

Uniform redundancy allocation maximizes the robustness of flow networks against cascading failures

Omur Ozel, Bruno Sinopoli, and Osman Yağın

Department of ECE, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA



(Received 29 April 2018; published 8 October 2018)

We study the robustness of flow networks against cascading failures under a *partial* load redistribution model. In particular, we consider a flow network of N lines with initial loads L_1, \dots, L_N and free spaces (i.e., redundant space) S_1, \dots, S_N that are independent and identically distributed with joint distribution $P_{LS}(x, y) = \mathbb{P}[L \leq x, S \leq y]$. The *capacity* C_i is the maximum load allowed on line i and is given by $C_i = L_i + S_i$. When a line fails due to overloading, it is removed from the system and $(1 - \varepsilon)$ fraction of the load it was carrying is redistributed *equally* among all remaining lines in the system. The rest (i.e., ε fraction) of the load is assumed to be *lost* or *absorbed*, e.g., due to advanced circuitry disconnecting overloaded power lines or an *interconnected* network or material absorbing a fraction of the flow from overloaded lines. We analyze the robustness of this flow network against *random* attacks that remove a p fraction of the lines. Our contributions include (i) deriving the *final* fraction of alive lines $n_\infty(p, \varepsilon)$ for all $p, \varepsilon \in (0, 1)$ and confirming the results via extensive simulations; (ii) showing that partial redistribution might lead to (depending on the parameter $0 < \varepsilon \leq 1$) the order of transition at the critical attack size p^* changing from first to second-order; and (iii) proving analytically that flow networks achieve maximum robustness [quantified by the area $\int_0^1 n_\infty(p, \varepsilon) dp$] when all lines have the same free space regardless of their initial load. The optimality of equal free-space allocation is also confirmed on real-world data from the UK National Power Grid.

DOI: [10.1103/PhysRevE.98.042306](https://doi.org/10.1103/PhysRevE.98.042306)

I. INTRODUCTION

Flow network abstractions have been extensively used to analyze complex phenomena occurring in power line networks, financial networks, transportation networks, and biological ecosystems. Current trends in technological development such as online social media, cyber-physical systems, and the internet of things have enabled a plethora of new applications involving dynamical interactions over and within flow networks. In this respect, current research on networks covers phenomena such as information dissemination and influence propagation [1–5], percolation [6–11], and robustness [12–17].

This paper focuses on the robustness of flow networks. In particular, we are interested in understanding how the random failure of a fraction lines might trigger failures of other lines leading to what has been identified in the literature as *cascading failures*. Real-life phenomena such as blackouts in power networks [18] and crises in financial networks [19] occur as a consequence of failures that follow one after the other. Consequently, understanding the possible mechanisms of such failures is of paramount importance; e.g., see Ref. [20]. The robustness of flow network structures against cascading failures is an active research topic and has received recent interest from many researchers [12–14]. For example, Refs. [21–25] consider power networks where the failure mechanism is the equal redistribution of load upon the failure of a power line. We consider a similar phenomenon in flow networks. In particular, we build our analysis upon the well-known fiber bundle model [26]. Fiber bundle models have

been used in a wide range of applications, including fatigue [27], failure of composite materials [28], and landslides [29].

Flow networks typically include built-in systems to counter cascade formations and alleviate the spread of the adverse effect of failures over other lines in the network [17]. In this paper, our objective is to obtain a unified understanding of cascading failures for networks with a specific mechanism to counter such failures. In particular, we study the robustness under *partial* load redistribution in a *democratic fiber bundle like* model. Our problem setting is as follows: We consider N lines whose initial loads L_1, \dots, L_N and free spaces S_1, \dots, S_N have joint distribution $P_{LS}(x, y) = \mathbb{P}[L \leq x, S \leq y]$ and are independent and identically distributed along lines. The maximum flow allowed on a line i defines its *capacity* and is given by $C_i = L_i + S_i$. When a line fails due to overloading, it is removed from the system and $(1 - \varepsilon)$ fraction of the load it was carrying (at the moment of failing) gets redistributed *equally* among all remaining lines in the system; hence, we refer to this as the *partial* load redistribution model. The rest (i.e., ε fraction) of the load is assumed to be *lost* or *absorbed*, e.g., due to advanced circuitry disconnecting overloaded power lines or an *interconnected* network or material absorbing a fraction of the flow from overloaded lines. Throughout the paper, we refer to ε as the loss or absorption, interchangeably.

We study the robustness of the flow network described above against *random* attacks or failures. In particular, we develop a complete analytic framework that unravels the dynamics of cascading failures initiated by a random attack removing a p fraction of the lines. This is done by computing

recursively the fraction of surviving lines and the additional load distributed to each surviving line at every step of the cascade. We show how these recursive relations can be utilized to derive the *final* fraction $n_\infty(p, \varepsilon)$ of alive lines, under any attack size $0 < p < 1$ and any absorption factor $\varepsilon \in (0, 1)$. These results are also confirmed through an extensive numerical study.

An interesting observation from our results concerns the *order* of phase transition that the system exhibits at the *critical* attack size p^* ; this is the smallest attack size p that collapses the entire system. It is widely known [23,26] that flow networks exhibit only first-order phase transitions; this phenomenon is referred to as the *abrupt rupture* in the context of fiber bundles. However, we show in this paper that under the *partial* redistribution model, the order of the phase transition might be first or second order depending on the absorption factor ε . In particular, we demonstrate that the order of the phase transition changes from first to second-order after ε exceeds a certain *tricritical* point. This is reminiscent of the behavior observed in percolation studies of inter-connected networks [6], where reducing the coupling strength between two inter-connected networks might lead to a second-order transition [30].

Utilizing the analytic framework we developed, we then seek to answer a very fundamental question concerning flow networks: Under given constraints on the total load and capacity available to N lines, how should these be allocated so that the resulting flow network has the maximum possible robustness against random attacks? This question is particularly challenging under the *partial* flow redistribution model owing to the fact that the system's robustness may not be fully characterized by the critical attack size p^* . Simply put, it is possible for an allocation to lead to a *large* p^* but a very poor performance [in terms of final system size $n_\infty(p, \varepsilon)$] for even small attack sizes p . To meaningfully assess the robustness of a flow network under attacks of all possible scale, we thus consider the robustness metric proposed in Ref. [31] given as

$$\mathcal{R}(\varepsilon) := \int_0^1 n_\infty(p, \varepsilon) dp.$$

It can be seen that $\mathcal{R}(\varepsilon)$ measures the *area* that falls under the curve $n_\infty(p, \varepsilon)$ as p varies from zero to one. More importantly, $\mathcal{R}(\varepsilon)$ gives the *mean* fraction of lines that will survive an attack whose size p is random over $(0,1)$.

Although it is often very difficult (if not impossible) to analytically prove optimality results for any network property, here we are able to give a complete mathematical proof showing how the robustness metric $\mathcal{R}(\varepsilon)$ can be maximized. In particular, we show that for a given total initial load and capacity to be allocated to N lines, robustness $\mathcal{R}(\varepsilon)$ is maximized when all lines have the same free space, i.e., capacity minus initial load, regardless of how the initial loads are allocated. This result is shown to hold under any absorption factor $0 \leq \varepsilon \leq 1$.

The optimality of equal free-space allocation (in the sense of maximizing robustness) in flow networks is interesting for several reasons. In particular, it shows that the widely used [32–35] assumption that line capacity is a fixed proportion of its load (e.g., by setting $C_i = (1 + \alpha)L_i$ with α defining

the *tolerance factor* used for all N lines) leads to suboptimal robustness. In fact, our result shows that to maximize robustness, lines with higher initial load should be given a smaller tolerance factor (i.e., free space divided by load). Moreover, our results show that although a flow network's robustness might be improved by increasing its capability to *absorb* failed load (which often will require a more expensive or complicated design), the system can be made optimally robust by allocating all lines the same free space irrespective of the absorbing capability.

We test the above result under a new cascading failure model for flow networks that combine local and global redistribution approaches; e.g., a portion of the failed load is redistributed in the local neighborhood according to network topology while the rest is redistributed globally. For this model, we provide simulation results that suggest that the mean-field redistribution model analyzed here captures the qualitative behavior of system robustness well. In particular, we observe that uniformly allocating free spaces promises to be optimal also in this more general setting where (part of the) failed load gets redistributed locally according to a network topology.

Finally, we test our optimality result on real-world systems. In particular, we run numerical experiments using the UK National Grid data available in Ref. [36]. These experiments confirm our optimality result and suggest that some real-world systems do not exhibit optimal robustness against random failures. In fact, their robustness can be improved significantly by reallocating the total available capacity in such a way as to ensure that every line has the same free space. We believe that our results provide interesting insights into the dynamics of cascading failures in flow networks and call for a careful examination of the capacity allocations in existing real-world systems.

The rest of the paper is organized as follows. In Sec. II, we explain the system model in detail and present the problem definition. In Sec. III, we obtain an iterative dynamical relation for the extra load per alive line at every stage of the cascading failure process, under general load and free-space distributions. Section IV is devoted to numerical results that confirm the main findings of the paper for systems of finite size. In Sec. V, we present a complete analytic framework to establish the optimal distribution of load and free space (when mean values of both are fixed) that leads to maximum robustness. These findings are confirmed via extensive simulations in Sec. VC using synthetic data (with various commonly used distributions utilized to generate load and free-space values) as well as real-world data from the UK National Power Grid. The paper is concluded in Sec. VI.

II. MODEL AND PROBLEM DEFINITION

A. Partial load-redistribution model

We consider a network with N lines $\mathcal{L}_1, \dots, \mathcal{L}_N$ with initial loads L_1, \dots, L_N . The *capacity* C_i of a line \mathcal{L}_i defines the maximum power flow that can be carried by it and is expressed as

$$C_i = L_i + S_i, \quad i = 1, \dots, N, \quad (1)$$

where S_i denotes the *free space* that is assigned to line \mathcal{L}_i . Alternatively, the capacity of a line \mathcal{L}_i can be defined as a

factor of its initial load, i.e.,

$$C_i = (1 + \alpha_i)L_i, \quad (2)$$

with $\alpha_i > 0$ denoting its *tolerance* factor. Accordingly, the free space S_i is given in terms of the initial load L_i as $S_i = \alpha_i L_i$. Most existing works assume a *fixed* tolerance factor for all lines in the system, i.e., $\alpha_i = \alpha$ for all i ; e.g., see Refs. [32–35].

The main assumption of our model is that when a line fails due to *overloading*, i.e., due to its load exceeding its capacity, it is removed from the system and $(1 - \varepsilon)$ fraction of the load it was carrying (at the moment of failing) gets redistributed *equally* among all remaining lines in the system; hence we refer to this as the *partial* load redistribution model. The rest (i.e., ε fraction) of its load is assumed to be *lost* or *absorbed*.

The partial load redistribution model is motivated by several real-world scenarios. For instance, most real-world power systems are equipped with advanced protection circuits that immediately disconnect the overloaded lines from the rest of the grid [37,38]; the parameter ε would then represent the fraction of lines protected by such advanced circuitry. Alternatively, we can think of a flow network (respectively, a bundle of fibers) that is *inter-connected* with another network (respectively, material) that can absorb a fraction of the flow from overloaded lines.

Throughout we assume that the pairs (L_i, S_i) are independently and identically distributed (i.i.d.) with the joint distribution $P_{LS}(x, y) := \mathbb{P}[L \leq x, S \leq y]$ for each $i = 1, \dots, N$. The corresponding joint probability density function is given by $p_{LS}(x, y) = \frac{\partial^2}{\partial x \partial y} P_{LS}(x, y)$. We assume that the marginal densities $p_L(x)$ and $p_S(y)$ are continuous on their support. L_{\min} and S_{\min} denote the minimum values for load L and free space S , respectively, throughout the paper. We assume that $L_{\min}, S_{\min} > 0$.

Our load redistribution model is intimately related to the *democratic* fiber bundle model [22,23] where N parallel fibers with failure thresholds C_1, \dots, C_N share an applied total force F *equally*. In this line of literature, it has been of interest to study the dynamics of recursive failures in the bundle as the applied force F increases; e.g., see Refs. [39–41]. This model was used by Ref. [21] in the context of power line networks, with F corresponding to the total load shared *equally* by N power lines. See also Ref. [42] for the latest developments on the democratic fiber bundle model relating to a power line network. The relevance of the equal load-redistribution model for power systems stems from its relation to Kirchhoff's law in the mean-field sense. Our current partial load redistribution model builds upon that in Ref. [42]. Even though our model involves a *global* redistribution of a fraction of the extra load due to line failures, the capability to partially capture the failures is, in effect, due to the ability to keep a portion of the extra load due to failed lines in a *local* level, possibly in a secondary network. While our current work does not consider the interaction of local and global behavior, our framework helps build a step towards this direction.

B. Problem definition

Our main goal is to study the robustness of the flow network under the partial load redistribution rule described

above. We consider a *random* attack or a random failure that leads to p fraction of the lines to be removed from the system. We assume that the load of these initially failed lines are redistributed *in full* to the nonattacked lines, with each nonattacked line receiving an equal portion of the total load failed. Our motivation in distinguishing these initial failures (resulting from an attack) from failures due to overloading of lines is twofold. First, in the case of a physical attack to the system, we would expect any advanced circuitry or interconnection to other networks to be damaged along with the failed lines, making it impossible for ε fraction of the failed load to be absorbed. Second, this assumption ensures that a random attack against p fraction of the lines is equivalent (in the mean-field sense) to a disturbance caused by increasing the initial load of every line (or, force applied to every fiber) by $p\mathbb{E}[L]/(1 - p)$. This in turn enables our analysis to provide insights on the robustness of the system against both random attacks (as commonly considered in the context of power systems) as well as the increase of total applied load or force (as commonly considered in the context of fiber bundles).

After the initial load redistribution, the amount of load on each alive line will be given by its initial load plus its share of the total load of the failed lines. This, in turn, leads to the failure of additional lines due to the updated flow exceeding their capacity. In the ensuing stages of this process, the network is assumed to have the capability of *absorbing* (i.e., removing from the system) ε fraction of the load from lines who fail due to overloading. Put differently, if a line \mathcal{L}_i fails due to its load $L_i(t)$ at time t exceeding its capacity, then only $(1 - \varepsilon)L_i(t)$ amount of load will be redistributed, in an equal manner, to the remaining lines; as mentioned before, the system is assumed to absorb the portion $\varepsilon L_i(t)$ either by help of advanced circuitry or by means of shedding that portion of load.

In the most general scenario for partial redistribution, we could allow ε to depend on time t and the extra load incurred per line at that time. However, our main goal for analysis in this paper is to understand the case when ε is constant throughout time. The load redistribution process continues recursively until no further failures occur, potentially generating a *cascade of failures*. Our goal is to understand the limits associated with this process. We let $n_\infty(p, \varepsilon)$ denote the *final* (i.e., steady-state) fraction of alive lines when a p fraction of lines is randomly attacked initially; as before, ε denotes the fraction of flow from *overloaded* lines that will be lost (and thus will not be redistributed to the remaining lines) at each stage of the cascade process.

We derive expressions for $n_\infty(p, \varepsilon)$ for all attack sizes $0 < p < 1$ and any $0 \leq \varepsilon \leq 1$ to understand the *robustness* of the network under the partial load redistribution model. We will be particularly interested in understanding the *critical* attack size p^* at which $n_\infty(p, \varepsilon)$ drops to zero, and in developing design guidelines (in the sense of choosing the distribution p_{LS}) to optimize network robustness under given constraints.

With regard to the notation in use: Probabilistic statements are made with respect to probability measure \mathbb{P} , and we denote the corresponding expectation operator by \mathbb{E} . The indicator function of an event A is denoted by $\mathbf{1}[A]$.

III. ANALYTIC RESULTS

In this section, we provide a mean-field analysis for the cascading failures of lines under the model described in Sec. II. We start by deriving recursive relations concerning the fraction f_t of lines that are *failed* at time stage $t = 0, 1, \dots$. The number of links that are still alive at time t is then given by $N_t = N(1 - f_t)$ for all $t = 0, 1, \dots$. The cascading failures start with a random attack that targets a fraction p of lines. Hence, we have $f_0 = p$. Upon the failure of these f_0 lines, their load will be redistributed to the remaining $(1 - f_0)N$ lines, with each remaining line receiving an equal portion of the failed load. Since the pN lines that have been attacked are selected uniformly at random, the mean total load that will be redistributed to the remaining lines is given by $\mathbb{E}[L]pN$. The resulting extra load per alive line, Q_0 , is thus given by

$$Q_0 = \frac{\mathbb{E}[L]pN}{(1 - p)N} = \mathbb{E}[L] \frac{f_0}{1 - f_0}. \quad (3)$$

In the next stage, a line i that survives the initial attack fails when its new load reaches its capacity. For convenience, we assume that a line also fails when its load equals its capacity, i.e., when

$$L_i + Q_0 \geq L_i + S_i,$$

or, equivalently $S_i \leq Q_0$. Therefore, at stage $t = 1$, an additional fraction $\mathbb{P}[S \leq Q_0]$ of the lines that were alive at the end of stage 0 fail. This yields

$$f_1 = f_0 + (1 - f_0)\mathbb{P}[S \leq Q_0] = 1 - (1 - f_0)\mathbb{P}[S > Q_0].$$

Now, to compute the extra load per alive line at stage 1, which we denote by Q_1 , we need to sum the total load of the lines failed precisely at stage 1, multiply it by $(1 - \varepsilon)$ to account for the load absorbed or lost as explained in Sec. II, divide it by the new system size $(1 - f_1)N$ and add this fraction on top of the existing extra load. With \mathcal{A} denoting the initial set of lines attacked, this gives

$$\begin{aligned} Q_1 &= Q_0 + \frac{1 - \varepsilon}{(1 - f_1)N} \mathbb{E} \left[\sum_{i \notin \mathcal{A}: S_i \leq Q_0} (L_i + Q_0) \right] \\ &= Q_0 + \frac{1 - \varepsilon}{(1 - f_1)N} \sum_{i \notin \mathcal{A}} \mathbb{E}[(L_i + Q_0)\mathbf{1}[S_i \leq Q_0]] \\ &= Q_0 + (1 - \varepsilon)(1 - f_0) \frac{\mathbb{E}[(L + Q_0) \cdot \mathbf{1}[S \leq Q_0]]}{1 - f_1}, \end{aligned}$$

where the last step uses $|\mathcal{A}|/N = p = f_0$.

At the stage $t = 2$, the following two conditions are needed for a line to still stay alive: (i) it should not have failed until this stage, which happens with probability $1 - f_1$ and necessitates its free space to satisfy $S > Q_0$; and (ii) its free space should also satisfy $S > Q_1$ so that its capacity is still larger than its current load. Thus, the fraction of failures f_2 at this stage is given by

$$f_2 = 1 - (1 - f_1)\mathbb{P}[S > Q_1 | S > Q_0].$$

However, we can calculate the total load that is redistributed to the remaining lines as before:

$$Q_2 = Q_1 + (1 - \varepsilon)(1 - f_0) \frac{\mathbb{E}[(L + Q_1) \cdot \mathbf{1}[Q_0 < S \leq Q_1]]}{1 - f_2}.$$

The form of the recursive equations for each $t = 0, 1, \dots$ can now be seen to be as follows:

$$\begin{aligned} f_{t+1} &= 1 - (1 - f_t)\mathbb{P}[S > Q_t | S > Q_{t-1}], \\ Q_{t+1} &= Q_t + (1 - \varepsilon)(1 - f_0) \\ &\quad \times \frac{\mathbb{E}[(L + Q_t) \cdot \mathbf{1}[Q_{t-1} < S \leq Q_t]]}{1 - f_{t+1}}, \quad (4) \\ N_{t+1} &= (1 - f_{t+1})N, \end{aligned}$$

where $f_0 = p$, $N_0 = N(1 - p)$, and $Q_0 = \mathbb{E}[L] \frac{p}{1 - p}$. For convenience, we also let $Q_{-1} = 0$. From Eq. (4) we see that cascades stop and a steady state is reached, i.e., $N_{t+2} = N_{t+1}$, if

$$\mathbb{P}[S > Q_{t+1} | S > Q_t] = 1. \quad (5)$$

We now work toward simplifying the recursion on Q_t to obtain a better understanding of the condition Eq. (5) needed for cascading failures to stop. To this end, we apply the first relation in Eq. (4) repeatedly to see that

$$\begin{aligned} 1 - f_{t+1} &= (1 - f_t)\mathbb{P}[S > Q_t | S > Q_{t-1}], \\ 1 - f_t &= (1 - f_{t-1})\mathbb{P}[S > Q_{t-1} | S > Q_{t-2}], \\ &\vdots \\ 1 - f_2 &= (1 - f_1)\mathbb{P}[S > Q_1 | S > Q_0], \\ 1 - f_1 &= (1 - f_0)\mathbb{P}[S > Q_0]. \end{aligned}$$

Applying these recursively starting from the last equality, we find that

$$1 - f_{t+1} = (1 - f_0) \prod_{\ell=0}^t \mathbb{P}[S > Q_\ell | S > Q_{\ell-1}],$$

where we set $Q_{-1} = 0$ as before. Since Q_t is monotone increasing in t , i.e., $Q_{t+1} \geq Q_t$ for all t , we further obtain

$$\begin{aligned} 1 - f_{t+1} &= (1 - f_1) \frac{\mathbb{P}[S > Q_t]}{\mathbb{P}[S > Q_{t-1}]} \cdot \frac{\mathbb{P}[S > Q_{t-1}]}{\mathbb{P}[S > Q_{t-2}]} \cdots \frac{\mathbb{P}[S > Q_1]}{\mathbb{P}[S > Q_0]} \\ &= (1 - p)\mathbb{P}[S > Q_t]. \quad (6) \end{aligned}$$

This last expression confirms the intuitive result that the fraction of alive lines at stage $t + 1$ is simply given by the fraction of lines who survive the initial attack and have more free space than the extra load Q_t that is distributed on every alive line at stage t .

Using Eq. (6) in Eq. (4), it is now understood that the dynamics of cascading failures is fully governed and understood by the recursions on Q_t given by

$$Q_{t+1} = Q_t + (1 - \varepsilon) \frac{\mathbb{E}[(L + Q_t) \cdot \mathbf{1}[Q_{t-1} < S \leq Q_t]]}{\mathbb{P}[S > Q_t]} \quad (7)$$

for each $t = 0, 1, \dots$ [with Q_0 given at Eq. (3)], with the condition for reaching the steady-state still being Eq. (5). Let t^* be the stage at which steady-state is reached, i.e., the first t for which Eq. (5) holds. Then, the final system sizes $n_\infty(p, \varepsilon)$ defined as the fraction of alive lines at the steady state can be computed simply from [viz. Eq. (6)]

$$n_\infty(p, \varepsilon) = (1 - p)\mathbb{P}[S > Q_{t^*}]. \quad (8)$$

Throughout, we will be particularly interested in the critical attack size p^* defined as the largest attack that the system can sustain [in the sense of having a positive final size $n_\infty(p, \varepsilon)$]; i.e., for given ε we let

$$p^*(\varepsilon) = \sup\{p > 0 : n_\infty(p, \varepsilon) > 0\}.$$

To demonstrate the impact of the absorption factor ε together with the number of stages needed to reach a steady state, we find it useful to further simplify Eq. (7). By simple algebra, we get

$$\begin{aligned} Q_{t+1}\mathbb{P}[S > Q_t] &= Q_t\mathbb{P}[S > Q_t] + (1 - \varepsilon)Q_t\mathbb{P}[Q_{t-1} < S \leq Q_t] \\ &\quad + (1 - \varepsilon)\mathbb{E}[L \cdot \mathbf{1}[Q_{t-1} < S \leq Q_t]] \\ &= Q_t\mathbb{P}[S > Q_{t-1}] - \varepsilon Q_t\mathbb{P}[Q_{t-1} < S \leq Q_t] \\ &\quad + (1 - \varepsilon)\mathbb{E}[L \cdot \mathbf{1}[Q_{t-1} < S \leq Q_t]], \end{aligned}$$

which is equivalent to the following difference relation on the sequence $Q_{t+1}\mathbb{P}[S > Q_t]$:

$$\begin{aligned} Q_{t+1}\mathbb{P}[S > Q_t] - Q_t\mathbb{P}[S > Q_{t-1}] &= (1 - \varepsilon)\mathbb{E}[L\mathbf{1}[Q_{t-1} < S \leq Q_t]] \\ &\quad - \varepsilon Q_t\mathbb{P}[Q_{t-1} < S \leq Q_t]. \end{aligned} \quad (9)$$

We see from Eq. (9) that higher values of ε have suppressing effect on the growth of Q_t , leading to steady-state being reached *faster* (i.e., in small number of steps), and with a larger final system size in view of Eq. (8).

It is desirable to obtain a closed-form solution for Q_{t^*} by solving the difference Eq. (8); in view of Eq. (8) this would lead to a closed-form expression for the final system size $n_\infty(p, \varepsilon)$. However, applying Eq. (9) recursively leads to a telescoping sum given by

$$\begin{aligned} Q_{t+1}\mathbb{P}[S > Q_t] - Q_0 &= (1 - \varepsilon) \sum_{i=0}^t \mathbb{E}[L \cdot \mathbf{1}[Q_{i-1} < S \leq Q_i]] \\ &\quad - \varepsilon \sum_{i=0}^t Q_i\mathbb{P}[Q_{i-1} < S \leq Q_i]. \end{aligned} \quad (10)$$

It is now clear that unless $\varepsilon = 0$ or $\varepsilon = 1$, a direct expression for Q_t (for arbitrary t) can not be obtained without going through the recursion Eq. (9) and obtaining each one of Q_1, Q_2, \dots, Q_{t-1} . Therefore, it is also not possible to derive a closed-form expression for Q_{t^*} and $n_\infty(p, \varepsilon)$.

IV. NUMERICAL RESULTS

In this section, we confirm our theoretical findings via numerical simulations. We focus on three commonly known

distributions for the load and free-space variables: (i) Uniform, (ii) Pareto, and (iii) Weibull. The probability density functions corresponding to these distributions are given below for a generic random variable L .

(1) Pareto Distribution: $L \sim \text{Pareto}(L_{\min}, b)$. With $L_{\min} > 0$ and $b > 0$, the support set is $x \geq L_{\min}$ and the density is given by

$$p_L(x) = L_{\min}^b b x^{-b-1}.$$

We also enforce $b > 1$ to ensure that $\mathbb{E}[L] = bL_{\min}/(b-1)$ is finite. The Pareto family distributions are also known as *power-law* distributions and have been extensively used in many fields.

(2) Uniform Distribution: $L \sim U(L_{\min}, L_{\max})$. The support set is $L_{\min} \leq x \leq L_{\max}$ and the density is given by

$$p_L(x) = \frac{1}{L_{\max} - L_{\min}}.$$

(3) Weibull Distribution: $L \sim \text{Weibull}(L_{\min}, \lambda, k)$. With $\lambda, k, L_{\min} > 0$, the support set is $x \geq L_{\min}$ and the density is given by

$$p_L(x) = \frac{k}{\lambda} \left(\frac{x - L_{\min}}{\lambda} \right)^{k-1} e^{-\left(\frac{x - L_{\min}}{\lambda}\right)^k}.$$

The case $k = 1$ corresponds to the exponential distribution, and $k = 2$ corresponds to Rayleigh distribution. The mean load is given by $\mathbb{E}[L] = L_{\min} + \lambda\Gamma(1 + 1/k)$, where $\Gamma(\cdot)$ is the γ function $\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt$.

Next, we confirm our results presented in Sec. III concerning the response of the system to attacks of varying sizes; i.e., concerning the final system size $n_\infty(p, \varepsilon)$. We are particularly interested in the transition behavior around the critical attack size p^* . In all simulations, we fix the number of lines at $N = 10^6$, and for each set of parameters being considered (e.g., the distribution $p_{LS}(x, y)$ and attack size p) we run 200 independent experiments. In all figures below, the symbols represent the *empirical* value of the final system size $n_\infty(p, \varepsilon)$ (obtained from simulations by averaging over 200 independent runs for each data point), and solid lines represent the analytic results computed from Eq. (8) with Q_{t^*} obtained by iterating Eq. (7) while checking the condition Eq. (5) at each iteration step.

We start our numerical results with Pareto distribution. In Figs. 1 and 2, we let L_1, \dots, L_N be drawn from Pareto distribution with $b = 2$, $L_{\min} = 10$ and $S = 0.7L$. In Fig. 1, we plot the ratio of surviving lines $n_\infty(p, \varepsilon)$ as a function of the attack size p for various ε values. We already know from the analysis in Ref. [3] that for $\varepsilon = 0$, Pareto distribution always fosters an abrupt first-order transition behavior at p^* . We observe that this first order transition behavior continues to hold as ε is increased from 0 to 0.8. In fact, up until that point, the system's ability to absorb ε fraction of the failed load at each stage does not affect the final system size $n_\infty(p, \varepsilon)$. Only after $\varepsilon = 0.8$ the behavior of $n_\infty(p, \varepsilon)$ starts to change as shown in Fig. 1. For a complementary visualization, we plot the behavior of the ratio of surviving node $n_\infty(p, \varepsilon)$ as a function of ε for different values of initial attack size p . We see that the transition no longer fosters any sharp behavior and continuously improves as it reaches $\varepsilon = 1$. We next provide the mathematical justification for this observation.




FIG. 1. The breakdown of the system is demonstrated, where L_1, \dots, L_N are drawn from a Pareto distribution with $b = 2$, $L_{\min} = 10$. We also have $S = 0.7L$ so that S_1, \dots, S_N are such that $S_i = 0.7L_i$. We plot the ratio of surviving lines $n_\infty(p, \varepsilon)$ as a function of the attack size p for various ε values.

When $\varepsilon = 1$, an overloaded line will be removed from the system *without* its load being redistributed to the remaining lines. Put differently, the only time redistribution will take place is stage 1, where Q_0 gets redistributed to each of the $(1 - p)N$ lines that have not been attacked. Therefore, a line that is not in the initial attack will be included in the final system size as long as its free space is larger than Q_0 . This leads to having

$$n_\infty(p, \varepsilon) = (1 - p)\mathbb{P}[S > Q_0], \quad (11)$$

where $Q_0 = \frac{p}{1-p}\mathbb{E}[L]$. Therefore, $n_\infty(p, \varepsilon)$ is continuous in p whenever the marginal distribution of S is continuous, a condition satisfied in all distributions considered in the




FIG. 2. In the setting of Fig. 1, we plot the relative final size $n_\infty(p, \varepsilon)$ as a function of the loss or absorption factor parameter ε for different attack sizes p .




FIG. 3. The breakdown of the system is demonstrated, where L_1, \dots, L_N are drawn independently from a uniform distribution with $L_{\min} = 10$ and $\mathbb{E}[L] = 20$, and S_1, \dots, S_N are drawn independently from a uniform distribution with $S_{\min} = 10$ and $\mathbb{E}[S] = 35$. We plot the relative final size $n_\infty(p, \varepsilon)$ as a function of the attack size p for different ε .

numerical results. Therefore, in the case of perfect ε , $n(p, \varepsilon)$ hits zero and the system fully collapses in a continuous fashion. We further note that the distribution of S under fixed $\mathbb{E}[L]$ and $\mathbb{E}[S]$ can be selected so that the critical attack size p^* can be made arbitrarily close to 1. Selecting S as a binary random variable at $\{s_1, s_2\}$ with $s_1 < \mathbb{E}[S] < s_2$ independent from L such that $p_1s_1 + p_2s_2 = \mathbb{E}[S]$ is sufficient to see this phenomenon. For fixed $s_1 < \mathbb{E}[S]$, one can select s_2 arbitrarily large with $p_1, p_2 \neq 0$ such that $p_1s_1 + p_2s_2 = \mathbb{E}[S]$ and this proves that $\mathbb{P}[S > Q_0]$ can be made nonzero *irrespective of the value of the initial attack size* p , i.e., p^* can be made arbitrarily close to 1.

In Fig. 3, we plot the ratio of surviving lines $n_\infty(p, \varepsilon)$ with respect to p when L_1, \dots, L_N are drawn from the load L and extra space S is independent and uniformly distributed with $L_{\min} = S_{\min} = 10$ and $\mathbb{E}[L] = 20$ and $\mathbb{E}[S] = 35$. We observe that the transition behavior has a failure with preceding divergence and the critical threshold p^* migrates from $p^*(\varepsilon) = 0.375$ at $\varepsilon = 0$ to $p^*(\varepsilon) = 0.75$.

In Fig. 4, we plot the ratio of surviving lines $n_\infty(p, \varepsilon)$ with respect to p when L_1, \dots, L_N are drawn from Weibull distribution with $k = 0.4$, $\lambda = 100$, $L_{\min} = 10$, $S = 1.74L$. We see that the Weibull distribution gives rise to a richer set of possibilities for the transition of $n_\infty(p, \varepsilon)$. Namely, we see that an abrupt rupture, a rupture with preceding divergence as well as a first-order transition followed by a second-order transition that is followed by an ultimate first-order breakdown are all possible in this case. As the parameter ε is increased, the transition behavior gets smoother.

We finally examine the phase diagrams corresponding to cases presented above and reveal the emergence of tricritical points. In Fig. 5, we plot the loss factor ε and attack size p pairs for which the phase transition occurs in either first or second order. We denote a first-order transition by straight line




FIG. 4. The breakdown of the system is demonstrated, where L_1, \dots, L_N are drawn from a Weibull distribution with $k = 0.4$, $\lambda = 100$, $L_{\min} = 10$, and free space for each lines is given by $S = 1.74L$. We plot the ratio of surviving lines $n_{\infty}(p, \varepsilon)$ as a function of the attack size p .

and a second-order transition by dashed line. In Fig. 5, we refer to the case of $L \sim \text{Pareto}(2,10)$ and $S = 0.7L$ as Case 1, the case of independent L, S with $L \sim \text{Uniform}([10,30])$ and $S \sim \text{Uniform}([10,60])$ as Case 2 and the case of $L \sim \text{Weibull}(10,100,0.4)$ and $S = 1.74L$ as Case 3. We observe that Case 1 fosters a very sharp phase diagram in that a first-order phase transition occurs at the same attack size for all loss factors smaller than a certain ε value, and a second-order transition is seen after that ε is exceeded. In contrast, Case 3 has a smoother phase diagram as the phase transition switches from first order to second order after a certain tricritical point




FIG. 5. The phase diagrams and the resulting tricritical behavior is shown for Cases 1, 2 and 3. For each value of the loss or absorption factor ε , we show the corresponding *critical* attack size p . Solid curves represent cases where the transition at the corresponding critical attack size is first order, while dashed curves stand for cases where a second-order phase transition occurs.




FIG. 6. We plot the ratio of surviving lines $n_{\infty}(p, \varepsilon)$ as a function of the attack size p in the setting when L and S are independent and distributed uniformly as $U[0, 120]$ and $U[10, 60]$, respectively, as well as the setting when S is Dirac- δ distributed as $\delta(x - \mathbb{E}[S])$.

[6]. We also observe that the transition in Case 2 is always of first order and the corresponding phase diagram is smooth.

V. OPTIMIZING ROBUSTNESS

A. Quantifying robustness

A typical metric to assess network robustness is the percolation threshold p^* at which the system fully collapses as a result of the cascading failures [3,42]. In particular, we know from earlier works [3,42] that for $\varepsilon = 0$, $p^*_{\text{optimal}} = \frac{\mathbb{E}[S]}{\mathbb{E}[S] + \mathbb{E}[L]}$, and it is achieved by a Dirac- δ distribution at $\mathbb{E}[S]$. We observed in the numerical results in Sec. IV that the presence of the parameter ε changes p^* . For example, we have seen that if more load is absorbed from the overloaded lines by increasing ε , p^* may increase or decrease with respect to the case $\varepsilon = 0$. In cases when p^* increases $n_{\infty}(p, \varepsilon)$ decreases. It is, therefore, difficult to assess the robustness using only the p^* metric; e.g., see Fig. 6 where the distribution that leads to the highest p^* does not maximize the final system size $n_{\infty}(p, \varepsilon)$ uniformly across all attack sizes $0 \leq p \leq 1$.

To quantify the overall robustness of the network under all possible attack sizes, we consider a metric that measures the area under $n_{\infty}(p, \varepsilon)$ over $0 \leq p \leq 1$. Namely, we let

$$\mathcal{R}(\varepsilon) = \int_0^1 n_{\infty}(p, \varepsilon) dp. \quad (12)$$

The metric $\mathcal{R}(\varepsilon)$ was introduced in Ref. [31], and can be seen in Ref. [43] to represent the *expected* final system size in response to an attack whose size p is selected uniformly at random over $[0,1]$. It is in this spirit that the metric $\mathcal{R}(\varepsilon)$ quantifies the overall system robustness under a range of attack sizes for fixed ε . Alternative metrics can also be defined where the attack size p is drawn from an arbitrary distribution $F(p)$, e.g., to account for the fact that attacks of certain size might be more likely than others. In that case, we would again

compute the mean of the final system size, i.e.,

$$\mathbb{E}_p[n_\infty(p, \varepsilon)] = \int_0^1 n_\infty(p, \varepsilon) dF(p). \quad (13)$$

Clearly, $\mathcal{R}(\varepsilon)$ defined in Eq. (12) is recovered when $F(p)$ is the uniform distribution over $[0, 1]$.

B. Maximizing the robustness metric $\mathcal{R}(\varepsilon)$

Our next result conclusively establishes that for any ε , the Dirac- δ distribution [44] for free space S optimizes the robustness of the system with respect to the metric $\mathcal{R}(\varepsilon)$ among all possible distributions $P_{LS}(x, y)$ with fixed $\mathbb{E}[S]$ and $\mathbb{E}[L]$. First, we note that if $p_{LS}(x, y) = p_L(x)\delta(y - \mathbb{E}[S])$, then the final system size is independent of ε . This is because after the initial attack, either all lines will fail [if $\mathbb{E}[S] \leq p\mathbb{E}[L]/(1-p)$], or they will all survive and the cascades will not continue. Thus, the final system size under the Dirac- δ distribution of free space is given [42] by

$$n_{\delta, \infty}(p) = \begin{cases} 1-p & \text{if } p < p_\delta^* \\ 0 & \text{if } p \geq p_\delta^* \end{cases}, \quad (14)$$

where, the critical attack size p_δ^* is given by

$$p_\delta^* = \frac{\mathbb{E}[S]}{\mathbb{E}[S] + \mathbb{E}[L]}.$$

We will show that the distribution $p_{LS}(x, y) = p_L(x)\delta(y - \mathbb{E}[S])$ maximizes the robustness metric $\mathcal{R}(\varepsilon)$ among all p_{LS} with mean values for L and S fixed at $\mathbb{E}[L]$ and $\mathbb{E}[S]$, respectively. In view of Eq. (14), this will follow if we show that

$$\int_0^1 n_\infty(p, \varepsilon, p_{LS}) dp \leq \int_0^{p_\delta^*} (1-p) dp, \quad (15)$$

where $n_\infty[p, \varepsilon, p_{LS}(x, y)]$ denotes the final system size under attack size p , when load and free-space values of the lines are generated independently from the distribution $p_{LS}(x, y)$ (with fixed $\mathbb{E}[L]$, $\mathbb{E}[S]$).

From Eq. (8) we note that

$$\begin{aligned} n_\infty[p, \varepsilon, p_{LS}(x, y)] &\leq (1-p)\mathbb{P}[S \geq Q_0] \\ &= (1-p)\mathbb{P}\left[S \geq \frac{p}{1-p}\mathbb{E}[L]\right], \end{aligned} \quad (16)$$

due to the fact that $Q_0 \leq Q_t$ for all $t = 1, 2, \dots$. Therefore, we will get the desired result Eq. (15) if we show that

$$\int_0^1 (1-p)\mathbb{P}\left[S \geq \frac{p}{1-p}\mathbb{E}[L]\right] dp \leq \int_0^{p_\delta^*} (1-p) dp,$$

or equivalently, that

$$\begin{aligned} &\int_{p_\delta^*}^1 (1-p)\mathbb{P}\left[S \geq \frac{p}{1-p}\mathbb{E}[L]\right] dp \\ &\leq \int_0^{p_\delta^*} (1-p)\mathbb{P}\left[S < \frac{p}{1-p}\mathbb{E}[L]\right] dp. \end{aligned} \quad (17)$$

Since $1-p$ is monotone decreasing over the range $0 \leq p \leq 1$, and both $\mathbb{P}[S \geq \frac{p}{1-p}\mathbb{E}[L]]$ and $\mathbb{P}[S < \frac{p}{1-p}\mathbb{E}[L]]$ are

nonnegative, Eq. (17) will follow if we show that

$$\int_{p_\delta^*}^1 \mathbb{P}\left[S \geq \frac{p}{1-p}\mathbb{E}[L]\right] dp \leq \int_0^{p_\delta^*} \mathbb{P}\left[S < \frac{p}{1-p}\mathbb{E}[L]\right] dp,$$

or equivalently, that

$$\int_0^1 \mathbb{P}\left[S \geq \frac{p}{1-p}\mathbb{E}[L]\right] dp \leq \int_0^{p_\delta^*} dp = \frac{\mathbb{E}[S]}{\mathbb{E}[S] + \mathbb{E}[L]}. \quad (18)$$

To establish Eq. (18), we make a change of variables $x = \frac{p}{1-p}\mathbb{E}[L]$ and write

$$\begin{aligned} &\int_0^1 \mathbb{P}\left[S \geq \frac{p}{1-p}\mathbb{E}[L]\right] dp \\ &= \int_0^\infty \mathbb{P}[S \geq x] d\left(\frac{x}{x + \mathbb{E}[L]}\right) \\ &= \mathbb{P}[S \geq x] \frac{x}{x + \mathbb{E}[L]} \Big|_{x=0}^\infty - \int_0^\infty \frac{x}{x + \mathbb{E}[L]} d(\mathbb{P}[S \geq x]) \\ &= \mathbb{E}\left[\frac{S}{S + \mathbb{E}[L]}\right] \\ &\leq \frac{\mathbb{E}[S]}{\mathbb{E}[S] + \mathbb{E}[L]}, \end{aligned} \quad (19)$$

where we use integration by parts in Eq. (19) and apply Jensen's inequality in Eq. (20) for the function $\frac{x}{x + \mathbb{E}[L]}$ that is concave in x . This establishes Eq. (18) and the desired result Eq. (15) follows in view of the preceding arguments.

This result shows that the system's robustness with respect to the metric $\mathcal{R}(\varepsilon)$ [defined at Eq. (12)] is maximized under the constraints of fixed $\mathbb{E}[L]$ and fixed $\mathbb{E}[S]$ (and hence fixed $\mathbb{E}[C]$), by giving each line an equal free space $\mathbb{E}[S]$, *irrespective of how the initial loads are distributed and the redistribution process in the later stages*. In other words, the robustness is maximized by choosing a line's capacity C_i through $C_i = L_i + \mathbb{E}[S]$ no matter what its load L_i is.

Note that in light of the results in Sec. IV, the critical attack size $p^*(\varepsilon)$ may be greater or less than $\frac{\mathbb{E}[S]}{\mathbb{E}[S] + \mathbb{E}[L]}$ depending on the specific distribution p_{LS} used (with fixed $\mathbb{E}[L]$ and $\mathbb{E}[S]$). In particular, we have seen that it is always possible to choose a marginal distribution with $\mathbb{E}[S]$ such that p^* is arbitrarily close to 1. Therefore, one might think that the area under $n_\infty(p, \varepsilon)$ while swiping all possible p will be maximized when ε is increased and the system has the capability to absorb the extra load coming from the failing lines and eradicate the potentially detrimental effect of their failure to the overall system. Our result shows firmly that this intuition is incorrect and the metric $\mathcal{R}(\varepsilon)$ is instead maximized when the distribution of S is the Dirac- δ function centered at $\mathbb{E}[S]$ (irrespective of ε and the distribution of L).

We note that our argument follows from the facts that the extra load due to the initial attack $Q_0 = \frac{p}{1-p}\mathbb{E}[L]$ is monotone increasing, continuous, and convex in the initial attack size p . These properties are expected to hold in a large set of instances of this problem. Therefore, the optimality of the Dirac- δ distribution of S is likely to hold under more general cases where these properties hold. For instance, if the prior randomness $F(p)$ on the attack size p in Eq. (13) has a

monotone decreasing derivative, i.e., if higher attack sizes are *less likely*, then the Dirac- δ distribution of free space S is still optimal with respect to the resulting metric in Eq. (13). Such cases occur in situations where the malicious attacker is more likely to choose smaller attack sizes due to resource or time constraints. We similarly observe that if the support of the prior distribution of p is contained in the interval $[0, \frac{\mathbb{E}[S]}{\mathbb{E}[S]+\mathbb{E}[L]}]$, then the Dirac- δ function at $\mathbb{E}[S]$ is optimal with respect to $\mathcal{R}(\varepsilon)$. Finally, the optimality of the Dirac- δ distribution of S [in the sense of maximizing $\mathcal{R}(\varepsilon)$] holds irrespective of the value of ε . As such, this optimality prevails under any time variation in ε or possible dependence of ε on the instantaneous extra load per line Q_t .

While it is hard to generalize the optimality of Dirac- δ function for $0 < \varepsilon < 1$ and any prior distribution on p , we observe that if p is deterministic and $\varepsilon = 1$, then the robustness metric in Eq. (13) is given by

$$\mathbb{E}_p[n_\infty(p, \varepsilon)] = n_\infty(p, 1) = (1 - p)\mathbb{P}[S > Q_0]. \quad (21)$$

Then, to maximize $n_\infty(p, \varepsilon)$, it suffices to minimize $\mathbb{P}[S > Q_0] = \mathbb{P}[S > p\mathbb{E}[L]/(1 - p)]$ subject to $\mathbb{E}[S]$ being fixed. If the known attack size satisfies $p \leq \frac{\mathbb{E}[S]}{\mathbb{E}[S]+\mathbb{E}[L]}$, then this is achieved when the distribution of S is given by the Dirac- δ function centered at $\mathbb{E}[S]$. If, however, we have $p > \frac{\mathbb{E}[S]}{\mathbb{E}[S]+\mathbb{E}[L]}$, then the optimal distribution of S consists of two Dirac- δ functions centered at 0 and $\frac{p}{1-p}\mathbb{E}[L]$, respectively, with appropriate probabilities selected such that the mean value is $\mathbb{E}[S]$. In other words, if the attack size is larger than what can be resisted by giving each line an equal amount $\mathbb{E}[S]$ of free space, then it is optimal to give a fraction of lines *zero* free space, while giving each of the other lines an equal amount of $\frac{p}{1-p}\mathbb{E}[L]$ free space; i.e., just enough so that they can handle the additional load of Q_0 that will be distributed to them after the attack.

C. Simulations with synthetic and real-world data

In Fig. 6, we plot the fraction of surviving lines $n_\infty(p, \varepsilon)$ as a function of the attack size p for several ε values, in the setting when L and S are independent and uniformly distributed as $U[0, 120]$ and $U[10, 60]$, respectively. We compare the corresponding final system size $n_\infty(p, \varepsilon)$ when S is Dirac δ distributed as $\delta(x - \mathbb{E}[S])$. We observe that the area under $n_\infty(p, \varepsilon)$ is larger under the Dirac- δ distribution compared to other cases. In Fig. 7, we provide a comparison of the metric $\mathcal{R}(\varepsilon)$ for a family of Weibull distributions with fixed $\mathbb{E}[L]$ and $\mathbb{E}[S]$ while varying the scale parameter k of the distribution. We observe that $\mathcal{R}(\varepsilon)$ is monotone increasing in k and is maximized in the limiting case $k \rightarrow \infty$. This observation is in perfect agreement with our result given that the Weibull distribution approaches to a Dirac- δ function as k goes to infinity.

We now test our analytic results concerning the optimization of robustness on real-world power system data from the UK National Grid [36]. In particular, we pick $N = 75$ loads L_1, \dots, L_{75} from the Great Britain two degree power flow diagram for 2017/18 shown explicitly in Fig. C.1 in Ref. [[36], Appendix C]. These power flows represent the power units in




FIG. 7. L_1, \dots, L_N are drawn from Weibull distribution with $L_{\min} = 1$ and k, λ such that $\mathbb{E}[L] = 2$ and $S = 1.74L$. We plot the metric $\mathcal{R}(\varepsilon)$ in Eq. (12) as a function of the parameter k .

MVA that flow through lines in the utility grid in the scale shown in Ref. [[36], Appendix C].

Our goal is to compare the robustness of this power system, as quantified through the area metric, under different allocations of line capacities. In particular, we will compare the performance of the equal free-space allocation with respect to the schemes described below:

(I) Proportional free-space allocation: For each line i , we set $S_i = \alpha L_i$, where α is the *tolerance* factor, set to be the same across all N lines.

(II) Equal capacity allocation: $S_i + L_i = C$ for all i if $L_i < C$ and $S_i = 0$ otherwise.

In order for this comparison to be fair, the total free space available to all 75 lines will be fixed; since loads are already given from the UK Power-grid data set, this will amount to having the total capacity of 75 lines fixed. The total load of the 75 lines given in this data set is 60 221 power units. We assume that a total free space of 168 620 power units is available to these lines and will be allocated according to one of three free-space allocation schemes to be considered. For the proportional free-space allocation, this leads to using a tolerance-factor of $\alpha = 2.8$ for all lines, while for the equal capacity allocation scheme, this leads to fixing the capacity of all lines at $C = 3 051$ power units. In the case of equal free-space allocation, this amounts to assigning the capacity of every line at their load plus $168\,620/75 \simeq 2\,248$ units of free space. These three schemes are compared with respect to the resulting area robustness metric defined here as

$$\mathcal{R}(\varepsilon) = \frac{1}{N} \sum_{i=1}^N n_\infty(p_i, \varepsilon), \quad (22)$$

where $p_i = \frac{L_i}{N}$.

The results are presented in Fig. 8, where each data point corresponds to the average over 5 000 independent experiments. We observe that the equal free-space allocation performs significantly better than the other two capacity allocation schemes, in the sense leading to a larger value of the




FIG. 8. Comparison of robustness performances of equal free-space allocation, equal capacity allocation and proportional free-space allocation for $N = 75$ lines with load data obtained from the UK National Grid [[36], Appendix C]. We observe that the robustness is better for equal free-space allocation with respect to the other two benchmarks.

robustness metric Eq. (22). For instance, when a random attack occurs whose scale p is selected uniformly at random from $(0,1)$, we see that close to 50% of the lines are expected to survive under the optimal choice of equal free-space allocation. We remind that under Eq. (22), the metric $\mathcal{R}(\varepsilon)$ is always less than 0.5, meaning that the performance under the equal free-space allocation is close to the theoretically possible maximum robustness level. This statement holds irrespective of the absorption factor ε and in particular even if the system has no capability of absorbing the failed load. However, if we use the widely assumed setting of free space being proportional to load, i.e., $C_i = (1 + \alpha)L_i$, then we can expect only about 35% of the nodes to survive the same scenario if the system has no load absorption capability. We see that even if the system is designed with full load absorption capability (i.e., $\varepsilon = 1$), the expected number of surviving nodes increases only to about 41%, still worse than the optimal case. The equal-capacity allocation leads to a higher robustness than the proportional free-space case, but the achieved robustness is still inferior to the case of where all lines have the same free space. It is only when the system has almost perfect flow absorption capability (i.e., $\varepsilon = 1$) that the performances of equal-capacity and equal free-space allocations become similar; for small ε we see that equal free space is significantly better. Although not reported here for brevity, our extensive simulation results indicate that similar conclusions apply for a wide range of α and C values.

D. Simulations under a topology-based redistribution model

The results presented in earlier sections of this paper were based on a model with a mean field assumption in that when a line fails, its load gets redistributed globally and equally among all active lines in the network. This mean field assumption may not hold for networks with possible local redistribution behavior such as power networks. However, it is well known that redistribution models with only local components fail to capture the long-range nature of Kirchhoff's Law.

Consequently, a model where the failed load is redistributed both locally and globally would be a more suitable one. To this end, we present in this section simulation results under a topology based redistribution model. We do so with an eye towards revealing whether the optimality of equal free-space distribution prevails when failed load is redistributed (at least in part) locally according to a network topology.

In the spirit of the redistribution models presented in Refs. [32,34], we consider a topology-based redistribution model that combines local and global redistribution behaviors. This extended model gauges the locality of the redistribution by the parameter $\gamma \in [0, 1]$ and the network topology is generated as an Erdős-Rényi graph $\mathbb{G}(n, N)$. At each stage, the portion of the load that is not absorbed from each failed line is divided into two parts: γ fraction is redistributed locally among neighboring lines (with each neighbor receiving an equal portion), and $(1 - \gamma)$ fraction is redistributed equally among *all* surviving lines (irrespective of topology). In this model, setting $\gamma = 0$ recovers the mean-field model introduced in Sec. II, while setting $\gamma = 1$ gives a merely topology based redistribution model.


For this new redistribution model, we have run simulations where we set the number of nodes as $n = 250$ and number of edges as $N = 8400$, and created a random network according to the Erdős-Rényi $\mathbb{G}(n, N)$ model. This leads to each line having on average 132 neighbors. The loads $\{L_1, \dots, L_N\}$ are drawn from i.i.d. Weibull distribution (with $L_{\min} = 1$ and k, λ such that $\mathbb{E}[L] = 2$) and the free spaces are set as $S_i = 1.74L_i$. We run 100 independent experiments for each parameter set, and report the average value of the robustness metric defined in Eq. (22).

The results are presented in Figs. 9(a), 9(b), and 9(c) for $\gamma = 0.25$, $\gamma = 0.6$, and $\gamma = 1$, respectively; we remark that the setting in Fig. 7 corresponds to the case with $\gamma = 0$. The most important observation is that Figs. 9(a)–9(c) obtained under (partially) *local* load redistribution are qualitatively the same with the results given in Fig. 7 obtained under mean-field analysis (i.e., global redistribution). For instance, we observe in all figures that as the parameter k is increased, the robustness of the network increases monotonically with Dirac- δ distribution of free space S leading the maximum robustness. Thus, we conclude that the optimality of equal free-space allocation prevails even when the redistribution of the loads on failing lines are performed based on the topology of the network. Another observation we make from Figs. 9(a), 9(b), and 9(c) with comparison to Fig. 7 is that increasing the locality parameter γ decreases robustness. This is in line with our understanding that introducing heterogeneity in the redistribution model causes further loss in robustness.

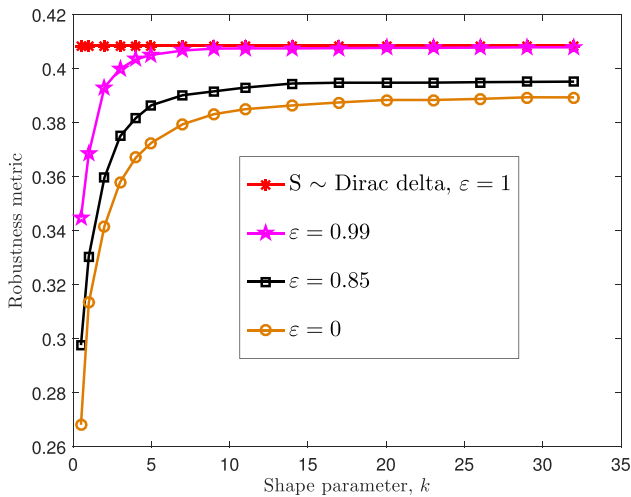
We finally test the effect of node degree on the robustness of the network. We consider the Erdős-Rényi $\mathbb{G}(n, N)$ model with the same number of nodes $n = 250$ and a significantly smaller number of edges $N = 625$. In this case, each line has 8 neighbors on average. The results are presented in Figs. 10(a), 10(b), and 10(c) for $\gamma = 0.25$, $\gamma = 0.6$, and $\gamma = 1$, respectively. We observe the same qualitative behavior with Fig. 9 in that as the shape parameter of the Weibull distribution gets larger, the robustness of the network improves irrespective of locality parameter γ and the absorption parameter ε . Additionally, we observe that robustness is affected by the



(a)




(b)




(c)


FIG. 9. We let L_1, \dots, L_N be obtained from Weibull distribution with $L_{\min} = 1$ and k, λ such that $\mathbb{E}[L] = 2$ and $S = 1.74L$. For (a) $\gamma = 0.25$, (b) $\gamma = 0.6$, and (c) $\gamma = 1.0$, the plots show the variation of the metric $\mathcal{R}(\varepsilon)$ in Eq. (22) with the parameter k for various ε values.



(a)



(b)



(c)

FIG. 10. We repeat the experiments in Fig. 9 for (a) $\gamma = 0.25$, (b) $\gamma = 0.6$, and (c) $\gamma = 1.0$, under a different Erdős-Rényi graph with significantly smaller average number of neighbors.

connectivity of the network in different ways for different γ and ε . In particular, robustness significantly drops as the average number of neighbors decreases when ε is smaller while the effect of the number of neighbors on robustness is limited for higher ε . This numerical observation points to an intricate relation among load absorption, the average number of neighbors and locality in the load redistribution process for their effects on survival and robustness against cascading failures. Overall, the numerical results in Figs. 9 and 10 confirm the promise of uniformly allocating the free spaces to each node and gives reason to extend our results to more general network and redistribution models [45].

VI. CONCLUSION

We studied the robustness of flow networks consisting of N lines against random attacks under a partial load redistribution model. In particular, when a line fails due to overloading, it is removed from the system and $(1 - \varepsilon)$ fraction of the load it was carrying gets redistributed equally among all remaining lines while the remaining ε fraction is assumed to be lost or absorbed. We derive recursive relations describing the dynamics of cascading failures for any attack size p , and identify the *final* fraction of surviving lines when the cascades stop. These findings are confirmed via extensive simulations. Among other things, we show that unlike the full redistribution case (i.e., $\varepsilon = 0$), partial redistribution might lead to the order of transition at the critical attack size p^* changing from first to second order, and a tricritical point emerges with respect to attack size p and absorption/loss factor ε .

One of the most interesting findings of this paper is concerned with how system robustness can be *maximized* by properly choosing the distribution p_{LS} that generates the initial load and free-space values of each line. We consider this problem when the mean load $\mathbb{E}[L]$ and mean free space $\mathbb{E}[S]$ are fixed. First, we show that unlike the full redistribution case (i.e., when $\varepsilon = 0$), the critical attack size p^* is not necessarily maximized by assigning every line the same free space $\mathbb{E}[S]$; depending on the fraction ε of the load that is absorbed at each stage, we see that distributions other than Dirac- δ for S may lead to higher critical points p^* . Next, we consider the robustness metric proposed in [31] that computes the *area* under the final system size $n_\infty(p, \varepsilon)$ over all possible attack sizes

$0 \leq p \leq 1$; this amounts to computing the average response of the network to initial attacks of different size. We show that the system is most robust in the sense that the area metric is maximized, when the variation among the free space of lines is minimized. In other words, the Dirac- δ distribution of free space leads to the optimum robustness, irrespective of ε and how load L is distributed. Additionally, we test the robustness performance with respect to this area metric in finite flow networks with load data obtained from a real power network. We observe that allocating the free space equally among the lines yields significantly higher performance when compared to well-known benchmarks despite the small size chosen for the network. Finally, we tested this robustness result in an extended topology-based redistribution model over an Erdős-Rényi graph and we observed that uniformly allocating free spaces promises to be optimal in this more general setting as well.

There are many open problems one can consider for future work. For instance, the analysis can be extended to the case where the redistribution parameter ε is not the same for all lines, but follows a given probability distribution. Similarly, ε could be a time-varying parameter or it could depend on the extra load per line in the current stage. Such possibilities would allow us to obtain further understanding on the dynamical properties of the cascading failures and the mechanisms that could lead to a smooth failure. Additionally, it would be interesting to see if the robustness is still maximized with a Dirac- δ -type distribution on the free space under various combinations of possibilities on ε . Variability in the redistribution as in [46] and stochastic behaviour in the loads as in Ref. [47] also constitute interesting avenues for future research. In light of these directions, we will investigate relations among ε , the average number of neighbors and locality in redistribution for their effects on robustness against cascading failures. Finally, it would be interesting to study the partial redistribution model under *targeted* attacks rather than random failures.

ACKNOWLEDGMENTS

This research was supported in part by National Science Foundation Grants CCF No. 1422165 and CCF No. 1646526, by the Department of Energy Grant No. DE-OE0000779, and by the Department of Electrical & Computer Engineering at Carnegie Mellon University.

-
- [1] S. Tang, J. Yuan, X. Mao, X.-Y. Li, W. Chen, and G. Dai, in *Proceedings of IEEE INFOCOM* (IEEE, Shanghai, China, 2011), pp. 2291–2299.
 - [2] H. Wu, A. Arenas, and S. Gómez, *Phys. Rev. E* **95**, 012301 (2017).
 - [3] O. Yağan and V. Gligor, *Phys. Rev. E* **86**, 036103 (2012).
 - [4] Y. Zhuang and O. Yağan, *IEEE Trans. Netw. Sci. Eng.* **3**, 211 (2016).
 - [5] O. Yağan, D. Qian, J. Zhang, and D. Cochran, *IEEE J. Sel. Areas Commun.* **31**, 1038 (2013).
 - [6] R. Parshani, S. V. Buldyrev, and S. Havlin, *Phys. Rev. Lett.* **105**, 048701 (2010).
 - [7] S.-W. Son, G. Bizhani, C. Christensen, P. Grassberger, and M. Paczuski, *Europhys. Lett.* **97**, 16006 (2012).
 - [8] B. Min, S. D. Yi, K.-M. Lee, and K.-I. Goh, *Phys. Rev. E* **89**, 042811 (2014).
 - [9] K.-M. Lee, C. D. Brummitt, and K.-I. Goh, *Phys. Rev. E* **90**, 062816 (2014).
 - [10] C. Wu, S. Ji, R. Zhang, L. Chen, J. Chen, X. Li, and Y. Hu, *Europhys. Lett.* **107**, 48001 (2014).
 - [11] F. Radicchi, *Nat. Phys.* **11**, 597 (2015).
 - [12] O. Yağan, D. Qian, J. Zhang, and D. Cochran, *IEEE Trans. Parallel Distrib. Syst.* **23**, 1708 (2012).
 - [13] X. Huang, J. Gao, S. Buldyrev, S. Havlin, and H. E. Stanley, *Phys. Rev. E* **83**, 065101 (2011).

- [14] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin, *Nature* **464**, 1025 (2010).
- [15] W. Li, A. Bashan, S. V. Buldyrev, H. E. Stanley, and S. Havlin, *Phys. Rev. Lett.* **108**, 228702 (2012).
- [16] J. Gao, S. V. Buldyrev, S. Havlin, and H. E. Stanley, *Phys. Rev. Lett.* **107**, 195701 (2011).
- [17] C. D. Brummitt, R. M. D'Souza, and E. Leicht, *Proc. Natl. Acad. Sci. USA* **109**, E680 (2012).
- [18] I. Dobson, B. A. Carreras, V. E. Lynch, and D. E. Newman, *Chaos: Interdisc. J. Nonlin. Sci.* **17**, 026103 (2007).
- [19] M. Elliott, B. Golub, and M. O. Jackson, *Amer. Econ. Rev.* **104**, 3115 (2014).
- [20] G. Zhang, Z. Li, B. Zhang, and W. A. Halang, *Physica A: Stat. Mech. Appl.* **392**, 3273 (2013).
- [21] S. Pahwa, C. Scoglio, and A. Scala, *Sci. Rep.* **4**, 3694 (2014).
- [22] H. Daniels, *Proc. R. Soc. London, Ser. A* **183**, 405 (1945).
- [23] J. V. Andersen, D. Sornette, and K.-T. Leung, *Phys. Rev. Lett.* **78**, 2140 (1997).
- [24] Y. Yang, T. Nishikawa, and A. E. Motter, *Science* **358**, eaan3184 (2017).
- [25] Y. Yang and A. E. Motter, *Phys. Rev. Lett.* **119**, 248302 (2017).
- [26] O. Yağan, *Phys. Rev. E* **91**, 062811 (2015).
- [27] W. Curtin, *J. Mech. Phys. Solids* **41**, 217 (1993).
- [28] F. Kun, M. Costa, R. Costa Filho, J. Andrade Jr, J. Soares, S. Zapperi, and H. Herrmann, *J. Stat. Mech.: Theory Exp.* (2007) P02003.
- [29] D. Cohen, P. Lehmann, and D. Or, *Water Resour. Res.* **45**, W10436 (2009).
- [30] For the model studied in this paper, the effect of increasing ε can be thought as a reduction in the dependency (or coupling) among the N lines forming the flow network.
- [31] C. Schneider, A. Moreira, J. Andrade Jr., S. Havlin, and H. Herrmann, *Proc. Natl. Acad. Sci. USA* **108**, 3838 (2011).
- [32] A. E. Motter and Y.-C. Lai, *Phys. Rev. E* **66**, 065102 (2002).
- [33] W.-X. Wang and G. Chen, *Phys. Rev. E* **77**, 026101 (2008).
- [34] B. Mirzasoleiman, M. Babaei, M. Jalili, and M. A. Safari, *Phys. Rev. E* **84**, 046114 (2011).
- [35] P. Crucitti, V. Latora, and M. Marchiori, *Phys. Rev. E* **69**, 045104 (2004).
- [36] U.K. National Power Grid Data (2017), technical notes available online at <https://www.nationalgrid.com/uk/publications/electricity-ten-year-statement-etys>.
- [37] U.S.-Canada Power System Outage Task Force (2004), Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations. Available online at <https://www.energy.gov/sites/prod/files/oeprod/DocumentsandMedia/BlackoutFinal-Web.pdf>.
- [38] P. Palensky and D. Dietrich, *IEEE Trans. Industr. Info.* **7**, 381 (2011).
- [39] D. Sornette, K.-T. Leung, and J. V. Andersen, arXiv preprint [arXiv:cond-mat/9712313](https://arxiv.org/abs/cond-mat/9712313) (1997).
- [40] S. Pradhan and B. K. Chakrabarti, *Int. J. Mod. Phys. B* **17**, 5565 (2003).
- [41] C. Roy, S. Kundu, and S. S. Manna, *Phys. Rev. E* **91**, 032103 (2015).
- [42] Y. Zhang and O. Yağan, *Sci. Rep.* **6** (2016).
- [43] The metric in Ref. [[31], Eq. (1)] is introduced as the sum of surviving nodes over all possible numbers of attacked lines. One can show by the almost sure convergence of the probabilities to the fraction of line failures due to the law of large numbers and the bounded convergence theorem that the metric $\mathcal{R}(\varepsilon)$ in Eq. (12) represents the same metric in Ref. [31] as the number of lines N grows.
- [44] We note that Dirac- δ distribution can be viewed as a limiting distribution for a family of continuous distributions and therefore it can be treated within the framework of our paper that assumed continuous $P_{LS}(x, y)$.
- [45] Our intuition is that equal free-space allocation will be optimal under any model where the (load, free-space values) of lines are drawn independently from a common distribution. After all, it should be the higher connectivity or higher load of neighboring lines that might warrant allocating a line a higher free space, instead of the load of the line itself as suggested by the linear free-space allocation scheme $S = \alpha * L$.
- [46] R. C. Hidalgo, Y. Moreno, F. Kun, and H. J. Herrmann, *Phys. Rev. E* **65**, 046148 (2002).
- [47] D.-H. Kim and A. E. Motter, *New J. Phys.* **10**, 053022 (2008).