Suppressing cascades in a self-organized-critical model with non-contiguous spread of failures

Heiko Hoffmann and David W Payton
HRL Laboratories, LLC, Malibu, CA 90265, USA

Abstract

Many complex systems that produce cascading events are thought to be self-organized critical (SOC). So far, models of SOC treat a cascade as a spread strictly between adjacent nodes, while in many real systems, e.g., the power-grid or the brain, this restriction is invalid. Here, we demonstrate for the first time SOC behavior in a model for which the spread is non-contiguous, i.e., not restricted to neighboring nodes. We illustrate our results in a circuit model obeying Kirchhoff's laws and demonstrate mitigation strategies that avoid large-scale cascades. We found that the following two unconventional strategies break SOC: (1) upgrade lines at random in addition to fixing failures and (2) upgrade a tripped line with one that has a random trip threshold. These results enhance our understanding about the conditions under which SOC can occur and may lead to insights that help avoid catastrophic events in real-world systems.

Keywords: Self-Organized Criticality, Cascading Failures, Transmission Network, Self-Organized Criticality Control

1. Introduction

Cascading events are ubiquitous in nature, e.g., they are found in power grids failures, disease propagation, forest fires, brain seizures, and earthquakes. The size of such cascading events was found in many systems to be approximately distributed according to a power law [1, 2, 3, 4]. Self-organized criticality (SOC) has been suggested as a possible mechanism that produces such a distribution of event sizes [5]. Here, self-organization implies that even without parameter tuning, a system becomes critical, i.e., produces scale-free event sizes, which are power-law distributed. However, SOC has been described more as a phenomenon [6], and the criteria that lead to SOC are still unknown.

So far, most model systems that study self-organized criticality have been limited to a local-spread of a cascade, i.e., the propagation of a cascade happens contiguously between neighboring or directly-connected nodes of the system. However, in real systems this restriction can be invalid. For example, in the power-grid, voltages and currents redistribute very quickly after a line trips (in practice, voltage ripples travel at more than 1000 miles per second through the grid [7]). Thus, it is possible that lines trip that are not directly connected to a previously tripped line. Brummitt et al. [8] called this behavior a non-local spread of a failure and criticized that many stylized models [9, 10] of power-grids assume that failures spread locally through adjacent nodes and therefore provide little insight into power-grid failures.

Self-organized criticality tends to be robust to perturbations, and thus, SOC systems are challenging to control in a way that alters the distribution of event sizes [9]. Controlling or modifying a system such that event sizes are not power-law distributed and instead decay faster (e.g., exponentially) could decrease the frequency of large events by several orders of magnitude. Thus, these modifications are beneficial when the cost of a large event is high. For example, for the power-grid, the cost of an outage depends non-linearly on its size: smaller outages can be compensated by backup systems, while large events lead to lost revenue, spoiled goods, and other cost in addition to repairs—the 2003 Northeast Blackout was estimated to reduce US earnings by $6.4 Billion [11].

In this article, we demonstrate for the first time evidence for self-organized criticality in a model of non-contiguous cascading failures. In our transmission-network model, voltages in the network instantaneously redistribute according to Kirchhoff's laws. Moreover, we contribute successful strategies for suppressing power-law-distributed event sizes.

In recent work on controlling an SOC system, the authors used a chip-firing model [9]. Such models are crafted after the original Bak-Tang-Wiesenfeld sandpile model [5], where chips are added to random nodes, and a node redistributes its chips to its neighbors after the number of chips at the node reaches a threshold. Here, we present a model that is closer to a physically-real system. However, we do not attempt to model accurately real power-grids. Instead, we aim for the optimal trade-off between model realism and simplicity, which captures the essence of non-contiguous cascades, while providing a model suitable for exploration. Our main motivation is to shed light on the conditions leading to self-organized criticality. The model allows us to limit the dependence on parameters, to visualize the spread of a cascade, and to study the impact of the grid's graph structure.

2. The Model System

Our model consists of a network of power lines connecting nodes, direct-current generators at each node, and loads connecting each node to ground (Fig. 1). Unlike real power-grids, which use alternating currents, to simplify our model, we focus only on direct-current circuits. Our current generators vary in strength, with values S_i randomly distributed uniformly in the interval $[0; I_0]$ (here, $I_0 = 1$)—the exact probability distribution of source sizes appeared to have little effect on our results.

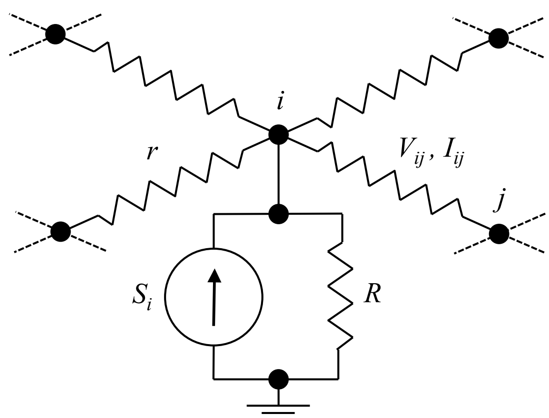

Figure 1: A circuit model. Each node i is connected to a current source S_i and a load resistor R. The grid lines have a resistance of r.

For the grid network, we assume an arbitrary graph of n nodes. The graph is defined given a symmetric connectivity matrix $\{M_{ij}\}$, where $M_{ij} = 1$ if node i is connected to node j, and $M_{ij} = 0$ otherwise. If a line trips, the corresponding elements in M are set to zero.

For each node i, the source current S_i gets distributed to the local load R and to the neighboring nodes connected through lines with resistance r.

Thus, according to Kirchhoff's laws,

$$S_i - \frac{V_i}{R} - \sum_{j \in N_i} \frac{V_i - V_j}{r} = 0\,, \tag{1}$$

where V_i is the voltage and N_i the set of neighbors of node i. Given the connectivity matrix M of the graph, we can write the above equations as follows,

$$\left(\frac{1}{R} + \frac{k_i}{r}\right) V_i - \sum_j \frac{M_{ij}}{r} V_j = S_i \,, \tag{2}$$

where k_i is the degree of node i, $k_i = \sum_j M_{ij}$. We can solve this set of linear equations for the voltage distribution V_i (we used a sparse linear solver [12]). The solution to (2) depends on the ratio of r/R, which we set to 0.001 ($r = 0.01$ and $R = 10$) unless otherwise noted, i.e., the resistance in the transmission lines is low compared to the load at each node. Apart from the graph structure, this ratio is the only parameter on which our model depends.

Given V_i, we evaluate the current for each line

$$I_{ij} = \frac{|V_i - V_j|}{r} \tag{3}$$

and compute the maximum normalized current above threshold T_{ij},

$$\Delta I_{\max} = \max_{ij} \left(\frac{I_{ij} - T_{ij}}{T_{ij}}\right) . \tag{4}$$

If $\Delta I_{\max} \geq 0$, the corresponding line trips, and we re-compute the voltage distribution V_i. In our model, the redistribution of voltages happens instantaneously after each line failure. We call the consecutive failure of lines until $\Delta I_{\max} < 0$ an avalanche.

At the beginning of a simulation run, all T_{ij} are set to T_0 (here, $T_0 = 10$). To put stress on our network and bring lines to their limit, we let all T_{ij} decay exponentially until a line breaks and between avalanches. This decay is a slow process. After the end of each avalanche, we immediately repair all broken lines (reset $M_{ij} = 1$ and $T_{ij} = T_0$). A related but opposite mechanism, a cycle of gradual increase in demand and line upgrade, has been suggested for a different kind of power-transmission model [13].

In our simulation, we used an outer and inner iteration loop to simulate the slow and fast processes. The outer loop iterates over avalanches ($500,000$ times), and the inner loop computes the cascade of an avalanche, iteratively solving (2) and removing lines until all currents are below threshold. In each iteration step t of the outer loop, we set $T_{ij}^{t+1} = \lambda \, T_{ij}^{t}$ for all lines and compute λ such that the next line that would break is exactly at threshold to simulate an infinitely-gentle decay, $\lambda = 1 + \Delta I_{\max}$. Such an infinitely-gentle drive has been also used in a version [6] of the Olami-Feder-Christensen earthquake model [14]. In our model, using an infinitely-gentle decay turned out to be important; otherwise, a discrete step (i.e., a fixed λ value) resulted in a characteristic length, which would break SOC.

This updated rule for T_{ij} furthermore makes the absolute values of r and I_0 irrelevant. However, we need to record λ, since the time interval between avalanches is proportional to $\Delta t = \ln(\lambda)$, and we evaluate the cost of repair per unit time, as described below.

We did our experiments with three different types of graphs: square and triangular lattices with $L \times L$ nodes and periodic boundary conditions (i.e., toroidal graphs) and a ladder graph with ends joined to form a ring. We chose periodic boundary conditions because we wanted all nodes to be geometrically indistinguishable to reduce limit-size effects. For our analyses, we skipped the first 2,000 avalanches (unless otherwise noted) to give the system sufficient time to reach a potential critical state (see Section 3.2).

3. Experiments and Analyses

In the following, we will demonstrate 1) the non-contiguous spread of cascading failures, 2) evidence for SOC behavior, 3) mitigation strategies to suppress power-law distributed avalanches, and 4) the impact of the graph structure.

3.1. Non-contiguous cascades

We found that indeed in our model an avalanche forms several distinct non-adjacent clusters of tripped lines (Fig 2). This behavior is possible since the voltage redistribution is faster than the cascading failure of lines. To distinguish a contiguous spread from a non-contiguous spread, we defined a spread to be contiguous if for each line failure that line was part of the same contiguous cluster of previously tripped lines. For a square lattice of size $L = 50$, we used this method to compute the probability of a contiguous spread as a function of avalanche size. As a result, this probability fell sharply with increasing size: p=0.54 for an avalanche of 3 lines and p=0 for more than 65 tripped lines; for the last case, not one out of 18,134 avalanches was a contiguous spread.

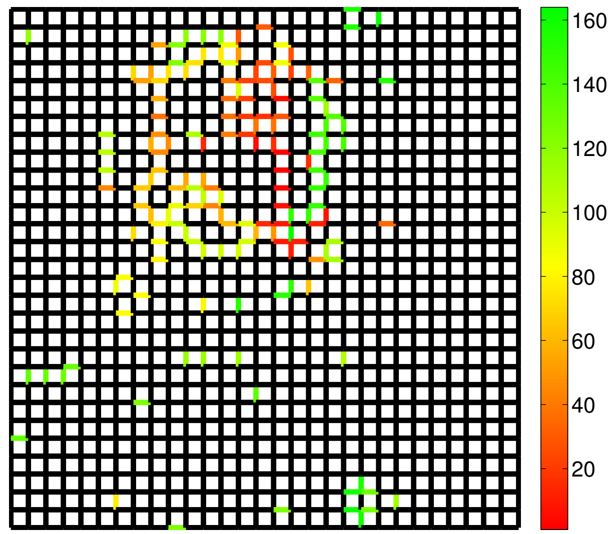

Figure 2: Sample of non-contiguous spread of an avalanche in a square lattice, $L = 30$, with periodic boundary condition. The color code depicts the order in which the lines tripped, with the first line to trip coded in red.

3.2. Evidence for self-organized criticality

Since a universally agreed definition of SOC has been missing [6] and the term "SOC" has been overused, we define for clarification a system to be SOC if it fulfills the following conditions

1. A complex system of multiple spatially separated components, where an event, like an avalanche, spans across a multitude of components.
2. The system reaches a critical state, defined by the size of events being free of scale and the average size approaching infinity.
3. Starting from a non-critical initial condition, the system self-organizes to the critical state without parameter tuning.

This definition is consistent with many models considered to be SOC [6]. The purpose of point 1 is to exclude some 1/f noise processes [15]. In our system, avalanches are spatially distributed.

Next, we demonstrate evidence that indicates that our system approached a critical state. On the square and triangular lattices, the probability of avalanche sizes followed a power law (Fig. 3). Thus, the sizes are free of scale (within the limits of the system size). For different L, the size distributions overlapped, while the cut-off from the power law increased with system size, which is an important condition for criticality.

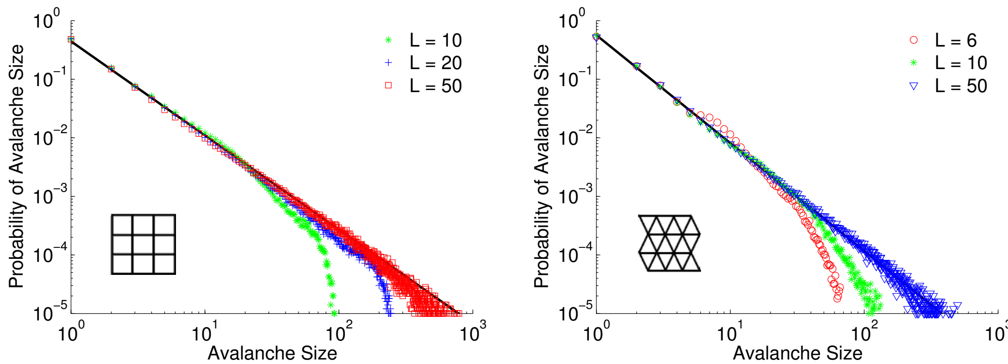

Figure 3: Scale-free behavior of avalanches in our circuit model for the square (Left) and triangular lattices (Right). Power laws were fitted to the data for $L = 50$.

Interestingly, on the triangular lattice, the slope of the power law was different. For the square lattice, the average slope was $\tau = -1.61 \pm 0.03$, and for the triangular lattice, the slope was $\tau = -1.88 \pm 0.03$ (mean$\pm$ SD, 4 simulations each with different random initialization, using $L = 50$). In comparison, simulation results on sand-pile models indicate that the slope is the same on square and triangular lattices [16]. In our case, having non-contiguous avalanches might make the difference. Consider the graph that links sites that could fail in a sequence. In contrast to the 2D sand-pile model where this graph is planar (e.g., square or triangular lattice), a non-contiguous spread creates cross links that may make the graph non-planar, which may result in a different slope. This hypothesis is consistent with the observation that a 3D sand-pile model has a different slope from a 2D model [5].

In our model, the slope for both lattice types was larger than -2, which implies that the average size of an avalanche diverges with increasing system size. The lack of a characteristic size is a criteria for criticality.

Finally, we demonstrate that our system started in a non-critical state and converged on its own to the critical state (Fig. 4). Thus, our system self-organized to criticality. During this convergence, the system oscillated between over-critical and sub-critical before settling into the critical state. So far, this convergence to criticality has been underreported. In contrast to our observation, previous work showed an overdamped convergence from sub-critical to critical [17]. In summary, our observations indicate that our system is self-organized critical according to the above definition.

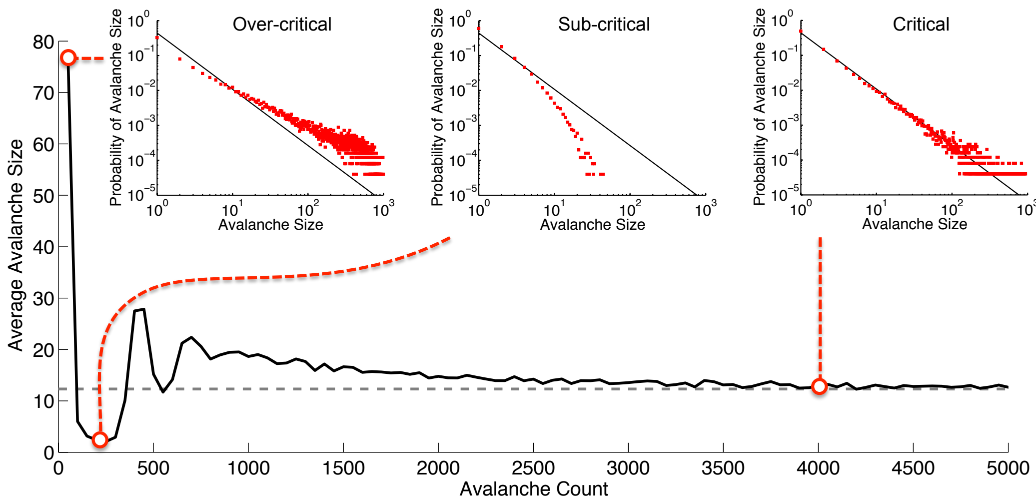

Figure 4: Convergence to criticality. The results are shown for the square lattice with $L = 50$. Data are averaged over time windows of 50 avalanches and over 500 separate simulation runs with different random initializations of S_i. The dashed line is the average avalanche size for all avalanches from count 8000 to 10000. Three sample avalanche-size distributions are shown, one over-critical that was obtained from avalanche count 1 to 50, one sub-critical from count 176 to 225, and one critical from count 3951 to 4000. The solid line in these graphs shows for comparison a power-law with slope $\tau = -1.61$.

In addition, we tested the dependence on the model parameter r/R. This parameter resulted also in a cut-off, which disappeared for $r/R \rightarrow 0$ (Fig. 5). Therefore, strictly speaking, our system could be self-organized critical only in the limit $r/R \rightarrow 0$. In a transmission network, this limit corresponds to negligible link resistances compared to load resistances.

3.3. Mitigation strategies to suppress power-laws

We found two strategies that suppressed a power-law distribution of avalanche sizes in our model. In our first strategy, we carried out additional (unnecessary) repairs after each avalanche. Apart from repairing broken lines, we chose a additional lines at random and upgraded their thresholds to T_0 independent of their value. As a result, the distribution of avalanche sizes moved away from the power law and the effect increased with increasing a (Fig. 6). Large avalanches became less likely, e.g., for a size of 100

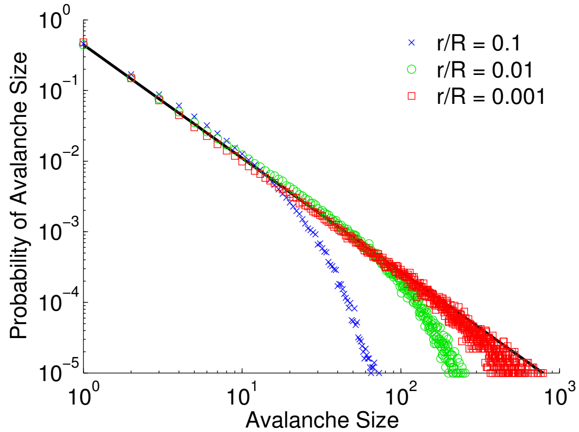

Figure 5: The relative resistance of the transmission lines, r/R, results in a finite size effect. The results are shown on a square lattice with $L = 50$.

lines, the probability dropped by more than 20x for the square lattice, $a = 2$, and $L = 50$. The reduction of large avalanches happened shortly after the beginning of a simulation run and not just after the system reached SOC.

We hypothesize the following mechanism to be responsible for the observed suppression of power laws. In the unperturbed case ($a = 0$), clusters form at the critical state such that a failing line can trigger the cascading failure of a whole cluster. When we upgrade a line at random, there is a probability $p \geq \epsilon > 0$ that a line in an otherwise critical cluster will not fail, where ϵ is the probability that the line just got upgraded. Thus, the probability that a whole cluster of m lines fails is upper-bounded by $(1 - \epsilon)^m$. This exponential decay as a function of cluster size leads to a characteristic length and prohibits a scale-free distribution of avalanche sizes, and thus, no power law. The value of ϵ increases with increasing a leading to a faster decay, which matches the observed behavior.

We evaluated if our repair strategy would actually reduce the total financial cost. To estimate the cost, we start with a lower cost bound that is linear in the total number of restored lines, i.e., $x + a$ per avalanche, if x is the size of the avalanche. However, the total cost likely increases faster than linear due to the extra cost of large blackouts. Here, we do not aim to have the exact relationship, but rather want to illustrate the impact of the nonlinearity, and therefore, we use the same relationship, x^α, as recently used in a related article [9]. In addition, we need to take into account differences in avalanche frequency, ν. Thus, our total cost, C, per unit time is

$$C = (x^\alpha + a)\nu\,, \tag{5}$$

where we set ν to the inverse of the average value of Δt during one simulation run (see above). As a result, if the non-linear increase was sufficiently large, we observed a cost benefit of our repair strategy (Fig. 7).

Specifically, we observed a trade-off between the cost of additional upgrades and their benefit, and the optimal number a depends on the form of the non-linearity (Noël et al. observed a similar behavior for their control strategy [9]).

In our second mitigation strategy, we upgraded each failed line to a random threshold instead of T_0. The random threshold was drawn uniformly from the interval $[T_0 - \mu/2; T_0 + \mu/2]$. This strategy also suppressed the power-law distribution (Fig. 8). For the square lattice with $L = 50$, the decay in probability for large avalanche sizes increased with increasing noise range μ and reached a plateau for about $\mu = 1.5T_0$. When computing the same cost, C, as above, the cost per unit time dropped by 75% (here, $a = 0$ and $\alpha = 1.5$).

The random update causes some lines to have a low threshold and thus fail prematurely. Apparently, this failure counterintuitively has the positive effect of reducing large-scale avalanches. It likely prevents the build up of a large cluster of lines that are close to critical threshold. In addition, lines that randomly get a high threshold can act as buffers to prevent a large-scale cascade.

3.4. Impact of the graph structure

On the triangular and square lattices, we observed the same effect of our mitigation strategies. Thus, our results are applicable beyond a specific kind of graph.

In addition, network graphs exist that fail to show SOC behavior in the first place: For example, on a ladder graph, we found that the probability of an avalanche decreases exponentially with its size (Fig. 9).

Here, the spread of a cascade is restricted mainly to one dimension. We hypothesize that the single dimension restricts the possible paths a cascade can travel. For any specific path, the probability of a cascade traveling along that path decays exponentially with path length. Thus, for an avalanche to be scale-free, this exponential decay has to be compensated by an exponential increase in the number of possible paths. Apparently, for the ladder graph, the number of possible paths does not increase exponentially despite the non-contiguous spread of a failure. Since the spread is non-contiguous, however, we lack a straightforward link between graph structure and the number of possible paths. The mechanism behind the impact of the graph structure on the occurrence of SOC and the slope of a power law is left for future investigation.

4. Conclusions

To conclude, we have demonstrated a circuit model obeying Kirchhoff's laws, which produced cascading failures that spread non-contiguously. To our knowledge, this model is the first with this property for which evidence for self-organized criticality has been shown. We demonstrated that our model converged to a critical state while oscillating between sub-critical and over-critical states. Moreover, the occurrence of SOC depended on the model's graph structure, which also impacted the slope of the power law. Using this model, we have shown two strategies that suppressed power-law distributed cascades: the first strategy restores lines at random to their original state independent of their condition. This strategy can be thought of as a form of preventive maintenance. The second strategy restores a failed line to a random trip threshold. We hypothesize that both strategies introduce an exponential decay in probability for the propagation of an avalanche. In contrast, criticality apparently requires the formation of large clusters of lines for which the probability of avalanche propagation decays slower than exponential. By preventing this formation, a system remains sub-critical. Such insights might be applicable to prevent catastrophic events in real-world systems that exhibit SOC behavior.

5. Acknowledgments

This work was funded by HRL Laboratories, LLC.

[1] I. Dobson, B. A. Carreras, V. E. Lynch, D. E. Newman, Complex systems analysis of series of blackouts: cascading failure, critical points, and self-organization, Chaos 17 (2007) 026103.

[2] B. D. Malamud, G. Morein, D. L. Turcotte, Forest fires: An example of self-organized critical behavior, Science 281 (5384) (1998) 1840–1842.

[3] G. A. Worrell, S. D. Cranstoun, J. Echauz, B. Litt, Evidence for self-organized criticality in human epileptic hippocampus, NeuroReport 13 (16) (2002) 2017–2021.

[4] B. Gutenberg, C. F. Richter, Earthquake magnitude, intensity, energy, and acceleration, Bulletin of the Seismological Society of America 46 (2) (1956) 105–145.

[5] P. Bak, C. Tang, K. Wiesenfeld, Self-organized criticality: An explanation of the 1/f noise, Physical Review Letters 59 (1987) 381–384.

[6] H. J. Jensen, Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems, Cambridge Lecture Notes in Physics, Cambridge University Press, 1998.

[7] S.-J. S. Tsai, L. Zhang, A. G. Phadke, Y. Liu, M. R. Ingram, S. C. Bell, I. S. Grant, D. T. Bradshaw, D. Lubkeman, L. Tang, Study of global frequency dynamic behavior of large power systems, in: Power Systems Conference and Exposition, Vol. 1, IEEE, 2004, pp. 328–335.

[8] C. D. Brummitt, P. D. H. Hines, I. Dobson, C. Moore, R. M. D'Souza, Transdisciplinary electric power grid science, Proceedings of the National Academy of Sciences of the USA 110 (2013) 12159.

[9] P.-A. Noël, C. D. Brummitt, R. M. D'Souza, Controlling self-organizing dynamics on networks using models that self-organize, Physical Review Letters 111 (2013) 078701.

[10] D. P. Chassin, C. Posse, Evaluating North American electric grid reliability using the Barabási Albert network model, Physica A 355 (2005) 667–677.

[11] P. L. Anderson, I. K. Geckil, Northeast blackout likely to reduce US earnings by $6.4 billion, Anderson Economic Group Working Paper 2.

[12] T. A. Davis, UMFPACK: unsymmetric multifrontal sparse LU factorization package, http://www.cise.ufl.edu/research/sparse/umfpack/.

[13] B. A. Carreras, V. E. Lynch, I. Dobson, D. E. Newman, Dynamics, criticality and self-organization in a model for blackouts in power transmission systems, in: Hawaii International Conference on System Sciences, IEEE, 2002.

[14] Z. Olami, H. J. S. Feder, K. Christensen, Self-organized criticality in a continuous, nonconservative cellular automaton modeling earthquakes, Physical Review Letters 68 (1992) 1244–1247.

[15] K. P. O'Brien, M. B. Weissman, Statistical signatures of self-organization, Physical Review A 46 (8) (1992) 4475–4478.

[16] S. S. Manna, Critical exponents of the sand pile models in two dimensions, Physica A 179 (1991) 249–268.

[17] C. M. Aegerter, K. A. Lörincz, M. S. Welling, R. J. Wijngaarden, Extremal dynamics and the approach to the critical state: Experiments on a three dimensional pile of rice, Physical Review Letters 92 (5) (2004) 058702.

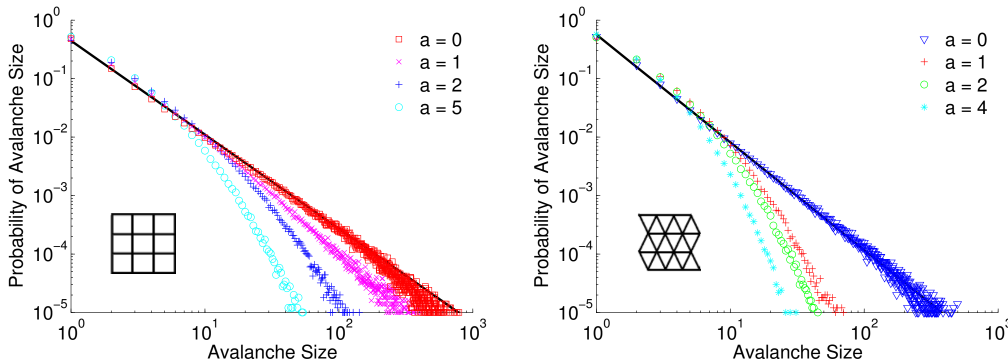

Figure 6: Restoring a additional lines at random locations at each iteration step (after each avalanche) breaks the SOC behavior. The results are shown for the square (Left) and triangular (Right) lattices with $L = 50$.

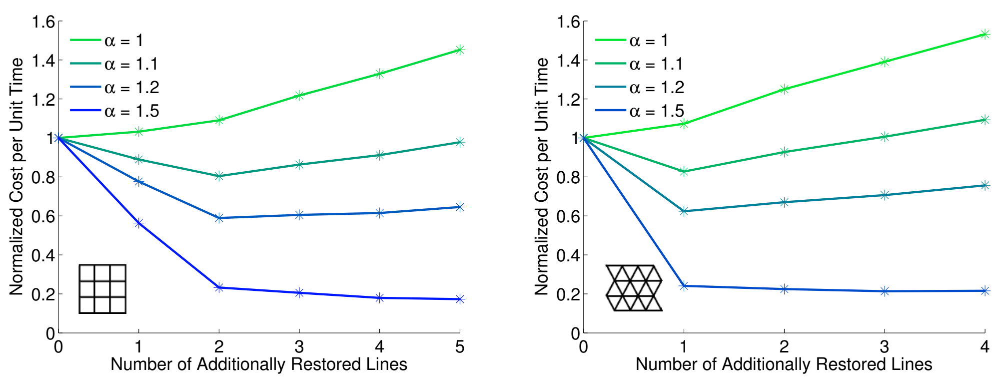

Figure 7: Cost per unit time (normalized to 1 for $a = 0$) as a function of the number of additionally restored lines for the square (Left) and triangular (Right) lattices with $L = 50$. The cost depends non-linearly on the avalanche size, x, while α controls the non-linearity, x^{α}.

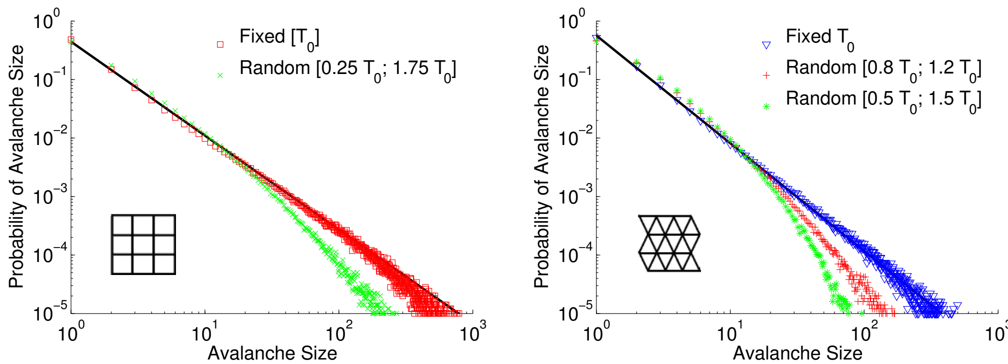

Figure 8: The power law breaks when updating each line after each avalanche to a random threshold (uniformly in an interval around T_0). The results are shown for the square (Left) and triangular (Right) lattices with $L = 50$.

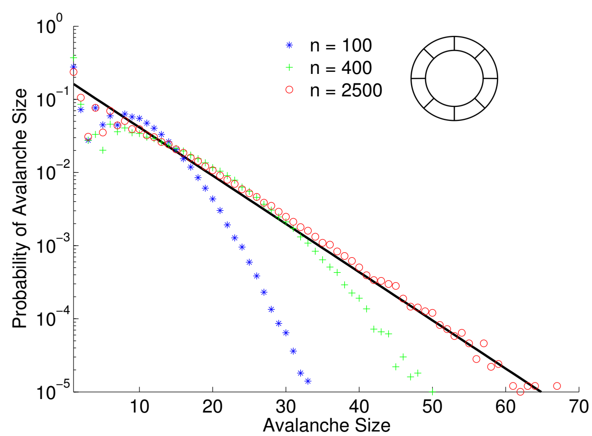

Figure 9: On a ladder graph with joined ends, the avalanches decay exponentially (for $n = 2500$, the black line shows a linear fit to the semi-log plot giving an exponent of -0.152 ± 0.002).