

Distribution of blackouts in the power grid and the Motter and Lai model

Yosef Kornbluth¹, Gabriel Cwilich², Sergey V. Buldyrev², Saleh Soltan^{3,*}, and Gil Zussman⁴

¹*Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA*

²*Department of Physics, Yeshiva University, New York, New York, USA*

³*Princeton University, Princeton, New Jersey, USA*

⁴*Department of Electrical Engineering, Columbia University, New York, New York, USA*



(Received 3 August 2020; accepted 9 February 2021; published 17 March 2021)

Carreras, Dobson, and colleagues have studied empirical data on the sizes of the blackouts in real grids and modeled them with computer simulations using the direct current approximation. They have found that the resulting blackout sizes are distributed as a power law and suggested that this is because the grids are driven to the self-organized critical state. In contrast, more recent studies found that the distribution of cascades is bimodal resulting in either a very small blackout or a very large blackout, engulfing a finite fraction of the system. Here we reconcile the two approaches and investigate how the distribution of the blackouts changes with model parameters, including the tolerance criteria and the dynamic rules of failure of the overloaded lines during the cascade. In addition, we study the same problem for the Motter and Lai model and find similar results, suggesting that the physical laws of flow on the network are not as important as network topology, overload conditions, and dynamic rules of failure.

DOI: [10.1103/PhysRevE.103.032309](https://doi.org/10.1103/PhysRevE.103.032309)

I. INTRODUCTION

Cascading failures in the power grids continue to happen in spite of efforts to make power grids more resilient [1–3]. The standard criterion of resiliency is the $N - 1$ criterion [4]: the grid must safely operate in the event of the failure of any single line. Carreras *et al.* have studied empirical data on the sizes of the blackouts in real grids [5] and modeled them with computer simulations using the direct current (dc) approximation [6–8]. They have found that the resulting blackout sizes are distributed as a power law and suggested that this is because the grids are driven to the self-organized critical (SOC) state [9–11]. In their model, they assume that at any stage of the cascade, one of the lines with loads exceeding the maximum values imposed by the $N - 1$ condition fails and immediately all the currents in the grids are redistributed adjusting to the new network topology. The motivation for this “one-by-one” failure rule of the cascade propagation comes from investigation of real blackouts. It is documented [1,12] that the failures of overloaded lines do not happen instantaneously but require a certain period of time, during which overloaded lines undergo heating expansion. When the expanded line touches the ground or foliage, the current in the line dramatically increases and the line breaks, after which the current in other lines changes almost instantaneously, so that if the current in a previously overloaded line reduces to normal, that line’s length reduces and it may eventually survive the cascade of failures. In a recent work [13] which models U.S.-south Canada power grid with the dc approximation and starts the cascade with the removal of $n_r \leq 4$ power lines,

the overheating has been modeled directly using a realistic model of temperature evolution [14]. This work reproduces the power-law distribution of blackouts, but more importantly finds a small set of vulnerable nodes responsible for the failure cascades.

Ren *et al.* [4] suggested that the power grid is driven to the SOC state by recursive upgrading of the power grid with constantly growing power demand, applying the $N - 1$ condition, or upgrading lines involved in the recent cascade of failures. They simulated this model of self-organization and found that the system converges in the infinite time limit to a steady state characterized by an exponential distribution of the blackout sizes, while the SOC models typically display power-law distribution of avalanche lengths associated with a second-order critical point as in percolation theory [15–17]. In contrast, recently [18] a power-law distribution of the blackouts has been related to Zipf’s law distribution of city sizes [19].

Other recent studies [20,21] have suggested that the distribution of cascades is bimodal, resulting in either a very small blackout or a very large one. In all these studies, the cascades of failures were started by a random failure of a single line, and the currents were computed from a given distribution of loads and generators using the dc current approximation. The difference between Refs. [6–8] and Ref. [20] was that the maximal loads in Ref. [20] were computed not by using $N - 1$ criterion, but by a uniform tolerance algorithm as in Refs. [22,23] and that the overloaded lines during each stage of the cascade were eliminated all at once [15,16,24], not one by one.

The difference in the outcome between the “one-by-one” rule and the “all-at-once” rule suggests that the cascade propagation in the overload models inherently depends on the

*Now at Amazon.com, New York, New York, USA.

dynamics of the cascade, which makes these models very different from the simple topological models of cascading failures in which the final outcome depends on only the network topology and initial failures and is the same for any sequence of removals of the unfunctional elements [25].

Here we reconcile the two approaches and show that the power-law distribution of cascades emerges for high protection level of the grid and also in cases where the network topology is close to the percolation point. We also show that the “one-by-one” removal rule significantly reduces the sizes of large blackouts and their probabilities, replacing the bimodal distribution of blackouts by an approximate power law for intermediate protection levels when the “all-at-once” rule still leads to a bimodal distribution. We show that these features are held for both the dc model of a power grid and the much simpler Motter and Lai model [22,26], suggesting that the exact physical laws of flow on the network (Kirchhoff’s laws versus minimal path rule) are not as relevant as the protection level, network topology, and cascade dynamics rules.

The rest of the paper is organized as follows: In Sec. II we first review phase transitions for cascading failures in topological and overload models with massive and single-element initial attacks. Next we interpret the cascading failure after a single-element initial attack as a branching process.

In Sec. III we first describe the models of power grids we simulate. Then we describe the implementation of the $N - 1$ condition, and uniform tolerance conditions for node failure and the one-by-one versus all-at-once cascade update rules for power grids. Next we investigate the effect of the $N - 1$ condition with the one-by-one and all-at-once rules. Finally, we investigate uniform tolerance condition with both update rules.

In Sec. IV we first describe the implementation of the betweenness centrality model and the implementation of the $N - 1$ condition and uniform tolerance conditions for node failure. Next we study the uniform tolerance condition with various tolerance levels and both cascade update rules. After this we discuss the effect of the average degree and the size of the system on the cascade distribution for the $N - 1$ condition and one-by-one update rule. Finally, we investigate such a cascade as a SOC dynamics.

In Sec. V we present the summary and conclusions.

II. POWER-LAW VERSUS BIMODAL DISTRIBUTIONS

A. Overview

In this section we review the conditions under which the distributions of cascade sizes in overload models are power law or bimodal, and compare these overload models with the topological models [27–34], in which the outcome of the cascade depends only on the initial topology of the network and location of the initially damaged elements. Topological or overload models can be investigated for three different classes of initial attacks. (The third class of initial attacks, characterized by a hybrid transition, [35,36] will be briefly reviewed in the Appendix).

In the first well-studied class [31,32,35–37], the cascade is started by a massive attack, after which only a fraction p of

elements survive. For this class, an interesting problem is to study the behavior of the order parameter $S(p)$, which is the fraction of the surviving nodes at the end of the cascade, as a function of the fraction p of the initially surviving elements. Many topological and overload models exhibit a first-order phase transition at $p = p_t$, at which the order parameter $S(p)$ exhibits a step discontinuity, or a second-order transition at $p = p_c$, at which the order parameter is continuous but its derivative with respect to p exhibits a step discontinuity. For finite systems, near $p = p_t$ the distribution of the order parameter becomes bimodal, and the transition point p_t is defined as the point at which the population of the two peaks is equal. At the second-order transition the distribution of the order parameter is always unimodal. For some topological and overload models, the line of the first-order transitions, p_t , in a plane of two parameters may transform into the line of second-order phase transitions, p_c , at a certain value of the second parameter [32,36–38].

B. Cascading failures as branching process

In the second class of initial attacks, such as in the above mentioned discussion of power grids [6–8,13,18] and a recent work on the Motter and Lai model [26], the cascade is started by the failure of a single element. Here it is informative to study the distribution of the “blackout” sizes, i.e., the number of nodes that failed during the cascade and how such a probability distribution depends on the model parameters. In this case the system before the attack exists in a state characterized by some set of parameters, but the parameter p describing the size of the initial attack is not applicable at all. In this case, in both topological and overload models, the cascading failures evolve as a branching process in which the failure of one element leads to the failure of b elements, which depend on the failed element in the topological models, or get overloaded due to the removal of the failed element in the overload models. The critical point of the simplest branching process with a fixed distribution of the branching factor $P(b)$ is defined by $\langle b \rangle = 1$ [39]. For the mean-field variant of the fixed distribution model, the distribution of the avalanche sizes is a power law with an exponential cutoff $P(s) \sim s^{-\tau} \exp(-s/s^*)$, where $s^* \sim (\langle b \rangle - 1)^{-1/\sigma}$. Here $\sigma = 1/2$ and $\tau = 3/2$ are the critical exponents [17,40]. These in turn give the values of the two other critical exponents, $\gamma = (2 - \tau)/\sigma = 1$ and $\beta = (\tau - 1)/\sigma = 1$, where γ governs the divergence of the average finite avalanche size for $\langle b \rangle \rightarrow 1$ and β governs the probability of a giant avalanche for $\langle b \rangle > 1$, i.e., the avalanche which in the thermodynamic limit $N \rightarrow \infty$ constitutes a finite fraction μ of the entire system. Note that μ plays the role of the order parameter of the phase transition. This phase transition is the second-order transition as in percolation. For $\langle b \rangle > 1$, $\mu > 0$, while for $\langle b \rangle \leq 1$, $\mu = 0$. In a finite network, the size of the giant avalanche fluctuates around μN and, thus, forms the second peak of the bimodal distribution of avalanche sizes. Other avalanches, which do not scale as a finite fraction of the network, form a first peak of the size distributions, which is separated from the second peak by a huge gap because there are no avalanches of size s , such that $s^* \ll s \ll \mu N$ (Fig. 1). Thus, this simple branching process model predicts the existence of regimes with bimodal ($\langle b \rangle >$




FIG. 1. Cumulative distributions of the avalanche sizes in the failure branching process with the Poisson branching degree distributions, with the average branching factors $\langle b \rangle = 0.9, 0.95, 1, 1.05, 1.1$ for finite system of $N = 100\,000$ and $N = 200\,000$ nodes. One can see that the bimodal distribution of avalanches, characterized by a wide horizontal plateau of the cumulative distribution, emerges for $\langle b \rangle > 1$. For $\langle b \rangle < 1$, the distribution has an exponential cutoff. Exactly at the critical point $\langle b \rangle = 1$, the distribution is a power law with a finite-system cutoff. The power-law region with the exponent $\tau - 1 = 0.5$ increases with the size of the system.

$1, \mu > 0$) and power-law distributions with exponential cutoff ($\langle b \rangle \leq 1, \mu = 0$), depending on the parameters of the model. Carreras *et al.* conjecture [5] that the power grid is a SOC system, which somehow drives itself to the critical point of this branching process, at which the distribution of blackouts is a power law. It is not clear how this SOC process works. It may be created by the long-term effect of the application of the $N - 1$ criterion [4] or, more likely, by some empirical compromise between safety and price of keeping $\langle b \rangle$ small: electric companies want to eliminate large catastrophic blackouts which exist for $\langle b \rangle > 1$, but, trying to save money on the infrastructure, are willing to accept finite blackouts. Obviously this compromise is achieved for $\langle b \rangle = 1 - \epsilon$. In the following sections we will check this failure-branching process model with computer simulations.

III. SIMULATIONS OF THE POWER GRIDS

A. Models of power grids

In order to verify that the $N - 1$ condition, in addition to the one-by-one failure rule, is responsible for the emergence of the power-law distributions of the blackouts, we perform simulations of several grid topologies including a simplified model of the U.S. Western Interconnect (USWI) [41], Learning Based Synthetic Power Grid (LBSPG) [42,43], Degree and Distance attachment model (DADA) [20], and two models of artificial topologies: a random 2D graph with given maximal line length (RML) and random regular 2D graph with given degree of a node (RR). The method of solving Kirchhoff's equations for the dc approximation is given in Ref. [20]. In all models there are n_p generator nodes which produce power, n_c loads which consume power, and n_t trans-

mitter nodes, which do not produce or consume power but redistribute the flow between several power lines. Within the dc approximation we assume that the power transmitted by a line connecting nodes i and j is simply its current I_{ij} . Each generator i is characterized by the power (current) $I_i^+ > 0$, while each load is characterized by the power (current) it consumes, $I_i^- < 0$. For transmitter nodes, $I_i^+ = I_i^- = 0$. In the USWI, LBSPG, and DADA models, the n_p , n_c , and n_t as well as their powers $I_i^\pm$ are given in Refs. [20,42,43]. In RML and RR models we assume that all generators or loads produce or consume equal power $I = I_i^+ = I_i^-$ and that $n_p = n_c = 1000$ and $n_t = 8000$. The power balance condition $\sum_i I_i^+ + \sum_i I_i^- = 0$ is satisfied in all models. The resistances of lines R_{ij} are proportional to their lengths. In the RML and RR models all $N = n_p + n_c + n_t$ nodes are randomly placed on a square of edge L and periodic boundaries. In the RR model, each node is connected to its k nearest neighbors. In order to make sure that each line is connected to exactly k nodes, we make a list of all nearest neighbor pairs and start to select pairs from this list in ascending order of length, creating a link between them if both of the nodes in the pair have less than k links. At the end of the process, the majority of nodes have k links with a few exceptions which have only $k - 1$ links. In the RML model, each node is surrounded with a circle of radius $r = L\sqrt{\langle k \rangle / (\pi N)}$ and is connected to all the nodes within this circle. Since the nodes are randomly distributed geographically, the degree distribution of nodes becomes a Poisson distribution with average degree $\langle k \rangle$. We used $k = 4$ for RR and $\langle k \rangle = 5.0$ for RML. If the grid does not form a single connected cluster, we discard all the nodes which do not belong to the largest connected cluster.

B. Implementation of the two schemes of protection: $N - 1$ condition and uniform tolerance

To implement the $N - 1$ condition for all of our models, we proceed as follows: Assuming that the grid contains n_L links, we obtain the current in a given link ij between nodes i and j after the removal of each one of the other $(n_L - 1)$ links, and we find the maximal current I_{ij}^* through that link ij obtained over the set of those $(n_L - 1)$ currents. We use this value I_{ij}^* as the maximal possible load that the link ij may sustain.

We also perform simulations for networks with uniform tolerance protection in which $I_{ij}^* = (\alpha + 1)I_{ij}$, where I_{ij} is the current in line ij in the unperturbed network.

C. Implementation of the cascade update rules

Under the $N - 1$ condition case, we start each simulation with the removal of a random pair of links, $i_a j_a$ and $i_b j_b$, while in the case of uniform tolerance we remove just one random link $i_a j_a$.

Having done that, we recompute the currents in all remaining links, I_{ij}^1 . If in some links $I_{ij}^1 > I_{ij}^*$ we find the link with maximal overload I_{ij}^1 / I_{ij}^* , we remove it and find the new currents I_{ij}^2 (one-by-one update rule). To implement the all-at-once update rule, at each stage k of the cascade, we delete all overloaded links with $I_{ij}^k > I_{ij}^*$ and compute the new currents in the remaining links. For both rules, we continue this process until, at the n th stage of the cascades, for all re-

maintaining links, $I_{ij}^n \leq I_{ij}^*$. If, at a certain stage, the grid splits into several disconnected clusters, we apply the power production-consumption equalization for each cluster using the minimal production-consumption rule described in Ref. [20]. We compute the blackout size as a fraction of the consumed power lost in the cascade $|I_0 - I_f|/I_0$, where $I_0 = \sum_i |I_i^-|$ for the intact grid and $I_f = \sum_i |I_i^-|$ at the end of the cascade. Another method of measuring the blackout size is to compute the fraction of the lines lost in the cascade $f_d = n_d/n_L$, where n_d is the number of lines removed from the grid due to overload.

In summary, for each network model, we study four cases: the $N - 1$ condition together with one-by-one or all-at-once update rules, and uniform tolerance together with one-by-one or all-at-once update rules.

D. Results for different models of power grid with the $N - 1$ condition

First, we consider the $N - 1$ condition, along with the one-by-one rule. Implementing the $N - 1$ condition significantly improves the robustness of the grid, as compared to the uniform tolerance model with small α . This is not surprising because for all four models of the power grid the effective tolerance of line ij under the $N - 1$ rule, $\alpha_{ij} \equiv I_{ij}^*/I_{ij} - 1$ (where I_{ij} is the initial current in the line ij), has a wide distribution with almost 5% of lines having $\alpha_{ij} > 1$. Thus, the fraction of runs that result in any additional line overload is about 0.03; for most pairs of initially attacked lines, no further lines are overloaded. For the tolerance model with a uniform α , this level of protection is achieved only for $\alpha > 1$. Thus, implementation of the $N - 1$ condition and the one-by-one rule reduces the sizes of blackouts and creates many small cascades that end after a few failures [Figs. 2(a) and 2(b)]. For all four models of the power grid we observe the emergence of the approximate power-law part in the size distribution of the small cascades. However, for each model except RML with $\langle k \rangle = 5.0$ (which is close to the RML percolation threshold for $\langle k \rangle = 4.5$), in addition to small blackouts characterized by an approximate power-law distribution, we still find some large blackouts with a significant fraction of power loss. Note also that the power-law part of this distribution is still practically absent in the case for the USWI and LBSPG models. Thus, the entire distribution remains bimodal with two peaks for small and large cascades and practically no cascades of an intermediate size. This absence can be observed on the cumulative distribution graph as a large plateau [Fig. 2(a)]. The value of the exponent characterizing the power-law part of the cumulative distribution of small cascades varies for different models but remains in the range between 0 and 1. The presence of large cascades is indicated by the sharp drop of the cumulative distribution functions [Figs. 2(a) and 2(b)] at blackout fractions $x = \mu$, for all cases except the RML model. This sharp drop corresponds to the maximum of the probability density function at $x = \mu$, which can be defined as the order parameter of the system. Note that the value of μ is different for different grid models.

We next consider the $N - 1$ condition, but with the all-at-once update rule. The power-law part of the blackout distribution disappears [Fig. 2(c)]. The distribution becomes strictly bimodal, with the positive order parameter $\mu > 0$,




FIG. 2. Cumulative distributions of blackout sizes, measured as a function of the fraction of the consumed power lost in the blackout for different grid topologies described in this paper, with the dc approximation and the $N - 1$ condition implemented. (a) One-by-one update rule. All distributions remain bimodal except for the RML model with $\langle k \rangle = 5.0$. (b) The same as (a) in double logarithmic scale. A power-law distribution with $\tau - 1 = 0.5$ emerges for the case of the RML model. (c) All-at-once update rule. All distributions remain bimodal.

even for the case of RML model with $\langle k \rangle = 5.0$, for which the $N - 1$ condition together with the one-by-one rule produces a power-law distribution of blackouts. Additionally, large blackouts become much larger with the all-at-once rule than with

the one-by-one rule; the system is more prone to failure, which is indicated by the larger values of the order parameter μ .

E. Comparison of different cascade models for the USWI and the LBSPG

Next, we systematically explore how the $N - 1$ condition, uniform tolerance, and different dynamic rules affect the blackout distributions in the USWI and LBSPG models (Fig. 3). First, we see that the behaviors of the two models are very similar. This is not surprising because the LBSPG has been designed to reproduce the topological features of USWI. Figure 3(a) shows that the $N - 1$ condition with all-at-once rule or with one-by-one rule leads to bimodal distributions. Although the chance to obtain a cascade of any length bigger than one is very small, once a cascade starts, the chance of obtaining a blackout of a significant fraction of the system size is quite large. The difference between the one-by-one and all-at-once rules is that for the latter, the typical maximal blackout corresponds to approximately 40% of all lines, while for the former it is only 6% of all the lines.

If, instead, we return to the uniform tolerance condition of Ref. [20] and still use the all-at-once rule for line removal as in that work, the shape of the blackout distribution depends only on the tolerance level α [Figs. 3(b) and 3(c)]. As before, we start a cascade with an initial removal of a single line. In agreement with the original results of Ref. [20], the bimodality disappears for $\alpha > 1$ and a power-law distribution appears instead. The introduction of the one-by-one update rule together with the uniform tolerance condition shows the emergence of the power law for lower values of the tolerance, $\alpha = 0.8$, and a significant reduction of the sizes of the largest blackouts [Fig. 3(b)]. Again, the positions of the sharp drops of the cumulative distribution functions indicate the value of the order parameter $\mu = x$. The disappearance of these drops at large protection level α indicates that $\mu = 0$.

F. Summary of the power grid simulations

All in all, our simulations suggest that neither the $N - 1$ condition nor the one-by-one update rule is necessary nor sufficient for the emergence of the power-law distribution of the blackouts. The important condition is an overall level of line protection which also can be efficiently achieved by increasing enough the uniform tolerance α . The $N - 1$ condition alone does not eliminate large catastrophic blackouts in the bimodal distribution; instead it significantly reduces the chance of any blackout caused by initial removal of two lines. The one-by-one update rule also reduces the sizes of largest blackouts (the value of the order parameter μ) and removes the bimodality at smaller values of α , but is not strictly necessary. Another important aspect is the network topology. For sparse homogeneous networks whose topologies are close to the percolation threshold such as our RML model, the power-law distribution of blackouts emerges for smaller levels of protection, while for the other models similar failure rules still yield a bimodal distribution.

IV. BETWEENNESS CENTRALITY OVERLOADS

As one can see, the models of the power grids we study are relatively complex and have many different parameters, the




FIG. 3. Cumulative distributions of blackout sizes, measured as a function of the fraction of failed lines, f_d , for the USWI and LBSPG models. (a) Both models for the $N - 1$ condition with all-at-once and one-by-one update rules. (b) USWI model and (c) LBSPG for different uniform tolerance α and both update rules. For sufficiently large values of the tolerance, $\alpha > 1.2$, the distribution becomes approximately power law with large values of $\tau - 1 \geq 1$. The implementation of the one-by-one update rule with the uniform tolerance model shows that the power-law distribution emerges for a smaller value of the tolerance, $\alpha = 0.8$.

role of which is difficult to clarify. Therefore, it is desirable to study a simpler model of overloads in which the role of

each parameter becomes clear. It is also interesting to see if the emergence of the power-law distribution is a universal phenomenon for different overload models. A good candidate for such a model is the betweenness centrality model, which displays a clear bimodal distribution of blackouts for the all-at-once removal rule and uniform tolerance condition employed by Motter and Lai [22,26]. However, we can expect that, similarly to the power grid models, the implementation of the $N - 1$ condition and one-by-one removal rule in the betweenness centrality model will also lead to the emergence of a power-law distribution of the blackout sizes. Testing of this hypothesis will allow us to better understand the general mechanism of the emergence of the power-law distribution of the sizes of blackouts in the overload models.

A. Implementation of the betweenness centrality model

We build the betweenness centrality model as in Ref. [37]. Namely, we create a randomly connected graph with a given degree distribution $P(k)$ and compute the shortest paths between each pair of nodes. The length of each edge is assumed to be equal to $1 + \epsilon$, where ϵ is a normally distributed random variable with a small standard deviation $\sigma_\epsilon \ll 1$. This precaution is taken in order to make sure that each pair of nodes ij has a unique shortest path connecting them. The betweenness centrality $b(k)$ of each node k is computed as the number of all shortest paths ij passing through node k , such that $k \neq i$, $k \neq j$.

B. Implementation of the $N - 1$ condition

The $N - 1$ condition in the betweenness centrality model should be understood in the following way: If any one node i is deleted, the rest of the nodes remain connected and the betweenness centrality of each remaining node j , $b_i(j)$ does not exceed its maximal possible load $b^*(j)$. Thus

$$b^*(j) = \max_{i \neq j} b_i(j), \quad (1)$$

and the original graph must be biconnected. Note that in the original Motter and Lai model with a uniform tolerance $\alpha > 0$,

$$b^*(j) = (1 + \alpha)b(j), \quad (2)$$

where $b(j)$ is the initial betweenness of node j in a completely intact network.

To implement this model, we first construct a randomly connected graph of N nodes with a given degree distribution $P(k)$ with the Molloy-Read algorithm [44] and find its largest biconnected component with N_b nodes using the Hopcroft-Tarjan algorithm [45]. After that, we compute the betweenness centrality of each node and repeat N_b simulations with each node removed in turn to find $b^*(i)$ for each node. To find the distribution of cascades of failures we remove a pair of random nodes and find the nodes i with overloads $z = b(i)/b^*(i) > 1$. In the first case of all-at-once removal, at each stage of the cascade we remove all the nodes with overloads, recompute the betweenness centrality of the remaining nodes and repeat this process until no overloaded nodes are left. In the second case of one-by-one removal, at each stage of the cascades we remove only one of the




FIG. 4. (a) Cumulative distributions of blackout sizes, S_b , measured as the number of failed nodes in the Motter and Lai model with different values of tolerance α and all-at-once update rule. The system consists of $N = 10\,000$ nodes. Black lines indicate graphs for values of $\alpha = 0.02, 0.03, \dots, 0.19$. (b) Comparison of all-at-once and one-by-one update rules for selected values of the tolerance. The emergence of the power-law distribution in the one-by-one removal case for much smaller values of tolerance is seen.

overloaded nodes, the one with the largest overload z . We then recompute all the centralities, and repeat this process until no overloaded nodes are left. We study several cases of Erdős Rényi (ER) networks with different average degrees $\langle k \rangle$ and different number of nodes in the biconnected component N_b . The size of the blackout is computed as the number of failed nodes at the end of the cascade, N_d , or as the fraction of nodes failed nodes, $S_b = N_d/N_b$. We also study the distribution of the number of nodes disconnected from the giant component.

C. Uniform tolerance condition for the betweenness centrality model

Again, the $N - 1$ condition is not necessary for the emergence of the power-law distribution of the cascades in the Motter and Lai model. For example, if we take an ER network with $\langle k \rangle = 1.5$, all-at-once update rule, and uniform tolerance, for $0.1 < \alpha < 0.2$, we see the transition from a bimodal distribution of the blackout sizes to a power-law distribution with an exponential cutoff [Fig. 4(a)] as α increases.




FIG. 5. Cumulative distributions of blackout sizes, S_b , measured as the number of failed nodes in the Motter and Lai model with the $N - 1$ condition and one-by-one update rule for different system sizes and different average degrees.

The value of the order parameter μ indicated by the sharp drops of the cumulative distribution function preceded by a plateau, as in Fig. 1, becomes smaller and disappears as α increases, suggesting a second-order transition. Replacing the all-at-once update rule by the one-by-one rule shifts the transition from bimodal to a power-law distribution to a smaller value of α [Fig. 4(b)], but both behaviors can still be seen.

Note that the first-order phase transition and the bimodal distribution of the blackouts have been recently observed in Motter and Lai model and in its multiplex variant [26]. In this model the cascade is started by elimination of a single node with the largest betweenness, and the dynamics of the cascade follows all-at-once rule. They use the size of the giant component of the network at the end of cascade as the order parameter. Here we show that when the initially eliminated node is random this transition turns into the second-order transition characterized by a gradual reduction of the order parameter μ (size of the giant blackout) as α increases, and by a power-law distribution of avalanches with an exponential cutoff for α above the critical value. This difference most probably arises from the difference in the type of the initial attack from a targeted one on the node with highest centrality in Ref. [26] to a random one in this article. This question deserves further investigation.

D. The effect of the average degree and system size

Simultaneous application of the $N - 1$ condition and the one-by-one removal rule leads to the emergence of the power-law distribution of the small cascades for small values of $\langle k \rangle$ (Fig. 5). For small $\langle k \rangle$, close to the percolation threshold and small system sizes, the bimodality of the blackout size distribution disappears, and the distribution of blackouts becomes an approximate power law with an exponential cutoff. As the size of the network increases, the exponential cutoff disappears and the bimodality emerges again, with a peak for large blackouts emerging. The beginning of the cumulative distribution for all values of $\langle k \rangle$ and all sizes of the system exhibits the same slope of $-1/2$, which is the mean-field value

for the SOC models [11]. For large $\langle k \rangle$, we observe only a small fraction of blackouts distributed as a power law, while the entire distribution remains bimodal. However, the increase of the system size leads to the relative increase of the range of the power-law behavior. As the system size increases, we observe similar changes in the distribution of the blackouts for all values of $\langle k \rangle$ studied (Fig. 5). For small system sizes, we observe a smaller initial negative slope ($> -1/2$) and a sharp exponential cutoff. For larger system sizes the slope increases to almost exactly $-1/2$, and the cutoff shifts to larger values. For even larger sizes, the slope of the initial part of the cumulative distribution increases even further ($< -1/2$) but an inflection point appears in the distribution, indicating the start of bimodality. As the system size continues to increase the negative initial slope continues to grow, and the plateau region of the cumulative distribution separating giant blackouts from small blackouts becomes longer and longer, and the average size of the giant blackouts indicated by the sharp drop of the distribution becomes proportional to the system size, N , where the proportionality coefficient μ can be regarded as the order parameter of the transition.

Overall, the behavior of the cumulative distribution of the blackout sizes when $\langle k \rangle$ decreases is similar to the behavior of the avalanche size distribution in the failure branching process when $\langle b \rangle$ decreases and thus resembles a second-order transition. However, in the branching process, the dependence on the system size N is simple (Fig. 1); i.e., for $\langle b \rangle < 1$ there is no dependence at all, while for $\langle b \rangle > 1$ the part of the distribution for small avalanches and its cutoff s^* is still independent of N , but the average size of the giant avalanches (and hence the length of the plateau) increases as μN . In contrast, in the betweenness centrality model with the $N - 1$ condition protection and a one-by-one failure rule (Fig. 5), the dependence on N is more complex. For each $\langle k \rangle > 0$, there exist $N = N_c(\langle k \rangle)$ below which the distribution is a power law with an exponential cutoff that increases with N , while above $N > N_c(\langle k \rangle)$ the distribution is bimodal, and the average size of giant blackouts increases as μN . Thus for each $\langle k \rangle > 1$ it exists an effective critical point at $N = N_c(\langle k \rangle)$ of the second-order transition, at which $\mu \rightarrow 0$. Note that $N_c(\langle k \rangle) \rightarrow \infty$ when $\langle k \rangle \rightarrow 1$.

E. Cascade dynamics as a SOC model

As we saw, Carreras *et al.* [5] suggested that the power-law distribution of the blackouts is due to the long-term evolution of the power grids, driven by simultaneous growth of power consumption and upgrades of the most vulnerable power lines [4]. However, the propagation of cascades in a power grid of fixed size, when applying the $N - 1$ condition and the one-by-one update rule, resembles SOC models, in which the most vulnerable element is removed, and the vulnerability of other elements changes instantaneously. To test this hypothesis, we perform a detailed study of the propagation of the cascade in a biconnected component with $N - 1$ protection and one-by-one update rule. One can see that during the cascade the number of overloaded nodes fluctuates near zero and never gets too large. In the mean-field variant of SOC, the simplest of the SOC models [11], an array of N real numbers, “fitnesses,” uniformly distributed between 0 and 1 is




FIG. 6. (a) Dependence of the number of overloaded nodes as function of the cascade step for the Motter and Lai model for $\langle k \rangle = 1.2$ and $N = 500\,000$, with the $N - 1$ condition and one-by-one update rule. (b) Dependence of the size of the biconnected component and the giant component as a function of the cascade step.

created. At each time step the smallest element in the array and m randomly selected elements are replaced by new values selected from the same uniform distribution. As a result, after a certain transition period, the majority of elements have fitness above the critical value $f_c = 1/(m + 1)$. The number of active elements (elements with fitness less than a certain value f) performs a random walk, and an avalanche ends when the number of active elements returns to zero. When the value of the selected fitness is equal to the critical value $f = f_c$, the random walk is unbiased, and the distribution of avalanche lengths coincides with the distribution of returns to the origin of a random walker, which has the exponent $\tau - 1 = 1/2$ [46]. The number of active elements for $f = f_c$ scales as N^{d_f} with $d_f = 1/2$. We see [Fig. 6(a)] that in our model the behavior of the number of overloaded nodes is similar to the behavior of the random walk, which is consistent with the observed value of $\tau \approx 0.5$. However, the detrended fluctuation analysis [47,48] of the number of overloaded nodes during the cascade gives the value of the Hurst exponent ≈ 0.1 , which is much smaller than the exponent for the uncorrelated random walk of the mean-field SOC model. The number of overloaded nodes scales as $\ln N_b$, or a very small power law. In our model, the fitness of a node is its relative overload z_i , and the level of z_i is not arbitrary as in the SOC models, but is defined initially as $z_i = 1$. When, after implementing the $N - 1$ condition, we initially remove two nodes, the system is placed near the critical point, because the number of overloaded nodes after the initial attack is usually very small. However, this is done not automatically, as in the SOC models, but by implementing the $N - 1$ condition and the initial removal of two nodes. This action results in the overload of very few nodes, because the combined effect of the removal of two nodes rarely creates overloads larger than the maximal overloads created by the separate removal of each of these nodes. This is due to a long-tail effect; the removal of a single node will greatly increase the load of a few nodes, but barely affect most of the other nodes. However, when more nodes are removed in the process

of the cascade, the topology of the network changes, and the system is slowly driven away from the critical point.

We also see that the disintegration of the biconnected component is approximately a random removal of nodes due to subsequent overloads, with a slight preference for the removal of nodes with $k = 2$ over nodes with $k = 3$. At the beginning of the cascade the removals take place in the biconnected component, reducing it by the length of a chain of nodes with $k = 2$ in between a pair of nodes with $k \geq 3$. Each such removal reduces the size of the giant component only by one. But as more and more nodes are removed, further removals may happen in the singly connected part of the network. Figure 6(b) shows the size of the biconnected component to which the removed node belongs as function of the cascade stage. If it belongs to a singly connected part of the remaining network, we plot zero instead of the size of the biconnected component. When a node that is part of the singly connected part fails, the network separates into disconnected parts and the size of the giant component drops significantly. At the beginning, the smaller part is usually a dangling end, and the drop in size is relatively small, but eventually two approximately equal parts get separated. At this point the betweenness of all nodes decreases dramatically (approximately by factor of 4) and the cascade ends. Thus, although the cascade dynamics has some similarities with SOC, (at each time step the node with the largest overload is removed and the overloads of the rest of the nodes are changed), our model does not stay at the critical state but moves away from it.

The comparison between a mean-field variant of SOC and our network overload dynamics is also instructive to help us understand when a bimodal distribution occurs and when a power-law distribution does. A power-law distribution, such as in a branching process, occurs when there is a series of failures, each of which has a chance of precipitating $b = 1$ additional failures. Since there is the possibility, at each stage of the cascade, of no new nodes failing and the process stopping, the chance that exactly one more node will fail decreases exponentially. At the critical point, the cascade begins with $b = 1$, but as the network disintegrates b changes. The all-at-once update rule causes the network to disintegrate in fewer steps, since multiple nodes are removed at each step, and thus b increases during the cascade of failures without many chances for the cascade to end prematurely. Thus, after a few steps $b > 1$, and the process will not stop itself with an intermediate level of destruction. However, there are two different conditions that will prevent b from increasing drastically as the network begins to disintegrate. One is a high level of overall protection, whether through the implementation of the $N - 1$ condition or through a high uniform α . High overall protection means that minor network destruction will not cause widespread overload and drive b away from one; only severe network destruction can do that. Thus, the overload process has a chance to stop spontaneously before it reaches that stage. The other condition is if the network is near the percolation point, where the failure of a single node has a high chance of separating a small, but sizable, component from the network. This will drastically decrease the overall load and thus stabilize the network, arresting the cascade of failures [37]. In conclusion, a one-by-one update rule and a high overall protection both contribute to the formation of a

power law. The effect of the system size discussed above can also be understood in light of the comparison to a branching process. The larger the system, the more nodes (as an absolute number, not a fraction) fail following the initial attack. Thus, the chance that none of the overloaded nodes will cause another node to overload (i.e., a spontaneous end to the cascade) becomes increasingly unlikely as the number of overloaded nodes increases, even if b remains close to one for many steps of the cascade. Thus, the chance that the cascade of failures will end with an intermediate amount of damage becomes more unlikely when the size of the original system increases.

V. DISCUSSION AND CONCLUSION

We have shown that the dc model of the power grid has remarkable similarities to the Motter and Lai model of betweenness centrality, suggesting that the exact physical laws governing the flux do not play as big a role as the network topology and the dynamics of the cascade propagation. We have shown that both overload models have similarities and differences with the topological models [27–34]. In the topological models the outcome of the cascade depends only on the topology of the network and the location of the initial damage, while in the overload models the outcome significantly depends on the dynamics of the cascade, which itself depends on the order of removal of overloaded elements and the relative speed of failure of overloaded elements and redistribution of flux after the failure. In general, if the failure is slow compared to the redistribution of flux (one-by-one removal rule) the cascades become smaller, while for fast failures of overloaded elements (the all-at-once removal rule) the blackouts can be larger by an order of magnitude than in the one-by-one case. But the one-by-one removal rule is neither necessary nor sufficient for the elimination of “giant” blackouts engulfing a finite fraction of the entire system.

We also find that the $N - 1$ condition is neither necessary nor sufficient for the disappearance of large cascades destroying a finite fraction of the elements in the system, but it reduces the probability of any cascade by an order of magnitude compared to the universal tolerance condition even for a large tolerance α , if two elements are deleted simultaneously at the beginning of the cascade.

The general picture of the behavior of the overload models is consistent with the simple model of the failure branching process, which demonstrates a second-order phase transition. For large protection levels (small $\langle b \rangle$) the distribution of blackout sizes is a power law with exponential cutoff, while for low protection levels (large $\langle b \rangle$) the distribution of blackout sizes becomes bimodal, when a fraction of large blackouts engulfing a finite fraction of the entire system emerges. This fraction and the size of these “giant” blackouts (which serve as the order parameter of the problem μ) continuously decrease when the protection level increases and vanishes when $\langle b \rangle$ reaches its critical value $\langle b \rangle = 1$.

The protection level can be increased by several factors: (1) increase of the tolerance α , (2) introduction of the one-by-one removal rule, (3) the reduction of the network connectivity, i.e., bringing the network closer to the percolation threshold, by reducing the average degree $\langle k \rangle$, and (4) by decreasing the system size. These methods have an additive effect, i.e., when

applied simultaneously they will achieve larger protection levels than when applied separately. For example, application of the one-by-one removal rule alone may not change the bimodal distribution of blackouts to unimodal; however, when implemented together with the uniform tolerance condition, it will help to achieve the unimodal distribution for smaller values of the tolerance α .

The empirical observation of the power-law distribution of the blackout sizes in real power grids indicates that these grids are near the critical point of the correspondent failure branching process. However, this is unlikely to be caused by a SOC mechanism in which the model is driven to the critical point by a certain dynamic rule. One hypothesis would be that it is the $N - 1$ criterion and the one-by-one update rule that brings the grid to the SOC state. However, we observe here that these two conditions are neither necessary nor sufficient for the existence of a power-law distribution of blackout sizes. Most likely the power grids are brought to the vicinity of the critical point by a risk-price compromise, in which the utilities, with the firm goal of avoiding the catastrophic blackouts that engulf a significant fraction of the entire system, increase the level of protection up to the critical point of the branching process $\langle b \rangle < 1$, but, in order to minimize costs under this condition, allow $\langle b \rangle$ to be as close as possible to the critical point from below, so that limited but occasionally very large blackouts are still possible. It is still desirable to develop a model of long time evolution of power grids similar to Ref. [4], which reproduces the power-law distribution of avalanches. The observation of power-law distribution of the blackouts in the sparse networks close to the percolation point also emphasizes the role of islanding for preventing large blackouts.

A plausible reason for the emergence of the power-law distribution of cascades is presented in a recent work of Nesti *et al.* [18], who relate it to the power-law distribution of city sizes and hence the distribution of loads (“sinks”) in the dc approximation of the grid. They observe $\tau > 2$ in agreement with empirical observations [5]. However, our study suggests that the distribution of loads is irrelevant for the emergence of the power-law distribution of blackouts for large protection levels. In fact, the USWI model has a lognormal distribution of the loads I_i^- with a sharp cutoff, while the LBSPG model has a power-law distribution $P(|I_i^-|) \sim |I_i^-|^{-2.2}$ of the loads, but their behavior is very similar in our study, and in both models the transition from a bimodal distribution to a power-law distribution of the blackouts occurs at high level of protection (high tolerance). The appearance of the power-law distribution of blackouts also in the Motter and Lai model suggests that this feature might be a very general phenomenon not necessarily related to the power-law distribution of the loads.

ACKNOWLEDGMENTS

This material is based upon work supported in part by the U.S. Department of Energy’s Office of Energy Efficiency and Renewable Energy (EERE) under the Solar Energy Technology Office, ASSIST Initiative Award No. DE-EE0008769, as well as by Defense Threat Reduction Agency grants HDTRA1-14-1-0017, HDTRA1-19-1-0016. S.V.B. and G.C.

acknowledge the partial support of this research through the Dr. Bernard W. Gamson Computational Science Center at Yeshiva College.

APPENDIX

As mentioned in Sec. II A there is a third class of initial attacks in which the system first undergoes a cascade of failures caused by a massive initial attack with $p > p_t$ and the order parameter stabilizes at $S(p)$. This is analogous to the first class of initial attacks. Afterwards, the avalanche is started by the removal of one additional element. In the well-studied topological models, it is known [35,36] that

$$S(p) - S(p_t + 0) \sim (p - p_t)^{1/2}, \quad (\text{A1})$$

where $S(p_t + 0)$ is the limit of $S(p)$ for $p \rightarrow p_t$ above the step discontinuity. Because in the topological models the order of removal elements does not play any role, the removal of one element after the system is stabilized at p is identical to the initial removal of $pN + 1$ elements, where N is the total number of elements in the system. In either situation, the average avalanche size under this

initial conditions is $s = N[S(p + 1/N) - S(p)] \approx \partial S(p)/\partial p \sim (p - p_t)^{-1/2}$. Hence, the avalanche size diverges above the transition point p_t , with the divergence characterized by the critical exponent $\gamma = -1/2$.

In network theory this transition is called a hybrid transition, but in fact this behavior is completely equivalent to the mean-field behavior near the spinodal of the first-order phase transition [49], such as the behavior of the isothermal compressibility near the spinodal of the gas-liquid phase transition. This follows from the fact that the isothermal compressibility is equivalent to the susceptibility in the Ising model, which in turn is equivalent to the average cluster size at the percolation transition.

The divergence takes place only for $p > p_t$, while for $p < p_t$ the avalanches remain of finite size or do not exist at all, since $S(p) = 0$ for $p < p_t$. For $p \rightarrow p_t + 0$, the distribution of the avalanche sizes develops a power-law behavior with an exponential cutoff S^* : $P(s) \sim s^{-\tau} \exp(-s/S^*)$ with the mean-field exponent $\tau = 3/2$ [35,36], and $S^* \sim (p - p_t)^{-1/\sigma}$ with $\sigma = -1$; they, together with $\gamma = -1/2$, satisfy a usual percolation scaling relation $\gamma = (2 - \tau)/\sigma$, but with different γ and σ than in the percolation theory [17,40].

-
- [1] U.S.-Canada Power System Outage Task Force: Final report on the August 14, 2003 blackout in the United States and Canada: Causes and recommendations (Canada, 2004), <https://eta-publications.lbl.gov/sites/default/files/2003-blackout-us-canada.pdf>.
 - [2] Report of the enquiry committee on grid disturbance in the Northern region on 30th July 2012 and in Northern, Eastern and North-Eastern region on 31st July 2012 (2012), https://powermin.nic.in/sites/default/files/uploads/GRID_ENQ_REP_6_8_12.pdf.
 - [3] Report on the Grid Disturbance on 30th July 2012 and Grid Disturbance on 31st July 2012 (2012), http://www.cercind.gov.in/2012/orders/Final_Report_Grid_Disturbance.pdf.
 - [4] H. Ren, I. Dobson, and B. A. Carreras, *IEEE Trans. Power Syst.* **23**, 1217 (2008).
 - [5] B. A. Carreras, D. E. Newman, and I. Dobson, *IEEE Trans. Power Syst.* **31**, 4406 (2016).
 - [6] B. A. Carreras, V. E. Lynch, I. Dobson, and D. E. Newman, *Chaos* **12**, 985 (2002).
 - [7] B. A. Carreras, V. E. Lynch, I. Dobson, and D. E. Newman, *Chaos* **14**, 643 (2004).
 - [8] I. Dobson and L. Lu, *IEEE Trans. CAS* **39**, 762 (1992).
 - [9] P. Bak and K. Sneppen, *Phys. Rev. Lett.* **71**, 4083 (1993).
 - [10] M. Paczuski, S. Maslov, and P. Bak, *Phys. Rev. E* **53**, 414 (1996).
 - [11] H. Flyvbjerg, K. Sneppen, and P. Bak, *Phys. Rev. Lett.* **71**, 4087 (1993).
 - [12] R. Albert, I. Albert, and G. L. Nakarado, *Phys. Rev. E* **69**, 025103(R) (2004).
 - [13] Y. Yang, T. Nishikawa, and E. Motter, *Science* **358**, eaan3184 (2017).
 - [14] M. Anghel, K. A. Werley, and A. E. Motter, in *Proceedings of the 2007 40th Annual Hawaii International Conference on System Sciences HICSS'07*, Big Island, HI, USA (IEEE, Piscataway, NJ, 2007), p. 113.
 - [15] H. Cetinay, S. Soltan, F. A. Kuipers, G. Zussman, and P. Van Mieghem, *IEEE Trans. Netw. Sci. Eng.* **5**, 301 (2017).
 - [16] S. Soltan, A. Loh, and G. Zussman, *IEEE Trans. Control Netw. Syst.* **5**, 1424 (2017).
 - [17] D. Stauffer and A. Aharony, *Introduction to Percolation Theory* (Taylor and Francis, Philadelphia, 1994). Note that Stauffer and Aharony define the distribution of cluster sizes s by counting the number of clusters of different sizes $N(s)$ in a finite network, which is characterized by the critical exponent $\tau = 5/2$. An alternative definition, consistent with the notation in the SOC literature, is the probability $p(s)$ of a randomly selected site to belong to a cluster of size s . This probability distribution is characterized by a power law with an exponent $\tau' = \tau - 1 = 3/2$.
 - [18] T. Nesti, F. Sloothak, and B. Zwart, *Phys. Rev. Lett.* **125**, 058301 (2020).
 - [19] G. K. Zipf, *Soc. Forces* **28**, 340 (1950).
 - [20] R. Spiewak, S. Soltan, Y. Forman, S. V. Buldyrev, and G. Zussman, *Netw. Sci.* **6**, 448 (2018).
 - [21] S. Pahwa, C. Scoglio, and A. Scala, *Sci. Rep.* **4**, 3694 (2014).
 - [22] A. E. Motter and Y.-C. Lai, *Phys. Rev. E* **66**, 065102(R) (2002).
 - [23] A. E. Motter, *Phys. Rev. Lett.* **93**, 098701 (2004).
 - [24] S. Soltan, D. Mazauric, and G. Zussman, *Proc. ACM e-Energy* **14**, 195 (2014).
 - [25] P. Hines, E. Cotilla-Sanchez, and S. Blumsack, *Chaos* **20**, 033122 (2010).
 - [26] O. Artime and M. De Domenico, *New J. Phys.* **22**, 093035 (2020).
 - [27] D. J. Watts, *Proc. Natl. Acad. Sci. USA* **99**, 5766 (2002).

- [28] J. P. Gleeson and D. J. Cahalane, *Proc. SPIE* **6601**, 66010W (2007).
- [29] G. J. Baxter, S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes, *Phys. Rev. E* **82**, 011103 (2010).
- [30] G. J. Baxter, S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes, *Phys. Rev. E* **83**, 051134 (2011).
- [31] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin, *Nature (London)* **464**, 1025 (2010).
- [32] R. Parshani, S. V. Buldyrev, and S. Havlin, *Proc. Natl. Acad. Sci. USA* **108**, 1007 (2011).
- [33] M. A. Di Muro, S. V. Buldyrev, H. E. Stanley, and L. A. Braunstein, *Phys. Rev. E* **94**, 042304 (2016).
- [34] M. A. Di Muro, L. D. Valdez, H. H. Aragao Rego, S. V. Buldyrev, H. E. Stanley, and L. A. Braunstein, *Sci. Rep.* **7**, 15059 (2017).
- [35] S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes, *Physica D* **224**, 7 (2006).
- [36] G. J. Baxter, S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes, *Phys. Rev. Lett.* **109**, 248701 (2012).
- [37] Y. Kornbluth, G. Barach, Y. Tuchman, B. Kadish, G. Cwilich, and S. V. Buldyrev, *Phys. Rev. E* **97**, 052309 (2018).
- [38] R. Parshani, S. V. Buldyrev, and S. Havlin, *Phys. Rev. Lett.* **105**, 048701 (2010).
- [39] T. E. Harris, *The Theory of Branching Processes*, Die Grundlehren der Mathematischen Wissenschaften Vol. 119 (Springer, Berlin, 1963).
- [40] A. Bunde and S. Havlin (eds.), *Fractals and Disordered Systems* (Springer, New York, 1996).
- [41] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman, in *Proceedings of the IEEE INFOCOM 2014 - IEEE Conference on Computer Communications, Toronto, ON, Canada* (IEEE, 2014), pp. 2634–2642.
- [42] S. Soltan, A. Loh, and G. Zussman, *IEEE Syst. J.* **13**, 625 (2019).
- [43] S. Soltan, A. Loh, and G. Zussman, Columbia University synthetic power grid with geographical coordinates, <https://egriddata.org/dataset/columbia-university-synthetic-power-grid-geographical-coordinates> (2018).
- [44] M. Molloy and B. Reed, *Combin. Probab. Comput.* **7**, 295 (1998).
- [45] J. Hopcroft and R. Tarjan, *Commun. ACM* **16**, 372 (1973).
- [46] W. Feller, *An Introduction to Probability Theory and its Applications*, Vol. I, 3rd ed. (John Wiley and Sons, Princeton, New Jersey, 1968).
- [47] C. K. Peng, S. V. Buldyrev, S. Havlin, M. Simons, H. E. Stanley and A. L. Goldberger, *Phys. Rev. E* **49**, 1685 (1994).
- [48] S. V. Buldyrev, A. L. Goldberger, S. Havlin, R. N. Mantegna, M. E. Matsu, C.-K. Peng, M. Simons, and H. E. Stanley, *Phys. Rev. E* **51**, 5084 (1995).
- [49] M. A. Di Muro, S. V. Buldyrev, and L. A. Braunstein, *Phys. Rev. E* **101**, 042307 (2020).