1 + \frac{\text{Li}_{\gamma}(\xi) - \zeta(\gamma)}{\zeta(\gamma - 1)}. \quad (\text{A2})$$

In this case, they showed that

$$v_s^{\text{Goh}} \sim \begin{cases} s^{-\gamma/(\gamma-1)} & (2 < \gamma < 3) \\ s^{-3/2} (\ln s)^{-1/2} & (\gamma = 3) \\ s^{-3/2} & (\gamma > 3) \end{cases} \quad (\text{A3})$$

(note that the distribution u_k^{Goh} cannot be normalized for $\gamma \leq 2$). In their demonstration, Goh *et al.* used the fact that the behavior of v_s in the asymptotic limit $s \rightarrow \infty$ can be obtained from the singular behavior of $U(\xi)$ near $\xi = 1$.

We now generalize this result. Define $U_{\text{ana.}}(\xi)$ as an analytic function in the neighborhood of $\xi = 1$ respecting $U_{\text{ana.}}(1) = 1$ and $U'_{\text{ana.}}(1) = 1$. Also define the (possibly singular) function $U_{\text{sing.}}(\xi)$ such that $U_{\text{sing.}}(1) = 1$ and $\lim_{\xi \rightarrow 1^-} U'_{\text{sing.}}(\xi) = 1$. For ρ respecting $0 < \rho \leq 1$, we see that $\tilde{U}(\xi) = (1 - \rho)U_{\text{ana.}}(\xi) + \rho U_{\text{sing.}}(\xi)$ respects $\tilde{U}(1) = 1$, $\lim_{\xi \rightarrow 1^-} \tilde{U}'(\xi) = 1$, and $\tilde{U}(z)$ has the same singular behavior in the neighborhood of $\xi = 1$ as $U_{\text{sing.}}(\xi)$. Moreover, the sequence $\tilde{u}_k = (1 - \rho)u_k^{\text{ana.}} + \rho u_k^{\text{sing.}}$ corresponding to this $U_{\text{sing.}}(\xi)$ is a probability distribution (i.e., $0 \leq \tilde{u}_k \leq 1$ for all $k \geq 0$). Hence, the branching process (A1) using the branching PGF $U_{\text{sing.}}(\xi)$ results in \tilde{v}_s that have the same behavior in the asymptotic limit $s \rightarrow \infty$ as one would obtain with any other choice of $0 < \rho \leq 1$ and/or $U_{\text{ana.}}(\xi)$ respecting the aforementioned conditions.

In particular, choosing $U_{\text{sing.}}(\xi) = U_{\text{Goh}}(\xi)$, the \tilde{u}_k obtained with such an analytic perturbation leads to the asymptotic behavior described by Goh *et al.* We do not claim that the model of Goh *et al.* amounts to an analytic perturbation of the actual BTW process on networks. We only want to point out that there is a continuum of models that belong to the same universality class as the model of Goh *et al.* for fixed $\gamma > 2$ and that it is not impossible that the BTW process on networks is one of them.

APPENDIX B: PROOFS FOR SEC. IV A

This section proves statements from Sec. IV A, namely, Theorems 1 and 2 and their associated Corollary 1. Before proving Theorem 1, we first consider the following lemma.

Lemma 1: Constraints on any cascade. In any cascade, the following statements hold for any positive integer n .

(i) A nonroot node cannot topple for the n th time before it receives n grains from a single neighbor.

(ii) No node may receive n grains from a single neighbor before the root topples an n th time.

Proof. To prove (i), first consider the case that a nonroot node v has degree k , begins with $k - 1$ grains before the cascade, and at some time t in the cascade has received exactly $n - 1$ grains from each of its k neighbors. Then v has $kn - 1$ grains initially on it and sent to it by time t , so v has toppled at most $n - 1$ times by time t . By construction, v must receive an n th grain from at least one neighbor before v can topple an n th time. If v received fewer than $n - 1$ grains from one or more of its neighbors by time t , then v still cannot topple an n th time until it receives an n th grain from a single neighbor. This argument proves (i).

To prove (ii), let t be the first time in a cascade at which a nonroot node receives an n th grain from a single neighbor. Let v be one such node. Then a neighbor u of v toppled an n th time before time t . Suppose for contradiction that u is not the root. In order to topple n times before time t , node u must have received at least n grains from one of its neighbors by time t , which contradicts the definition of t . Thus u must be the root and claim (ii) follows. ■

We use Lemma 1 to prove Theorem 1 of Sec. IV A.

Proof of Theorem 1. To show the first claim, note that if the root is at capacity then it topples at least once and if the root is not at capacity then it does not topple and the cascade ends there.

To prove the rest, fix $n \geq 1$. If (a) and (b) hold, then the number of grains initially on and received by the root (including the first grain dropped on it) by time t is in the interval $[kn, k(n + 1) - 1]$, so the root topples n times by time t .

Inversely, if (a) does not hold, then the root topples $0 < n$ times. It remains only to show that if (b) does not hold then the root does not topple n times by time t . There are three cases. First, if the root has received n grains from each of its neighbors by time t , then (by counting grains) we know that the root has toppled $n + 1$ times by time t . Second, if the root has received $m > n$ many grains from a neighbor by time t , then that neighbor (call it u) must have toppled at least m times by time t . Thus, u must have received at least m grains from at least one of its neighbors by time t . However, by Lemma 1(ii), no node can receive m grains from a single neighbor before the root topples m times. Hence, the root toppled at least $m > n$ times by time t .

In the third and final case, the root has received $m < n - 1$ grains from at least one neighbor by time t . Let t' be the time when the root topples for the n th time. Before t' , no neighbor of the root can have received n grains from a single neighbor by Lemma 1(ii) because the root has toppled $n - 1$ or fewer times. Hence, by Lemma 1(i), no neighbor of the root can have toppled n times before time t' . Thus, the root cannot have received n or more grains from a single neighbor by time t' . To conclude, the number of grains initially on and received by the root (including the first grain dropped on it) by time t' is $\leq kn + 1 + (m - n) < kn$, which contradicts the root toppling for the n th time at time t' . This concludes the proof. ■

Lemma 2 facilitates the proof of Theorem 2.

Lemma 2: Constraints on cascades that form a finite tree. The following statements hold for a cascade that forms a finite tree \mathcal{G}^\dagger .

(i) Let n be a non-negative integer and v be any node in \mathcal{G}^\dagger . No descendant of v may receive an n th grain from a single neighbor before v topples an n th time.

(ii) Let n be a non-negative integer and v be any nonroot node in \mathcal{G}^\dagger . Then node v cannot topple an n th time before receiving an n th grain from its parent.

(iii) Let n be a positive integer and let v be a nonroot node that has received n grains from its parent by time t . Then node v has toppled n times by time t if and only if it is initially at capacity and at the moment of toppling for the n th time it has received $n - 1$ grains from every one of its children and n grains from its parent.

(iv) Let n be a positive integer and let v be a nonroot node that has toppled n times by time t . Then v 's parent toppled at most $n + 1$ times by time t .

(v) Let v be any node in \mathcal{G}^\dagger . If v is not at capacity, then v topples zero times.

Proof. To show (i), first note that if v is the root, then the claim follows directly from the analogous (but weaker) result, Lemma 1(ii), for cascades that do not form trees. To finish proving (i), suppose v is not the root and assume (for contradiction) that a descendant of v receives an n th grain from a single neighbor before v topples an n th time. Let t be the first time a descendant of v receives n grains from a single neighbor; call such a descendant u , and let w be a neighbor of u from which u receives an n th grain at time t . Then w must have toppled for an n th time at time $t' < t$. We know that w is not v because v does not topple for the n th time before time t . Because w is a descendant of v , we know w is not the root, so by Lemma 1(i) we know that w must have received at least n grains from a single neighbor before time t' , which contradicts the definition of t . Thus, claim (i) follows.

We show (ii) by contradiction. Suppose (for contradiction) that v does not receive an n th grain from its parent by time t and yet v topples an n th time at time t . Before time t , v has toppled fewer than n times, so [by Lemma 2(i) applied to v] no children of v have received n grains from the same neighbor. Thus, no children of v have toppled n times by time t [by Lemma 1(i)], so v does not receive n grains from the same child by time t . However, by Lemma (i) and the assumption that v topples at time t , we know v must have received n grains from a single neighbor before time t , a contradiction. Thus, claim (ii) follows.

Claim (iii) follows from counting grains of sand. Let k be the degree of the nonroot node v that has received $n \geq 1$ grains from its parent by time t . Suppose that v topples for an n th time at some time $t' \leq t$. Before time t' , Lemma 2(i) guarantees that the children of v toppled at most $n - 1$ times, so v received at most $n - 1$ grains from each of its children by time t' . Moreover, v received at most n grains from its parent by time t' (because it received n grains from its parent by time $t \geq t'$). Because v topples for the n th time at time t' , the total number of grains initially on and received by v by time t' should be kn . From the preceding constraints, this is only possible if v is initially at capacity, v receives $n - 1$ grains

from each of its children by time t' , and v receives n grains from its parent by time t' .

Conversely, suppose that v is initially at capacity and suppose that t' is the first time v has received $n - 1$ grains from every one of its children and n grains from its parent. Then, by the previous grain-counting argument, v topples an n th time at time t' . It remains to be proven that v does not topple again by time t , which is guaranteed by Lemma 2(ii) because v received only n grains from its parent by time t . Thus claim (iii) holds.

We show (iv) by induction over the generation $g \geq 1$ of the nonroot v that topples $n \geq 1$ times by time t . Let u be the parent of v . Suppose u is the root (i.e., $g = 1$) and suppose (for contradiction) that u topples $m \geq n + 2$ times by time t . By Theorem 1, we know that u has received either m or $m - 1$ grains from each of its children by time t , including node v . However, by assumption, v topples $n \leq m - 2$ times by time t , so u receives $m - 2$ or fewer grains from v by time t , a contradiction. Thus, claim (iv) holds for $g = 1$.

Now suppose a node v at generation $g > 1$ topples $n \geq 1$ times by time t and assume that the claim holds at generation $g - 1$. Suppose (for contradiction) that v 's parent u toppled $m \geq n + 2$ times by time t and let $t' \leq t$ be the moment when u topples for the m th time. Before time t' , u toppled at most $m - 1$ times, so by the inductive hypothesis we know that u 's parent has toppled $\leq m$ times. Moreover, by Lemma 2(ii), u receives at least m grains from its parent by time t' because its parent topples an m th time. Thus, u receives exactly m grains from its parent by time t' and u topples an m th time at time t' , so we can apply Lemma 2(iii) to u to conclude that u must have received $m - 1$ grains from each of its children (including node v) by time t' . However, v has toppled $n \leq m - 2$ times by time t , so there is no time $t' \leq t$ at which u receives an $(m - 1)$ th grain from v , a contradiction. Thus, (iv) follows by induction on g .

Claim (v) is already shown if v is the root (Theorem 1); here we show the case in which v is nonroot by contradiction. If v is not at capacity, then it must receive at least two grains (from any source) before toppling. Consider (for contradiction) the time $t' \leq t$ at which t receives a second grain. Before t' , v has toppled zero times, so v cannot receive grains from its children [by Lemma 2(i)] and v cannot have received more than one grain from its parent [because v 's parent cannot have toppled more than once by Lemma 2(iv)]. Hence, v cannot receive a second grain by time t' , a contradiction. So claim (v) holds. ■

We use Lemmas 1 and 2 and Theorem 1 to prove Theorem 2 of Sec. IV A.

Proof of Theorem 2. We first prove the first claim concerning v toppling zero times. By 2(v), v topples zero times if v is not initially at capacity. Now suppose that the nonroot node v receives zero grains from its parent by time t . By Lemma 2(i) with $n = 0$, v cannot receive a grain from one of its children by time t . Thus, v receives a total of zero grains by time t and hence topples zero times by time t .

Inversely, if v is initially at capacity and v receives at least one grain from its parent by time t , then v clearly topples at least once by time t . This concludes the proof of the claim for v toppling zero times.

To prove the second claim, first suppose that conditions (a)–(c) hold. Let k be the degree of v . By counting grains, we see that the number of grains initially on v and received by v by time t is in the interval $[kn, k(n+1) - 1]$, so v topples n times by time t .

To show the converse, suppose one of (a), (b), or (c) does not hold and suppose (for contradiction) that v topples n times by time t . If (a) does not hold (i.e., if v is not initially at capacity), then v topples zero times by time t [Lemma 2(v)], a contradiction.

Next consider the two cases in which (b) does not hold. If v received fewer than n grains from its parent by time t , then by Lemma 2(ii) v topples fewer than n times by time t , a contradiction. If v received more than $n+1$ grains from its parent by time t and if v 's parent were the root, then by Theorem 1 node v would necessarily topple $n+1$ or more times by time t , a contradiction. Finally, if v received more than $n+1$ grains from its parent by time t and if v 's parent were not the root, then by part (c) of this theorem applied to the parent of v , we know that v topples $n+1$ or more times by time t , a contradiction.

To finish the proof, consider the three ways in which (c) may not hold. If v received fewer than $n-1$ grains from any of its children by time t , then v toppled fewer than n times by time t [by Lemma 2(iii)], a contradiction. If v received more than n grains from any of its children by time t , then v toppled more than n times by time t [because Lemma 2(ii) implies that such a child must have received more than n grains from its parent, i.e., from v], a contradiction. In the last case, v received $n+1$ grains from its parent and n grains from each of its children. Let k be the degree of v . By counting grains, we see that the number of grains initially on v and received by v by time t is $k(n+1)$, so v topples $n+1$ times by time t , a contradiction. Hence the second claim holds, which completes the proof. ■

Finally, we use Theorems 1 and 2 to prove Corollary 1 of Sec. IV A.

Proof of Corollary 1. We prove claims (i)–(ix) by letting t be some time after the cascade finishes and by applying Theorems 1 and 2.

We first consider the cases in which v topples zero times. Suppose v is the root. By Theorem 1, point (i) is necessary and sufficient for v to topple zero times.

Now suppose v is nonroot and topples zero times. Because v topples zero times, v 's parent receive zero grains from v , so the parent toppled at most one time (by Theorem 1 or 2 if the parent is the root or not, respectively), and at least one time (otherwise v would not be in \mathcal{G}). Thus, the parent of v has toppled exactly one time. Hence, the signature of v is $(1,0)$. By Theorem 2, v is not initially at capacity and/or v receives zero grains from its parent during the cascade. These conditions leave only three possibilities: v is not initially at capacity and receives the grain sent by its parent [point (ii)]; v is not initially at capacity and the grain dissipates [point (iii)]; or v is initially at capacity and the grain dissipates [point (iv)]. There are no other possibilities.

Next consider the cases in which v topples $n \geq 1$ times. If v is the root, then point (v) is equivalent to Theorem 1 and the fact that v 's children are the same as v 's neighbors.

Now suppose v is not the root. Let (n', m') be the signature of v ; let m be the amount of sand received by v from its parent; and let l_c be the amount of sand received by v from one of its children c [that child thus has signature (n, l_c)]. By Theorem 2, it is necessary and sufficient that v was initially at capacity, v received n or $n+1$ grains from its parent (so $m \geq n$), and v received from each of its children n or $n-1$ grains (so $n \geq l_c \geq n-1$ for every child c of v), except that v cannot receive n grains from all of its children if it received $n+1$ grains from its parent. Moreover, v cannot receive more grains from its parent than the number of times the parent toppled (so $n' \geq m$); v 's parent cannot receive more grains from v than the number of times v toppled (so $n \geq m'$); and v 's parent must receive at least $n'-1$ grains from v (by Theorem 1 or 2 if the parent is the root or not, respectively, so $m' \geq n'-1$). Grain exchanges between v and its parent may thus be summarized as $n' \geq m \geq n \geq m' \geq n'-1$, which leaves four possibilities (i.e., four possible positions of the $>$ symbol): $n' > m = n = m' = n'-1$ [the last grain sent by the parent of v toward v dissipated, point (viii)], $n' = m > n = m' = n'-1$ [no dissipation, v cannot receive n grains from all its children, point (vi)], $n' = m = n > m' = n'-1$ [the last grain sent by v toward its parent dissipated, point (ix)], and $n' = m = n = m' > n'-1$ [no dissipation, point (vii)]. There are no other possibilities. ■

APPENDIX C: PROOFS FOR SEC. IV B

This section proves statements from Sec. IV B, namely, Theorem 3 and its associated Corollary 2.

Proof of Theorem 3. If v is not initially at capacity, then it does not topple, so \mathcal{G}^\dagger contains only v (a valid tree), $\tilde{\mathcal{G}}'$ is empty, and $\mathcal{A} = 0$. By definition, v does not belong to $\tilde{\mathcal{G}}$ because v is not at capacity, so $\tilde{\mathcal{G}}'$ is also empty (a valid tree) and $\tilde{\mathcal{A}} = 0$.

Otherwise, suppose v is at capacity. Then v topples at least once in the cascades on both \mathcal{G} and $\tilde{\mathcal{G}}$, so v is in all of the graphs \mathcal{G}^\dagger , \mathcal{G}' , and $\tilde{\mathcal{G}}'$. By definition, the number of nodes in $\tilde{\mathcal{G}}'$ equals the area $\tilde{\mathcal{A}}$ of the cascade on $\tilde{\mathcal{G}}$. Since $\tilde{\mathcal{A}} < M$ and (\mathcal{G}, M, v) is good, we know that $(\tilde{\mathcal{G}}, \tilde{\mathcal{A}}, v)$ is good too. Also $(\tilde{\mathcal{G}}, \tilde{\mathcal{A}}, v)$ and $(\mathcal{G}', \tilde{\mathcal{A}}, v)$ are both good because $\tilde{\mathcal{G}}'$ is a subgraph of $\tilde{\mathcal{G}}$ and because $\tilde{\mathcal{G}}$ is a subgraph of \mathcal{G} . Since $\tilde{\mathcal{A}} < M$ and $(\tilde{\mathcal{G}}, \tilde{\mathcal{A}}, v)$ is good, we know that $\tilde{\mathcal{G}}'$ is a forest. Moreover, because the cascade spreads between adjacent nodes starting from the root v , $\tilde{\mathcal{G}}'$ is connected and thus $\tilde{\mathcal{G}}'$ is a tree.

It remains only to show that the cascade forms a finite tree \mathcal{G}^\dagger because that conclusion will imply the equalities $\mathcal{G}' = \tilde{\mathcal{G}}'$ and $\mathcal{A} = \tilde{\mathcal{A}}$. First we show that \mathcal{G}' is a subgraph of \mathcal{G}^\dagger . Observe that \mathcal{G}^\dagger can be obtained by the union of the root v , the graph \mathcal{G}' , the neighbors of nodes that topple, and the links joining nodes that topple to their neighbors. Since there are no links between nodes that do not topple, we may consider individually each additional node u adjacent to a node that topples such that u itself does not topple. Fix one such u . We define $\tilde{\mathcal{G}}''(u)$ to be the subgraph of \mathcal{G} induced by $\{u\} \cup \{\text{all nodes of } \tilde{\mathcal{G}}'\}$. Now $\tilde{\mathcal{G}}''(u)$ contains $\tilde{\mathcal{A}} + 1 \leq M$ nodes, so $(\tilde{\mathcal{G}}''(u), \tilde{\mathcal{A}} + 1, v)$ is good. Thus u is a leaf of the tree $\tilde{\mathcal{G}}'$. Hence the presence of node u in \mathcal{G}^\dagger (together with its single link to a node that topples) cannot

introduce a cycle in \mathcal{G}^\dagger , so \mathcal{G}^\dagger is a tree and the cascade forms a finite tree $\tilde{\mathcal{G}}^\dagger$. Finally, Theorem 2 guarantees that nodes not at capacity do not topple, so $\mathcal{G}' = \tilde{\mathcal{G}}'$ and thus $\mathcal{A} = \tilde{\mathcal{A}}$. ■

Proof of Corollary 2. The case $x = 0$ is trivial: $\mathcal{A} = 0$ if and only if v is not at capacity, which is true if and only if v belongs to a component of size 0 in $\tilde{\mathcal{G}}^{(1-\epsilon)}$.

Otherwise suppose v is at capacity. Then v topples at least once. Let $\tilde{\mathcal{C}}_v^{(1-\epsilon)}$ denote the connected component to which v belongs in $\tilde{\mathcal{G}}^{(1-\epsilon)}$. If the number of nodes in $\tilde{\mathcal{C}}_v^{(1-\epsilon)}$ is less than M , then because (\mathcal{G}, M, v) is good, we know that

$\tilde{\mathcal{C}}_v^{(1-\epsilon)}$ is a tree. Moreover, by Theorem 3, for $x < M$ we have $\mathcal{G}' = \tilde{\mathcal{G}}'$ and $\mathcal{A} = \tilde{\mathcal{A}}$. Thus, for $x < M$, $\tilde{\mathcal{C}}_v^{(1-\epsilon)}$ is built in the same way as the cascade grows on the subgraph $\tilde{\mathcal{G}}$ of nodes at capacity (with nodes not at capacity treated as sinks). Specifically, because $x < M$ and (\mathcal{G}, M, v) is good, both $\tilde{\mathcal{C}}_v^{(1-\epsilon)}$ and the subgraph \mathcal{G}' of \mathcal{G} induced by nodes that topple are built by adjoining at-capacity neighbors of leaves independently with probability $1 - \epsilon$. Thus, for $x < M$, the chance that $\tilde{\mathcal{C}}_v^{(1-\epsilon)}$ has x many nodes equals the chance that the area $\mathcal{A} = x$. ■

-
- [1] I. Dobson, B. A. Carreras, V. E. Lynch, and D. E. Newman, *Chaos* **17**, 026103 (2007).
- [2] X. Gabaix, P. Gopikrishnan, V. Plerou, and H. E. Stanley, *Nature (London)* **423**, 267 (2003).
- [3] D. E. Juanico and C. Monterola, *J. Phys. A* **40**, 9297 (2007).
- [4] T. Ribeiro, M. Copelli, F. Caixeta, and H. Belchior, *PLoS ONE* **5**, e14129 (2010).
- [5] A. Haimovici, E. Tagliazucchi, P. Balenzuela, and D. R. Chialvo, *Phys. Rev. Lett.* **110**, 178101 (2013).
- [6] A. Saichev and D. Sornette, *Phys. Rev. E* **70**, 046123 (2004).
- [7] S. Hergarten, *Nat. Hazard. Earth Syst. Sci.* **3**, 505 (2003).
- [8] G. L. Mamede, N. A. M. Araujo, C. M. Schneider, J. C. de Araújo, and H. J. Herrmann, *Proc. Natl. Acad. Sci. USA* **109**, 7191 (2012).
- [9] B. D. Malamud, G. Morein, and D. L. Turcotte, *Science* **281**, 1840 (1998).
- [10] P. Sinha-Ray and H. J. Jensen, *Phys. Rev. E* **62**, 3215 (2000).
- [11] E. T. Lu and R. J. Hamilton, *Astrophys. J.* **380**, L89 (1991).
- [12] M. Paczuski, S. Boettcher, and M. Baiesi, *Phys. Rev. Lett.* **95**, 181102 (2005).
- [13] P. Bak, C. Tang, and K. Wiesenfeld, *Phys. Rev. Lett.* **59**, 381 (1987).
- [14] P. Bak, C. Tang, and K. Wiesenfeld, *Phys. Rev. A* **38**, 364 (1988).
- [15] C. Barrett, K. Channakeshava, F. Huang, J. Kim, A. Marathe, M. V. Marathe, G. Pei, S. Saha, B. S. P. Subbiah, and A. K. S. Vullikanti, *PLoS ONE* **7**, e45406 (2012).
- [16] D. Helbing, *Nature (London)* **497**, 51 (2013).
- [17] L. de Arcangelis and H. J. Herrmann, *Physica A* **308**, 545 (2002).
- [18] J. Lahtinen, J. Kertész, and K. Kaski, *Physica A* **349**, 535 (2005).
- [19] E. Bonabeau, *J. Phys. Soc. Jpn.* **64**, 327 (1995).
- [20] S. Lise and M. Paczuski, *Phys. Rev. Lett.* **88**, 228301 (2002).
- [21] K.-I. Goh, D.-S. Lee, B. Kahng, and D. Kim, *Phys. Rev. Lett.* **91**, 148701 (2003).
- [22] D.-S. Lee, K.-I. Goh, B. Kahng, and D. Kim, *Physica A* **338**, 84 (2004).
- [23] K.-I. Goh, D.-S. Lee, B. Kahng, and D. Kim, *Physica A* **346**, 93 (2005).
- [24] C. D. Brummitt, R. M. D'Souza, and E. A. Leicht, *Proc. Natl. Acad. Sci. USA* **109**, E680 (2012).
- [25] K.-M. Lee, K.-I. Goh, and I. M. Kim, *J. Korean Phys. Soc.* **60**, 641 (2012).
- [26] P.-A. Noël, C. D. Brummitt, and R. M. D'Souza, *Phys. Rev. Lett.* **111**, 078701 (2013).
- [27] B. Dupoyet, H. Fiebig, and D. Musgrove, *Physica A* **390**, 3120 (2011).
- [28] M. Molloy and B. Reed, *Random Struct. Algor.* **6**, 161 (1995).
- [29] M. E. J. Newman, S. H. Strogatz, and D. J. Watts, *Phys. Rev. E* **64**, 026118 (2001).
- [30] M. Hoore and S. Moghimi-Araghi, *J. Phys. A* **46**, 195001 (2013).
- [31] P.-A. Noël, C. D. Brummitt, and R. M. D'Souza, *Phys. Rev. Lett.* **111**, 078701 (2013), Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.111.078701>.
- [32] K.-I. Goh, B. Kahng, and D. Kim, *Phys. Rev. Lett.* **87**, 278701 (2001).
- [33] K. Christensen and Z. Olami, *Phys. Rev. E* **48**, 3361 (1993).
- [34] M. E. J. Newman, *Phys. Rev. E* **66**, 016128 (2002).