

Attack vulnerability of power systems under an equal load redistribution model

Talha Cihad Gulcu, Vaggos Chatziafratis, Yingrui Zhang, and Osman Yağan

Abstract—This paper studies the vulnerability of flow networks against adversarial attacks. In particular, consider a power system (or, any system carrying a physical flow) consisting of N transmission lines with initial loads L_1, \dots, L_N and capacities C_1, \dots, C_N , respectively; the capacity C_i defines the maximum flow allowed on line i . Under an equal load redistribution model, where load of failed lines is redistributed equally among all remaining lines, we study the *optimization* problem of finding the best k lines to attack so as to minimize the number of *alive* lines at the steady-state (i.e., when cascades stop). This is done to reveal the worst-case attack vulnerability of the system as well as to reveal its most vulnerable lines. We derive optimal attack strategies in several special cases of load-capacity distributions that are practically relevant. We then consider a modified optimization problem where the adversary is also constrained by the *total* load (in addition to the number) of the initial attack set, and prove that this problem is NP-Hard. Finally, we develop heuristic algorithms for selecting the attack set for both the original and modified problems. Through extensive simulations, we show that these heuristics outperform benchmark algorithms under a wide range of settings.

Index Terms—Flow networks; cascading failures; vulnerability; optimal attack strategies

I. INTRODUCTION

Networks that carry or transport physical commodities, e.g., electricity/water/gas distribution networks and road/public transportation networks, have been an integral part of our daily lives for decades. For example, our quality of life largely depends on the continuous availability of an electrical power supply. This dependence is expected to be further amplified in the near future due to the increasing market share of electric vehicles and increasing integration of major national infrastructures to the power grid; e.g., water, transport, communications, etc. All of these point to a future where the reliability of the flow networks will be paramount, with the central research question being how we can design these networks in a robust and reliable manner. In the rest of the paper, we focus our attention on power systems for concreteness, but most of the discussion applies generally to any transport system.

This research was supported by the National Science Foundation through grant CCF #1422165 and a generous gift from *Persistent Systems, Inc.*

Preliminary version of some of the material was presented at the Information Theory and Applications Workshop (ITA 2016), San Diego, CA [4].

T. C. Gulcu is with TUBITAK Advanced Technologies Research Institute, 06800, Umitkoy, Ankara, Turkey. E-mail: tcgulcu@gmail.com

V. Chatziafratis is with the Computer Science Department, Stanford University, Palo Alto, CA 94305. Email: vaggos@stanford.edu

Y. Zhang and O. Yağan are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213. E-mail: {yingruiz, oyagan}@andrew.cmu.edu

A major problem with existing power systems is the seemingly unexpected large scale failures. Although rare, the sheer size of such failures has proven to be very costly, at times affecting hundreds of millions of people [2], [21], [28]; e.g., the recent blackout in India [34]. Such events are often attributed to a small initial shock getting escalated due to intricate dependencies within a power system [3], [27], [11]. This phenomenon, also known as cascade of failures, has the potential of collapsing an entire power system as well as other infrastructures that depend on the power grid [8], [16], [29]. Thus, understanding the dynamics of failures in power systems and mitigating the potential risks are critical for the successful development and evolution of many critical infrastructures.

In this paper, we continue our study of the robustness of power systems under a simple model based on *equal* redistribution of flow upon the failure of a power line. Namely, we consider a power system with N transmission lines with initial loads L_1, \dots, L_N and capacities C_1, \dots, C_N . If a line fails (for any reason), its load is assumed to be redistributed equally among all lines that are *alive*. Thus, the load carried by a line i may exceed its initial value L_i over time due to load redistribution. The capacity C_i defines the maximum flow allowed on the line i , meaning that if the load carried by i exceeds this capacity at any time, the line will be tripped (i.e., disconnected) by means of automatic protective equipments so as to avoid costly damages to the system. Subsequently, the load that was carried by line i before failure will be redistributed to remaining lines, which in turn may cause further failures, possibly leading to a *cascade* of failures.

The equal load redistribution model gets its appeal from its ability to capture the *long-range* [7], [19] nature of the Kirchhoff's law, at least in the mean-field sense, as opposed to the *topological* models [5], [14], [26] where failed load is redistributed only *locally* among neighboring lines. For example, it was suggested by Pahwa et al. [19] that equal load redistribution is a reasonable assumption¹ especially under the DC power flow model; the DC model is known [17], [24] to approximate the AC model well in many cases. With these in mind, an important goal is to understand the robustness of systems under the equal load redistribution model described above against *random* and *targeted* attacks. The former case was studied by Yağan [29] under the assumptions that initial loads L_1, \dots, L_N are independent and identically distributed

¹The equal load redistribution model is also similar in spirit to the well-studied CASCADE model introduced in [8], [9]. There, upon failure of a line, a fixed amount Δ of load is redistributed to all functional lines irrespective of the load being carried by the failed line or the number of remaining lines.

with $P_L(x) = \mathbb{P}[L \leq x]$ and that capacities are given by

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

where $\alpha > 0$ denotes the *tolerance factor*; in [29] all lines assumed to have the same tolerance factor. There, the robustness of the system against *random* attacks that target a p -fraction of the lines was studied; system robustness was quantified by the *final* (i.e., steady-state) fraction $n_\infty(p)$ of *non-failed* lines. Among other results, it was shown that the system robustness is maximized if all lines are given the same initial load, for a given fixed mean load $\mathbb{E}[L]$.

Recently, Zhang and Yağan [36] extended the results in [29] to the more general case where lines can have varying tolerance parameters. Namely, they let

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

with S_i denoting the free-space (or, redundancy) available at line i . The tolerance factor, given by $\alpha_i = S_i/L_i$, may now vary from one line to another. Under the assumption that load-‘free space’ pairs (L_i, S_i) are independent and identically distributed with $P_{LS}(x, y) = \mathbb{P}[L \leq x, S \leq y]$, they studied the robustness against random attacks that target a p -fraction of the lines. Their main result is that, with the mean values $\mathbb{E}[L]$ and $\mathbb{E}[S]$ are fixed, $n_\infty(p)$ is uniformly maximized for all p values if all nodes are given the same free space $\mathbb{E}[S]$, regardless of how the loads are distributed. This leads to the counterintuitive conclusion that lines with higher initial load shall be assigned smaller tolerance factors to maximize robustness and raises the possibility that the widely used (both in academy and industry) setting of equal tolerance factor is not optimal for system robustness (as far as the metric $n_\infty(p)$ is concerned).

With the case of random attacks being well-understood, we shift our attention in this paper to understanding the vulnerability of power systems under *targeted* attacks. As before, the main goal would be to derive design strategies (in the form of optimal load-‘free space’ distributions) that would lead to maximum robustness, this time against a knowledgeable adversary attacking a carefully selected set of lines. However, for this optimization problem to be well-defined one has to have a good understanding of the problem from an adversary’s perspective. With this in mind, this paper aims to develop *good* attack strategies that lead to maximal damage to the system for a given number of lines that can be attacked. The solution to the optimal attack problem will also help a system designer by i) revealing the worst-case attack vulnerability of the system which can help evaluate a given system design; and ii) revealing the most vulnerable lines in the system that will potentially be targeted by adversaries; this may then provide useful design guidelines for *improving* system robustness.

Formally, we consider the following optimization problem. Given N lines with loads L_1, \dots, L_N and free spaces S_1, \dots, S_N , we seek to find the optimal set A of k lines that the adversary should attack in order to minimize the final number $n_\infty(A)$ of alive lines. We provide optimal solutions via greedy algorithms in three special cases: i) when all lines have the same load; ii) when $S_i = \alpha L_i$ for each $i = 1, \dots, N$ (as commonly used in the literature [5], [11], [13], [14], [26], [29]); and iii) when all lines have the same free space, i.e.,

when $S_1 = \dots = S_N$. The last case is of particular interest as it is known to lead to maximum robustness against random attacks. For general cases of load-‘free-space’ distribution P_{LS} , we develop heuristic attack algorithms and show via simulations that these outperform benchmark attack strategies. In particular, we show that attacking the lines with the *largest* Load \times Free-space ($L * S$) is in general a good strategy. Parametrized attack strategies targeting the lines with the largest $L * S^\beta$ are also considered where β is in $[0, \infty]$, and optimum choices for β are provided in several cases.

We also consider a variation of the problem with an additional constraint on the total load of the lines attacked; i.e., when the adversary is further constrained with $\sum_{i \in A} L_i \leq Q$ for some Q . From a practical point of view, this might be the case if high-load carrying lines are *protected* better by the network owner and the *cost* of attacking them is proportional to their load. We show that this variation of the optimal attack problem is in fact NP-Hard, meaning that no polynomial-time algorithm can find the set A that minimizes $n_\infty(A)$, unless $P \equiv NP$. Our proof is based on a polynomial time reduction from the k -Subset Sum problem, i.e., the problem that seeks to find whether a sequence of integers has a subset of size k whose sum equals Q . For the modified optimization problem, we develop several heuristic algorithms and evaluate their performance in comparison with benchmarks through an extensive simulation study. In particular, we modify the previously developed heuristics with a *switch* that, when actuated during a sequential selection of lines to be attacked, changes the way algorithm makes the remaining selections; this idea is inspired from heuristics developed in [12] for the multi-dimensional 0-1 Knapsack problem. Among other things, we demonstrate via simulations that the max- $L * S$ algorithm with a *switch* performs well in a range of settings.

We believe that our results will shed some light on the vulnerability of power systems against powerful adversaries launching *targeted* attacks. In particular, we expect them to capture the qualitative properties of a power system well. This work may also have applications in other fields. In particular, the model considered here might be relevant for any flow network that is responsible for carrying (or, transporting) a physical quantity. A particularly interesting application is the study of traffic jams in roads [20], where the capacity of a line is regarded as the traffic flow capacity of a road.

The rest of the paper is organized as follows. We describe the system model in Section II. In Section III, we survey the results in [36] concerning robustness against random attacks. We start our discussion on optimal attack strategies in Section IV by demonstrating that certain greedy algorithms fail to give the optimal solution in general. In Section V we consider some special cases of interest where greedy algorithms are shown to find optimal attack sets. In Section VI we prove a hardness result showing that a variation of the optimal attack problem is NP-Hard. Finally, in Section VII, we develop heuristic algorithms for both the original and the modified optimization problems and present a detailed numerical study comparing the performance of our heuristics with benchmark algorithms such as max- S , max- L , max- C , and random attacks.

The random variables (rvs) under consideration are defined

on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where \mathbb{P} denotes the probability measure. We denote the expectation operator by \mathbb{E} . The indicator function of an event E is denoted by $\mathbf{1}[E]$. We write $|A|$ for the cardinality of a discrete set A .

II. SYSTEM MODEL

We consider a power system with N transmission lines $\mathcal{L}_1, \dots, \mathcal{L}_N$ with initial loads (i.e., power flows) L_1, \dots, L_N . The *capacity* C_i of a line \mathcal{L}_i defines the maximum power flow that it can sustain, and is given by

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

where S_i denotes the *free-space* (or, redundancy) available to line \mathcal{L}_i . The capacity of a line is typically written [14], [26], [13], [5] as a factor of the line's original load, i.e.,

$$C_i = (1 + \alpha_i)L_i$$

with $\alpha_i > 0$ defining the *tolerance* parameter for line \mathcal{L}_i . Put differently, the free space S_i is often given in terms of the initial load L_i with $S_i = \alpha_i L_i$. It is assumed that a line *fails* (i.e., outages) if its load exceeds its capacity at any given time. In that case, the load it was carrying before the failure is redistributed *equally* among all remaining lines.

The main question of interest is to characterize the robustness of this power system against i) *random* attacks that result initially with a failure of a (randomly selected) p -fraction of the lines; and ii) *targeted* attacks that initially fail a specific set A of lines. The initial set of failures lead to redistribution of power flows from the failed lines to *alive* ones (i.e., non-failed lines), so that the load on each alive line becomes equal to its initial load plus its equal share of the total load of the failed lines. This may lead to the failure of some additional lines due to the updated flow exceeding their capacity. This process may continue recursively, generating a *cascade of failures*, with each failure further increasing the load on the alive lines, and may eventually result with the collapse of the entire system.

Throughout, let $n_\infty(p)$ denote the *final* fraction of alive lines when a p -fraction of lines is randomly attacked. The robustness of the power system can be evaluated [29], [36] by the behavior of $n_\infty(p)$ as the attack size p increases, and particularly by the critical attack size p^* at which $n_\infty(p)$ drops to zero. In the case where a specific set A of lines are attacked, we define $n_\infty(A)$ as the final *number* of alive lines. The main goal of this paper is to seek *optimal* attack strategies, i.e., to find the set A of lines that minimizes $n_\infty(A)$ under certain constraints; e.g., the size $|A|$ being fixed.

The equal flow redistribution model is inspired by the democratic fiber bundle model [1], [6], where N parallel fibers with random failure thresholds C_1, \dots, C_N (i.e., capacities) drawn independently from $P_C(x)$ share equally an applied total force of F ; see also [20], [22]. This model has been recently adopted by Pahwa et al. [19] in the context of power systems with F corresponding to the total load that N power lines share equally. The formulation considered here, introduced in [29], [36], differs from the original democratic fiber-bundle model in that i) we do not assume that the total load of the system is fixed at F ; and ii) we allow for power lines to carry different initial loads; see also [18]. In addition,

[19] is concerned with failures in the power system that are triggered by increasing the total force (i.e., load) applied. Instead, our formulation allows analyzing the robustness of the system against external attacks or random line failures, which are known to be the source of system-wide blackouts in many interdependent systems [3], [32], [35].

III. RELEVANT WORK: EVALUATING AND OPTIMIZING ROBUSTNESS AGAINST RANDOM ATTACKS

We now survey the results obtained by Zhang and Yağan [36], [29] on the robustness of power systems under equal load redistribution of loads. These works consider the problem from a defender's perspective and provide means to characterize and optimize the robustness of the system, assuming that the adversary will launch a *random* attack to a certain fraction of lines. With the randomness involved in the attack model, as well as load-capacity values, [29], [36] rely on *mean-field* analysis to characterize the *mean* performance of the system in the asymptotic regime where N approaches infinity.

Throughout this section, assume that the pairs (L_i, S_i) are independently and identically distributed with $P_{LS}(x, y) := \mathbb{P}[L \leq x, S \leq y]$ for each $i = 1, \dots, N$. The corresponding probability density is given by $p_{LS}(x, y) = \frac{\partial^2}{\partial x \partial y} P_{LS}(x, y)$. We assume that $L \geq L_{min} > 0$ and $S > S_{min} > 0$ with probability one, and that the marginal densities $p_L(x)$ and $p_S(y)$ are continuous on their support.

A. Final system size as a function of the attack size

The main result in [36] characterizes the robustness of power systems under any distribution of initial load L and free space S , and any attack size p .

Theorem 3.1 ([36]): Let L and S denote generic random variables following the same distribution with initial loads L_1, \dots, L_N , and free space S_1, \dots, S_N , respectively. Then, with x^* denoting the smallest solution of

$$\mathbb{P}[S > x] (x + \mathbb{E}[L | S > x]) \geq \frac{\mathbb{E}[L]}{1-p} \quad (2)$$

over $x \in (0, \infty]$, the final system size $n_\infty(p)$ is given by

$$n_\infty(p) = (1-p)\mathbb{P}[S > x^*]. \quad (3)$$

An intuitive explanation of this result is that x^* , defined as the smallest solution of (2), gives the *extra load* each alive line has to carry (in addition to their initial load) at the steady-state; (3) is then understood easily since a line will still be alive at the steady-state if i) it is *not* targeted initially; and ii) it has enough free-space to handle the extra x^* amount of load.

For a graphical solution of $n_\infty(p)$, one shall plot $\mathbb{P}[S > x] (x + \mathbb{E}[L | S > x])$ as a function of x , and draw a horizontal line at $\mathbb{E}[L]/(1-p)$ on the same plot. The leftmost intersection of these two lines gives the operating point x^* , from which we can compute $n_\infty(p) = (1-p)\mathbb{P}[S > x^*]$. When there is no intersection, we set $x^* = \infty$ and understand that $n_\infty(p) = 0$; e.g., see Figure 1.

Theorem 3.1 enables computing $n_\infty(p)$ as a function of p , thereby helping understand the response of the system to attacks of varying magnitude. It also enables computing the

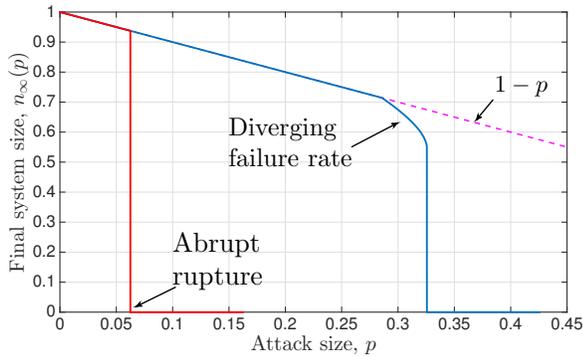


Fig. 1. We see the variation of final system size $n_\infty(p)$ as a function of the (random) attack size p when loads are drawn independently and uniformly from the range $[L_{\min}, L_{\max}] = [10, 50]$, and free spaces are given by $S_i = \alpha L_i$. Red (lower) curve stands for $\alpha = 0.2$, whereas blue (upper) curve represents $\alpha = 1.2$. For $\alpha = 0.2$, we observe an abrupt first-order transition of $n_\infty(p)$ as it suddenly jumps to zero at $p^* = 0.0625$, while decaying linearly as $1 - p$ up until that point. A slightly different behavior is seen when $\alpha = 1.2$ in that $n_\infty(p)$ exhibits a continuous divergence from the $1 - p$ line before complete system failure, which again takes place through a discontinuous transition.

critical attack size p^* , defined as the largest attack the system can sustain. More precisely, we define

$$p^* = \sup_{p > 0} n_\infty(p) > 0,$$

so that for any attack with size $p > p^*$, the system undergoes a complete breakdown leading to $n_\infty(p) = 0$. The next result provides a closed form expression for the critical attack size.

Theorem 3.2 ([36]): The critical attack size p^* is given by

$$p^* = 1 - \frac{\mathbb{E}[L]}{\max_x \{\mathbb{P}[S > x] (x + \mathbb{E}[L | S > x])\}}. \quad (4)$$

The operational usefulness of p^* derives in part from its ability to quantify the system robustness by a single scalar, paving the way to designing the system with optimal robustness by means of maximizing p^* ; see Section III-B below.

An interesting question here is whether $n_\infty(p)$ decays to zero continuously (i.e., through a second-order transition), or discontinuously (i.e., through a first-order transition). The practical significance of this is that continuous transitions suggest a more stable and predictable system behavior with respect to attacks, whereas with discontinuous transitions, system behavior becomes more difficult to predict based on past data. In [36], it was shown that the total breakdown of the system will always be through a *first-order* (i.e., discontinuous) transition under random attacks. More precisely, it was shown that $n_\infty(p^*) > 0$, while by definition it holds that $n_\infty(p^* + \varepsilon) = 0$ for any $\varepsilon > 0$ arbitrarily small. This shows that despite its simplicity, the equal flow redistribution model can capture real-world phenomena of unexpected large-scale system collapses; i.e., cases where seemingly identical attacks/failures leading to entirely different consequences. In addition, depending on system parameters, the first-order breakdown may have early indicators at smaller attack sizes such as a *diverging* failure rate leading to a non-linear decrease in $n_\infty(p)$. In other cases, the system *abruptly* transitions to a total collapse while perfectly resisting attacks until the critical point; see Figure 1 and [36, Fig. 2] for the rich transition behavior that can be observed under the equal load redistribution model.

B. Achieving Optimal Robustness

An important question from a system designer's perspective is concerned with deriving the *universally optimum* distribution of initial loads L_1, \dots, L_N and free spaces S_1, \dots, S_N when the mean values $\mathbb{E}[L]$ and $\mathbb{E}[S]$, respectively, are fixed. The results obtained in [36] concerning this are presented next.

Theorem 3.3 ([36]): For any initial load-free space distribution, it holds that

$$p^* \leq \frac{\mathbb{E}[S]}{\mathbb{E}[S] + \mathbb{E}[L]} = \frac{\mathbb{E}[S]}{\mathbb{E}[C]} = p_{\text{dirac}}^* \quad (5)$$

where p_{dirac}^* denotes the critical attack size when the free-space follows a Dirac-delta distribution, i.e., $P_{LS}(x, y) = P_L(x)\mathbf{1}[y \leq \mathbb{E}[S]]$, for any $P_L(x)$.

In words, this result states that the critical attack size can never be larger than the ratio of mean free space to mean capacity. In addition, this *optimal* p^* value can be achieved by assigning every line the same free-space, i.e., by setting $S_1 = \dots = S_N = \mathbb{E}[S]$, regardless of how loads are distributed.

Thus far, we have seen that the equal distribution of free space leads to the largest possible critical attack size, hereafter denoted p_{opt}^* . It is clear that the final system size after an attack of size p is at most $1 - p$. With this in mind, the next result establishes the optimality of the Dirac-delta distribution of free-space in the sense of maximizing the robustness of power systems *uniformly* over all attack sizes.

Theorem 3.4 ([36]): With $S_1 = S_2 = \dots = S_N$ we have

$$n_\infty(p) = \begin{cases} 1 - p & \text{for } p < p_{\text{opt}}^* \\ 0 & \text{for } p \geq p_{\text{opt}}^* \end{cases}$$

Hence, the Dirac-delta distribution of free-space S maximizes $n_\infty(p)$ over the entire range $0 \leq p \leq 1$.

These findings suggest that under the equal flow-redistribution model considered here, power systems with homogeneous distribution of redundant space are more robust against random attacks and failures, as compared to systems with heterogeneous distribution of redundancy. Interestingly, this suggests that the optimal robustness is achieved when the tolerance factor $\alpha_i = S_i/L_i$ decreases with increasing load, leading to the counter-intuitive conclusion that the lines carrying the highest load should have the smallest tolerance factor to achieve maximum robustness; see [36] for a more detailed discussion on this matter.

IV. THE OPTIMAL ATTACK PROBLEM AND INSIGHTS

While existing results in [36], [29] shed some light on the robustness under the equal flow redistribution model, many real-world threats to power systems would be expected to come from powerful adversaries targeting specific parts of the grid to inflict maximum damage on the overall system. In order to understand the system's robustness against such *targeted attacks* and to reveal its most vulnerable lines, we now consider the problem from an adversary's perspective and seek effective strategies for attacking the system under given constraints. In particular, we consider a scenario where the adversary has full information about the system and aims to find the best set of k of lines to attack so that it fails maximum number of lines as a result cascading failures. This optimization problem is formally introduced next.

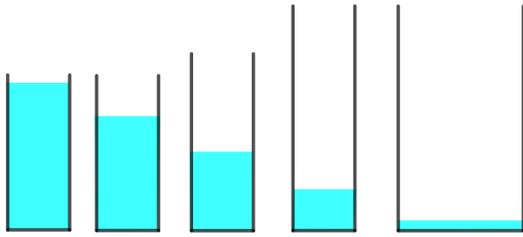


Fig. 2. In this example we have (load, capacity) values given by $(8, 8 + \epsilon)$, $(6, 8 + \epsilon)$, $(4, 4 + 14/3 + \epsilon)$, $(2, 11 + \epsilon)$, $(1, 21 + \epsilon)$ where $\epsilon > 0$ is arbitrarily small. The greedy maximum-load attack will need to attack $k = 5$ containers to fail all. It will start attacking the leftmost container with load $L_1 = 8$ which will not lead to any further failures. Then, it will continue with the second one from the left, again unable to trigger a cascade, and continue until attacking all containers directly. The optimal solution can be seen to be $k = 1$ by attacking the last container, which will trigger a cascading failure destroying the whole system. We can generalize this counterexample to the case with N containers with the greedy algorithm's output being $k = N$ while the optimal solution being $k = 1$.

A. The Main Optimization Problem: ER- k

The Equal Redistribution (ER) problem with k attacks is the *optimization problem*, denoted ER- k , that aims to find the set A of k lines such that attacking A leads to the maximum number of total line failures (as a result of load redistribution and cascading failures), among all possible attack sets with size k . Put differently, we seek to find A with $|A| = k$ that minimizes $n_\infty(A)$. Throughout, we find it useful to consider the *decision* version of this optimization problem (referred to as the ER- k - k' problem) formally defined as follows.

INPUT: N pairs of non-negative numbers in the form (L_i, C_i) indicating the load and the capacity of each line, and integers k and k' such that $0 < k < k' \leq N$. We assume $C_i > L_i$ so that no line fails initially at its own load.

OUTPUT: The answer to whether or not there is an attack set A with size k , such that at the end of the cascading failures the number of failed nodes is at least k' ; i.e., whether there exists A with $|A| = k$ and $n_\infty(A) \leq N - k'$.

B. Heuristic Algorithms that Fail

Here we will present three intuitive greedy algorithms and give concrete examples demonstrating their poor performance for the optimization problem described above. In doing so, we will focus on the special case where $k' = N$ meaning that the goal of the attack is to destroy the whole system, by attacking a minimum number k of lines.

In what follows, we find it useful to describe the problem in a simpler way, where we have N water containers with capacities C_1, \dots, C_N , and initial water levels L_1, \dots, L_N . As in the equal flow-redistribution model, when a container is "attacked" its content is redistributed equally to the remaining containers. Also, if the water level in a container exceeds its capacity, we assume that it has failed and redistribute its content, again equally, to the remaining containers. With this formulation, the goal of the attackers is to find the smallest number k of containers to target so that all containers get overloaded and fail eventually.

An important observation is that the following intuitive algorithms can deviate significantly from the optimal solution.

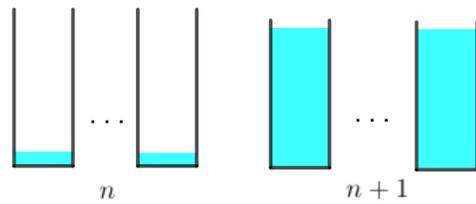


Fig. 3. Consider $2n + 1$ containers where (load, capacity) values are given by (ϵ, M) for the first n containers and $(M - 2\epsilon, M - \epsilon)$ for the last $n + 1$ containers; here $\epsilon > 0$ is arbitrarily small and $M > 2(n + 1)\epsilon$. The greedy max-capacity attack will need to attack $k = n + 1$ containers to fail the all containers; it will start attacking the first n containers but cascading failures will not take place. On the other hand, the optimal solution is $k = 1$ as it takes to attack only one of the containers with $(M - 2\epsilon, M - \epsilon)$ to trigger a cascading failure that will fail all.

a) *Greedy max-load attack:* This greedy strategy aims to maximize the load that will be redistributed in each attack round. Namely, it starts by attacking the container with the highest load, and proceeds similarly, waiting after every attack for a steady-state to be reached (meaning that all load redistribution and potential further failures end). The algorithm stops when all containers fail. This strategy is not optimal in general because it fails to recognize the opportunity to eliminate containers with very large capacities that will otherwise be difficult to fail by redistributing the load. The worst-case deviation from the optimal (in terms of the number of lines needed to be attacked for complete system failure) is $\Theta(n)$; e.g., see Figure 2.

b) *Greedy max-capacity attack:* This strategy is similar in spirit with the greedy max-load attack except that this time the container with the maximum capacity is attacked in each round. The idea here is that by taking out large containers, the remaining, supposedly small, containers will be destroyed due to load redistribution. This strategy is not in general optimal either, because there may be containers with *large* capacities but small (or, even almost zero) loads, rendering an attack to such containers very ineffective in terms of triggering failures by means of load redistribution. The worst-case deviation from the optimal is again $\Theta(N)$ as demonstrated in Figure 3.

c) *Greedy max-free-space attack:* It is clear from the previous two cases that the optimal attack strategy will be one that considers both the loads and capacities of the containers involved. The greedy approach that targets containers with largest free space (i.e., (capacity - load) difference) falls into this category, and is based on the fact that containers with largest free space will fail the *latest* in the course of a cascading failure; e.g., see Section IV-C for a discussion of this fact. Therefore, it is sensible to eliminate those containers with a direct attack. On the other hand, containers with small free space are already on the verge of failing and therefore can be taken down by means of redistribution of loads. Although this greedy strategy is intuitive (and in fact optimal in some special cases), it fails to be the optimal solution in general. The main reason is that this approach does not take into account the loads of the containers directly. For example, a container may have a large free space but its load may be negligible compared to other containers, rendering a direct attack on this container ineffective. The worst-case deviation from the optimal is again $\Theta(N)$ as demonstrated in Figure 4.

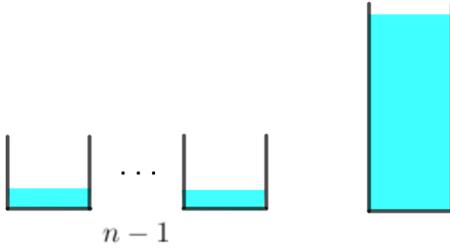


Fig. 4. In this example we have n containers with (load, capacity) values $(\epsilon, (n+1)\epsilon)$ for the first $n-1$ containers and $(M, M+(n-1)\epsilon)$ for the last container, where $\epsilon > 0$ is arbitrarily small and M satisfies $M > (n^2 - n)\epsilon$. The greedy max free-space ($C - L$) attack will output $k = n$ since it will start attacking the leftmost containers and no cascading failures will take place. The optimal solution is obviously $k = 1$ by attacking the last container.

C. Observations towards Designing a Smart Algorithm

We now present some observations that will be useful in designing a *smart* attack algorithm.

a) *The order of attack does not matter:* In the equal redistribution model, the order with which we launch an attack does not affect the final set of failed containers. This is because the load of the attacked nodes will be distributed to all of the remaining nodes so at the end an amount of $\sum_{i \in A} L_i$ will end up in the remaining containers (leading to new failures or not) irrespective of the order we chose to attack the containers in A . We remark that an attack strategy can still be designed in a greedy fashion, where the set A is determined one member at a time, waiting for cascades to stop after each attack.

b) *Order of failures during the cascading process:* Assume that containers are sorted by increasing free space, $S_i = (C_i - L_i)$. Given that any failed load is redistributed *equally* among the remaining containers, it is clear that this ordering will remain the same throughout the course of cascading failures; the containers that are attacked directly at the beginning are excluded from this argument. Therefore, in the process of recursive load redistribution, containers will fail (due to their free space diminishing to (below) zero) in this exact same order: the one with smallest free space will fail first, and so on and so forth.

V. OPTIMAL ATTACK STRATEGIES UNDER SPECIAL CASES

We now present three special cases of the ER- k problem and provide *optimal* attack strategies for each of them.

a) *Same Loads:* An interesting situation arises when initial loads are the same for all containers while capacities differ. This reflects situations in which all lines in the power system are given the same initial load, but have different capacities owing to the physical constraints or material used. We show that a greedy algorithm finds the optimal solution in this special case. The ER- k -Same Loads Problem is defined formally as follows.

INPUT: A non-negative rational number L for the common load and a list of N non-negative numbers $C_i > L, \forall i$ indicating the capacity of each line. The integer k represents the number of attacks we can launch.

OUTPUT: The set A of lines to be attacked that minimizes $n_\infty(A)$ under the constraint $|A| = k$.

The next result, proved in Supplementary Material, shows that the max- C -greedy algorithm finds the optimal solution.

Theorem 5.5: The max- C -Greedy Algorithm is optimal for the ER- k -Same Loads Problem.

b) *Same Free Spaces:* Sometimes it might be the case that the containers have arbitrary load and capacity but they have a fixed free space. In [36], this was in fact shown to be the optimal design that gives maximum robustness against *random* attacks. We refer to the corresponding problem as the ER- k -Same Free Spaces, formally defined as follows.

INPUT: A list of N non-negative rational numbers L_i indicating the load of each container and a positive rational number S indicating the common free space.

OUTPUT: Find the minimum number k of containers needed to be attacked in order to destroy the whole system.

We changed the output from having a fixed number of lines to be attacked to inflict the maximum damage, to the case where we aim to destroy the whole system with the minimum number of attacks. This is because in the case where every container has the same free space, there are no intermediate cascading failures. After an attack, the system will either fail completely, or no single line will fail other than those attacked directly. We show in Supplementary Material that the max- L -Greedy algorithm that targets lines with the largest loads leads to the optimal solution for this problem.

Theorem 5.6: The max- L -Greedy Algorithm is optimal for the ER- k -Same Free Spaces Problem.

c) *Capacities Proportional to Loads:* In many cases, the capacities and the loads of power lines are related in a particular way. Namely, the capacity of a line is often set to be proportional to its load. For example with $\alpha > 0$ denoting the *tolerance factor*, we have $C_i = (1 + \alpha)L_i$ for each line $i = 1, \dots, N$. In this variation, we will also show that there is a greedy algorithm achieving the optimal solution. The ER- k - $(C \propto L)$ Problem is defined formally as follows.

INPUT: A list of N non-negative numbers L_i indicating the load of each container and a positive number α such that container capacities are set to $C_i = (1 + \alpha)L_i$ for each i .

OUTPUT: The set A of lines to be attacked that minimizes $n_\infty(A)$ under the constraint $|A| = k$.

In this setting, load, free-space, and capacity of the lines will be ordered in the same way, and as we show in the Supplementary Material, max- L, C, S -Greedy algorithms that target lines with the largest load *and* free-space *and* capacity give the optimal solution to this problem.

Theorem 5.7: The max- L, C, S -Greedy Algorithms are optimal for the ER- k - $(C \propto L)$ Problem.

VI. A MODIFIED OPTIMAL ATTACK PROBLEM WITH TOTAL LOAD CONSTRAINTS

In this Section, we will prove that a variation of the decision problem ER- $k-k'$ is NP-Complete. In particular, we consider the ER- $k-k'-Q$ problem, defined formally as follows.

INPUT: N pairs of non-negative numbers in the form (L_i, C_i) indicating the load and the capacity of each line, integers k and k' such that $0 < k < k' \leq N$, and a positive number Q . We also assume $C_i > L_i$ for each $i = 1, \dots, N$.

OUTPUT: The answer to whether or not there is an attack set A with size k , and total sum of loads $\sum_{i \in A} L_i \leq Q$, such

that at the end of the cascading failures the number of failed nodes is at least k' ; i.e., whether there exists $A \subset \{1, \dots, N\}$ with $|A| = k$, $\sum_{i \in A} L_i \leq Q$, and $n_\infty(A) \leq N - k'$.

It is clear that the objective is two-fold here and that there is an inherent trade-off: by attacking lines with larger initial loads we can shed more load on other lines and have a better chance to trigger a cascade of failures that would destroy the whole system. However, the problem enforces a constraint on the total load of the attacked containers as well. This *knapsack-like* trade-off is what makes the problem NP-complete as we now show. Our proof is based on the reduction of the ER- k - k' - Q problem from the k -Subset Sum variant defined as follows: *Given a set of integers and a target sum Q , is there any subset of size k whose sum is Q ?*

Theorem 6.8: The ER- k - k' - Q Problem is NP-Complete.

Proof. First, we show that ER- k - k' - Q Problem is in NP: The certificate is a list of the k containers we choose to attack. We can check in polynomial time (e.g., see the ER-Attack Projection algorithm in [4]) whether at least k' lines in the system fail or not. Since we have a certificate that can be checked in polynomial time, ER- k - k' - Q is in NP!

Given an instance of the k -Subset Sum problem we will create an instance of the ER- k - k' - Q problem: Given a set of N integers a_1, a_2, \dots, a_N , the k -Subset Sum problem asks whether there exists k members of the set whose some equals Q . If $k = N$, we can check if $\sum_{i=1}^N a_i = Q$ and respond accordingly. From now on, we suppose $k < N$ and create an equivalent version of the ER- k - k' - Q problem in the following manner. Let lines $\mathcal{L}_1, \dots, \mathcal{L}_N$ have loads $L_1 = a_1, L_2 = a_2, \dots, L_N = a_N$ and consider the ER- k - k' - Q problem; i.e., we seek to find a set A of k lines such that $\sum_{i \in A} L_i \leq Q$ and that attacking A leads to failure of at least $k' > k$ lines in the system. We also set $C_i = L_i + S_i$ where the free space is $S_i = \frac{Q}{N-k}$ for each $i = 1, \dots, N$. This last constraint ensures two things. First, as discussed in Section V, when all lines have the same free space then attacking k lines can only have two consequences: either only those k lines that are attacked fail, or all N lines fail. In either case, there is no *cascade* of failures and the system reaches a steady-state immediately. Thus, with equal free space among all lines, the ER- k - k' - Q problem with $k' > k$ is equivalent to ER- k - N - Q problem. Secondly, under the enforced assumptions it is clear that a complete system failure will take place if and only if the total load failed by the initial attack A is larger than the sum of the free spaces of those that are not in the attack set A ; i.e., if and only if

$$\sum_{i \in A} L_i \geq \sum_{j \in \{1, \dots, N\} \setminus A} S_j = (N - k) \frac{Q}{N - k} = Q.$$

Here, the first equality follows from the facts that $|A| = k$ and $S_i = \frac{Q}{N-k}$ for each $i = 1, \dots, N$. Recalling further the constraint that $\sum_{i \in A} L_i \leq Q$, this leads to $\sum_{i \in A} L_i = Q$. Therefore, the created instance of the ER- k - k' - Q problem indeed seeks to find a subset A of $\{a_1, \dots, a_N\}$ such that $|A| = k$ and $\sum_{i \in A} L_i = Q$, rendering it equivalent to the k -Subset Sum instance that we have started with. For the reverse direction, assume that the ER- k - k' - Q problem has a

solution with k lines $\mathcal{L}^{(1)}, \dots, \mathcal{L}^{(k)}$. Then the loads of these lines constitute a solution to the k -Subset-Sum problem.

The above reduction can be constructed in polynomial time (more precisely, in linear time), so if there was a polynomial algorithm that could solve the ER- k - k' - Q , then the k -Subset Sum would be in P, which is wrong unless P=NP. Thus, we conclude that the ER- k - k' - Q Problem is NP-complete. ■

An important implication of the above result is that the *optimization* version of the ER- k - k' - Q problem, which seeks to find the set A of lines that minimizes $n_\infty(A)$ under the constraints $|A| = k$ and $\sum_{i \in A} L_i \leq Q$, is NP-Hard. This means that under these constraints, the adversary can not launch an optimal attack in polynomial time unless P=NP.

VII. HEURISTIC ALGORITHMS AND THEIR PERFORMANCE

Although, it is not known whether the original optimization problem of finding the best k lines to attack to minimize final system size is NP-Hard or not, the discussion in the preceding section indicates that the optimal attack problem is likely to be computationally challenging; in particular, we know that the problem is NP-Hard if we are further constrained by the total load of those we can attack. This prompts us to develop heuristic algorithms, for both the original and the modified optimization problems, that work in polynomial time and have competitive performance under arbitrary load-capacity distributions. The performance of these heuristics will then be compared with some benchmark heuristics such as max- L , max- C , max- S , and random attacks.

In the interest of brevity and concreteness, the discussion is restricted here to *non-greedy* algorithms that choose the attack set A *without* ever running the attack projection algorithm. In other words, the algorithms are not allowed to run the cascading failures initiated by a subset of A , and then continue making the remaining selections from the lines that survived the cascade. Of course, all heuristics considered here including the benchmarks can be modified to operate in a greedy fashion. One might also be tempted to use a greedy algorithm where, at each round, the line to be attacked is chosen in an optimal way; i.e., the line whose failure leads to smallest number of surviving lines is chosen from among all lines that are still alive. However, for the problem at hand, one can realize that unless $|A|$ is relatively large, the final system size equals $n_\infty(A) = N - |A|$ regardless of the set A of attacked lines. For example, with $|A| = k$, this will be the case whenever

$$k \leq \frac{S_{min}}{L_{max}} (N - k),$$

meaning that if $S_{min}/L_{max} > 0$, the greedy heuristic will have to deal with *ties* when making its choices for the next line to be attacked until it makes $\Omega(n)$ choices. Since resolving the ties by randomization for such a large number of selections is likely to lead to poor performance, one needs heuristic rules to resolve the ties. Even then, our preliminary simulation study indicated that the greedy versions of the heuristics considered here perform only slightly better than their non-greedy counterparts, and the comparison among the greedy heuristics provided no additional insight to what was already

observed from non-greedy algorithms; hence the decision to consider only non-greedy attack strategies here.

A. Heuristics for the original optimization problem

We first consider the original case where there is no constraint on the total load of the lines that can be attacked; i.e., we consider the ER- k problem. Let A be the set of lines to be attacked such that $|A| = k$. It is clear from the previous discussions that a good attack should aim to

$$\text{maximize } \sum_{i \in A} L_i, \text{ and} \quad (6)$$

$$\text{minimize } S_j, j \in \{1, \dots, N\} - A \quad (7)$$

In words, the attack should aim to find the lines with the largest free space while making the total load of the failed lines as large as possible. Thus, the attack should intuitively look for lines with *large* initial load *and* *large* free-space. Of course, most difficult situations arise when the load and free-space values of the lines are in reverse order; e.g., the highest load carrying line has the smallest free-space, etc. as in Figure 2.

Our main idea towards handling the trade-off described above is based on its similarities with the well-studied 0-1 Knapsack problem. In the 0-1 Knapsack problem, we are given a set of N items, $\{1, \dots, N\}$, each with a weight w_i and a value v_i , and the goal is to choose items such that their total value is maximized while the total weight is bounded by W ; i.e.,

$$\text{maximize } \sum_i v_i x_i \quad (8)$$

$$\text{subject to } \sum_i w_i x_i \leq W \quad \text{and} \quad x_i \in \{0, 1\} \quad (9)$$

The 0-1 Knapsack problem is known to be NP-Hard, but polynomial-time heuristics can still give close-to-optimal solutions. For example, a competitive heuristic is to order the items based on their "value per weight", i.e., v_i/w_i , and choose items according to this order, starting with the one with the highest v_i/w_i , until the total weight capacity W is reached. In fact, with a small modification to handle corner cases, this heuristic is known to yield at least 50% of the optimum value.

The optimal attack problem we consider, i.e., the ER- k problem, has some similarities with but is not equivalent to the 0-1 Knapsack problem. In particular, one can construct an analogy between the constraints of the 0-1 Knapsack problem and the ER- k problem by assigning all item weights as $w_i = 1$ and the total weight limit to be $W = k$. However, the objectives (6)-(7) of the ER- k problem are much more complex than the objective (8) of the Knapsack problem. Nevertheless, the two problems have some similarity in that their main difficulties lie in the trade-offs involved. In the Knapsack problem the trade-off is between the value and the weight of the item and it is desirable to pick items with high value and low weight, while in the ER- k , the trade-off lies between the possibly conflicting objectives of choosing lines with high load and high free-space. Inspired with the efficient heuristic for the Knapsack problem that is based on selecting items with the largest v_i/w_i ratio (i.e., items with the biggest *bang for the buck*), our first heuristic for the ER- k problem is based choosing lines with the highest load times free-space, i.e., with the highest $L_i * S_i$.

a) *Maximum Load×Free Space Attack*: In this algorithm, the load free space product, $L_i * S_i$ is computed for each line $i = 1, \dots, N$. After sorting the lines based on this product, the k lines to be attacked is chosen as the ones having the highest k load-free space product. As mentioned above, this is inspired by the 2-approximation heuristic for the Knapsack problem that orders items according to $v_i * \frac{1}{w_i}$ when the goal is to choose items with high v and low w . In the ER- k problem, we wish to choose lines with high L and high S , or equivalently, with high L and low $\frac{1}{S}$. Thus, constructing an analogy between value v_i and load L_i , and weight (or, cost) w_i and $1/S_i$, our heuristic chooses lines with the maximum

$$L_i * \frac{1}{1/S_i} = L_i * S_i.$$

The performance of this heuristic is demonstrated via several numerical examples in the next subsection along with a comparison with some benchmark heuristics.

Aside from its connection to a powerful heuristic in a relevant problem, the maximum $L * S$ heuristic has several advantages. First of all, this heuristic becomes equivalent to the *optimal* attack strategy in the three special cases considered in Section V; e.g., when all lines have the same load, it chooses ones with highest free-space (and hence capacity), or when all lines have the same free-space, it chooses lines with maximum L , etc. Secondly, considering the product $L * S$ is an effective way to favor lines with high load *and* free-space, while heavily penalizing load or free-space values close to zero; note that benchmark heuristics including the highest-capacity attack ($C = L + S$) fail to penalize small L or C values. In the optimal attack problem, this makes perfect sense given that a line with almost no load should never be attacked even if it has very high free-space since its failure will likely not affect any other line. Similarly, it may not be a good idea to directly attack a line with almost no free-space even if it has very high load, since the line will likely fail due to load redistribution *regardless* of which other lines are attacked.

b) *Maximum $L * S^\beta$ attack*: While maximum $L * S$ attack is intuitive and will be seen to be powerful in many cases, we observe that its performance can be further improved by a small modification. To this end, we propose a second heuristic as a modified version of the max- $L * S$ attack that allows adjusting the relative importance of load and free-space values of the lines. In particular, with β in $[0, \infty]$, we consider a heuristic that chooses k lines with the maximum $L_i * S_i^\beta$. An added benefit of this heuristic is that it contains several heuristics as special cases. In particular, the maximum $L * S$ algorithm described above is obviously a special case of this algorithm, corresponding to the case $\beta = 1$. Also, by setting $\beta = 0$, this heuristic reduces to the max- L attack, while setting $\beta = \infty$ (or, large enough) makes it equal to the max- S attack.

B. Numerical Comparison with Benchmark Heuristics

We now compare the heuristics we developed against some benchmark heuristic algorithms via numerical experiments. The benchmark heuristics we will consider are given below:

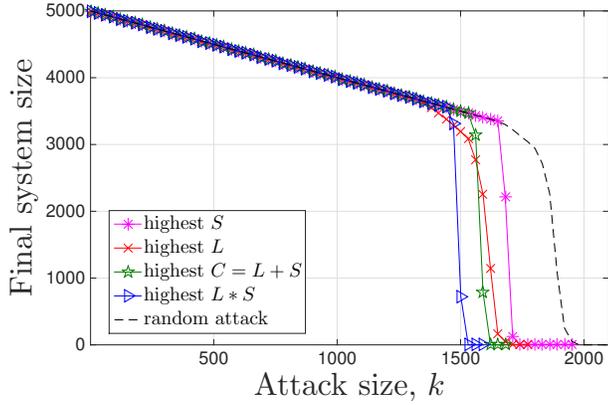


Fig. 5. The performance comparison of different heuristic algorithms for $L \sim U[10, 30]$, $S \sim U[10, 60]$, $N = 5000$.

a) *Random attack*: This is the most primitive attack strategy and considered here only for comparison purposes. The attack picks k lines to be attacked uniformly at random from amongst all N lines.

b) *Highest-L, highest-C, highest-S attacks*: These three attacks are based on sorting lines with respect to their initial load L_i , free-space S_i , or capacity $C_i = L_i + S_i$, respectively and attacking the top k lines with the highest value of the corresponding metric.

We fix the number of lines at $N = 5000$ for all experiments. First, we consider the case where each line is independently given an initial load from a uniform distribution, $U(L_{min}, L_{max})$, where we set $L_{min} = 10$ and $L_{max} = 30$. The free-space allocated to each line is generated independently from its load, again from a uniform distribution, $U(S_{min}, S_{max})$ with $S_{min} = 10$, $S_{max} = 60$. The capacity of a line L_i is given by the sum $C_i = L_i + S_i$. The independence of L and S leads to some lines having high load but small free-space, or vice versa, rendering the optimal attack problem *non-trivial*; e.g., with these choices, the realized load-capacity values will almost surely *not* fall into one of the special cases presented in Section V where an optimal solution is available.

Under this setup, we compute the final system size as a function of the number k of lines attacked, where the set of attacked lines are selected according to the heuristics considered. The results are given in Figure 5 where each data point is obtained by averaging over 100 independent runs. We already see in this simple setting that our attack strategy of targeting lines with the highest $L * S$ outperforms all other benchmarks (except a small interval of attack size where max- L attack seems to give the highest damage). In particular, we see that the highest $L * S$ attack is able to fail the whole system by targeting 90, 180, 210, and 450 fewer lines as compared to max- C , max- L , max- S , and random attacks, respectively.

Next, we check if this performance can be further improved by attacking lines with highest $L * S^\beta$ for some $\beta \geq 0$. To this end, we repeat the previous experiment as β varies from zero to ten. The results are demonstrated in Figure 6 and as expected show that with $\beta = 0$ or $\beta \gg 1$, we obtain the same performance with max- L and max- S attack, respectively. More interestingly, we see that the case $\beta = 1$ is indeed *not* the best one can do. For example, we see that when $\beta = 0.3$, the max-

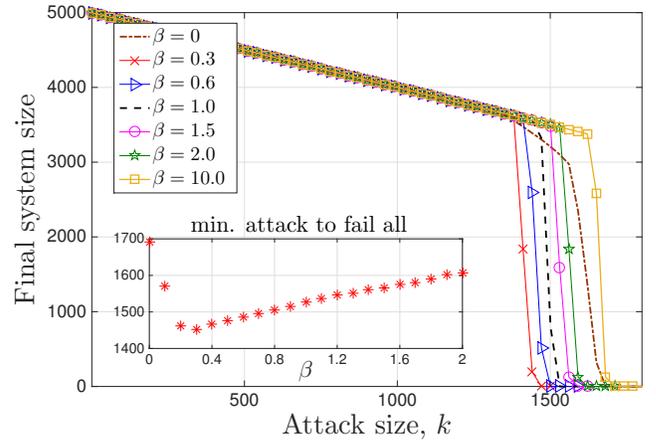


Fig. 6. The performance comparison of maximum $L * S^\beta$ algorithms for various β values for $L \sim U[10, 30]$, $S \sim U[10, 60]$, $N = 5000$. Inset: The minimum number of lines needed to be attacked to fail all lines for the maximum $L * S^\beta$ attack.

$L * S^\beta$ attack can fail the whole system by attacking 75 fewer lines the case for $\beta = 1$. To demonstrate this better, we plot in the inset of Figure 6 the minimum number lines needed to be attacked to fail all N lines. To compute this, we again run 100 independent experiments and pick the minimum attack size for which *all* 100 experiments led to entire system failure.

The performance of heuristic algorithms are known to vary significantly under different parameter settings, and our case is not expected to constitute an exception. To that end, we have tested the performance of the max- $L * S^\beta$ attack with $\beta \in [0, 4]$, for a wide range of possibilities for the distribution of L and S . In all cases we considered, we were able to identify a β value for which the max- $L * S^\beta$ attack was at least as good as all benchmark attacks (random, highest- C , L , S -attacks) showing its versatile performance.

As already mentioned, the most challenging cases arise when the load and free-space values are in reverse order. To that end, we close this section by demonstrating the performance of the max- $L * S^\beta$ attack in such cases. In particular, we start by generating L_1, \dots, L_N and S_1, \dots, S_N independently according to some distribution. Then, the load values (resp. free-space values) are sorted and re-arranged in increasing (resp. decreasing) order, leading to highest-load carrying line having the smallest free-space, and so on. To make the problem more challenging and interesting, we also consider *Pareto* distribution. Namely, a random variable X is said to follow Pareto distribution, written $X \sim \text{Pareto}(X_{min}, b)$ with $X_{min} > 0$ and $b > 0$, if its probability density is given by

$$p_X(x) = X_{min}^b b x^{-b-1} \mathbf{1}[x \geq X_{min}].$$

To ensure that $\mathbb{E}[X] = bX_{min}/(b-1)$ is finite, one must set $b > 1$, while the variance of X is finite only if $b > 2$.

The results for the case where L and S values are reverse ordered are depicted in Figure 7. Here, we show a small number of representative results that correspond to different behaviors for brevity. As before, all results correspond to the minimum attack size that led to an entire system collapse in all 100 experiments. The curves represent the results for the max- $L * S^\beta$ attack as β varies from zero to two. In each plot,

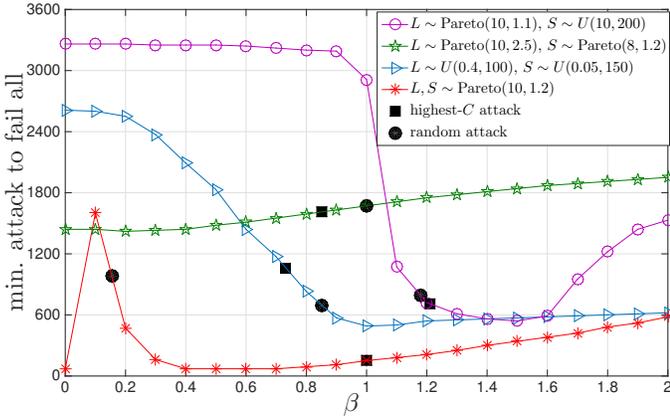


Fig. 7. Minimum number of lines needed to be attacked to fail all lines in the system is shown when load and free-space values are generated independently (from the distributions given at the figure legend) and then sorted in reverse orders; e.g., to ensure $L_1 \leq L_2 \leq \dots \leq L_N$, while $S_1 \geq S_2 \geq \dots \geq S_N$. Curves stand for the results obtained under the max- $L * S^\beta$ attack as a function of β . Corresponding results for the max- C attack are shown by filled square symbols, and those for the random attack by filled circles.

we add the corresponding results for the max- C attack (shown by a filled square) and random attack (shown by a filled circle) as well; for convenience, the x -axis values for these symbols are chosen such that they stay on the corresponding curve showing the results for max- $L * S^\beta$ attack. We note that the max- L attack is already demonstrated by the case $\beta = 0$ while $\beta = 2$ gives a good indication of the performance of the max- S attack, so these plots provide a comparison of the max- $L * S^\beta$ attack with all the benchmarks considered here.

The main observations from Figure 7 are as follows. We see that in all cases there is a particular β value for which the max- $L * S^\beta$ attack performs the best among all benchmarks; it is only the case of $L, S \sim \text{Pareto}(10, 1.2)$ where we see that the best performance of max- $L * S^\beta$ attack is attained when $\beta = 0$, or $\beta \in (0.4, 0.7)$ meaning that max- L attack matches the performance of the max- $L * S^\beta$ attack. Also, we see that the β value that leads to the best performance can be equal to, smaller than, or larger than one in different scenarios showing the importance of trying different values of β to get the best performance. Finally, while benchmark attacks (including the random attack) occasionally give results close to the best max- $L * S^\beta$ attack, we see examples for each benchmark where its performance is significantly worse than the best max- $L * S^\beta$ attack; these cases are summarized in Table I.

C. Heuristic attacks for the modified optimization problem

We now consider the modified optimization problem ER- $k-k'-Q$ where the attack set A is further constrained by $\sum_{i \in A} L_i \leq Q$, in addition to $|A| \leq k$. As shown in Theorem 6.8, in this case finding the attack set A that minimizes the final system size is NP-Hard, prompting us to develop heuristic strategies. With the additional constraint on the total load of the lines that we can attack, the trade-offs involved become more complicated and heuristics developed in the previous section may not be well-suited for the ER- $k-k'-Q$ problem. Ultimately, our strategy should be to choose an attack set that has k (or, very close to k) lines with total load equal (or, very close) to Q , and that have the highest free space among all

lines. This is because at the first stage of the cascades, any line \mathcal{L}_i that was not directly attacked will fail only if

$$S_i \leq \frac{\sum_{i \in A} L_i}{N - |A|}.$$

Thus, to facilitate failures it is desirable to make $\sum_{i \in A} L_i$ and $|A|$ as large as possible, while S_i as small as possible.

Given the multiple constraints involved, this problem shows similarity with the 2-dimensional 0-1 Knapsack problem [12], [23], [25]: Consider a collection of items, where each item i is given a value v_i , has weight w_i , and volume q_i . The objective is then to maximize $\sum_i v_i x_i$ subject to $\sum_i w_i x_i \leq W$ and $\sum_i q_i x_i \leq Q$ where $x_i \in \{0, 1\}$; i.e., we want to choose items with the maximum total value while the total weight is limited by W and total volume is limited by Q . As can be inferred from the discussion above, an important difference of the ER- $k-k'-Q$ problem is that while it is desirable to choose lines with high S_i (could be thought to be analogous to the “value” of the item) under given constraints, it is perhaps equally important to attain or be very close to the limits on both total load and total number of lines attacked.

With these in mind, our heuristics for the ER- $k-k'-Q$ problem are based on incorporating a *switch* to the previously developed heuristics that is actuated to ensure that the attack set A attains or gets close to both constraints on its cardinality and the total load. This idea is inspired from the greedy-like heuristic developed for the multi-dimensional Knapsack problem in [12]. This algorithm initially starts choosing items based on a given set of rules until one or more of the constrained resources become scarce, and then switches to a different set of rules that favor items that use very little of the scarce resource. Here, we propose to use a heuristic that chooses the lines to be attacked one at a time according to the previously developed max- $L * S$ strategy. After each selection, we check whether the switch needs to be activated. Namely, with A' denoting the set of k' lines selected so far, we check

- i) Is it still feasible to select all remaining $k - k'$ lines from the smallest load carrying lines available? Namely, with the remaining lines' loads sorted in ascending order $L^{(1)} \leq L^{(2)} \leq \dots \leq L^{(N-k')}$, we check if

Distribution of L and S	Minimum # of lines to attack to fail all				
	random	max- C	max- L	max- S	best β
$L \sim \text{Pareto}(10, 1.2)$ $S \sim \text{Pareto}(10, 1.2)$	981	151	71	2241	71
$L \sim U(0.4, 100)$ $S \sim U(0.05, 150)$	691	1061	2611	1021	491
$L \sim \text{Pareto}(10, 2.5)$ $S \sim \text{Pareto}(8, 1.2)$	1671	1611	1421	2111	1411
$L \sim \text{Pareto}(10, 1.1)$ $S \sim U(10, 200)$	791	711	3261	2221	541
L, S from the UK National Grid*	2371	1491	1611	2111	1441

TABLE I

Performance comparison of benchmark attacks with the best result of the max- $L * S^\beta$ attack. The first four rows are obtained from Figure 7, while the last row is obtained from simulations with UK National Grid data (see Section VIII for details). Values significantly worse (in the sense of needing to attack many more lines to fail all) than the best- $L * S^\beta$ attack are made bold.

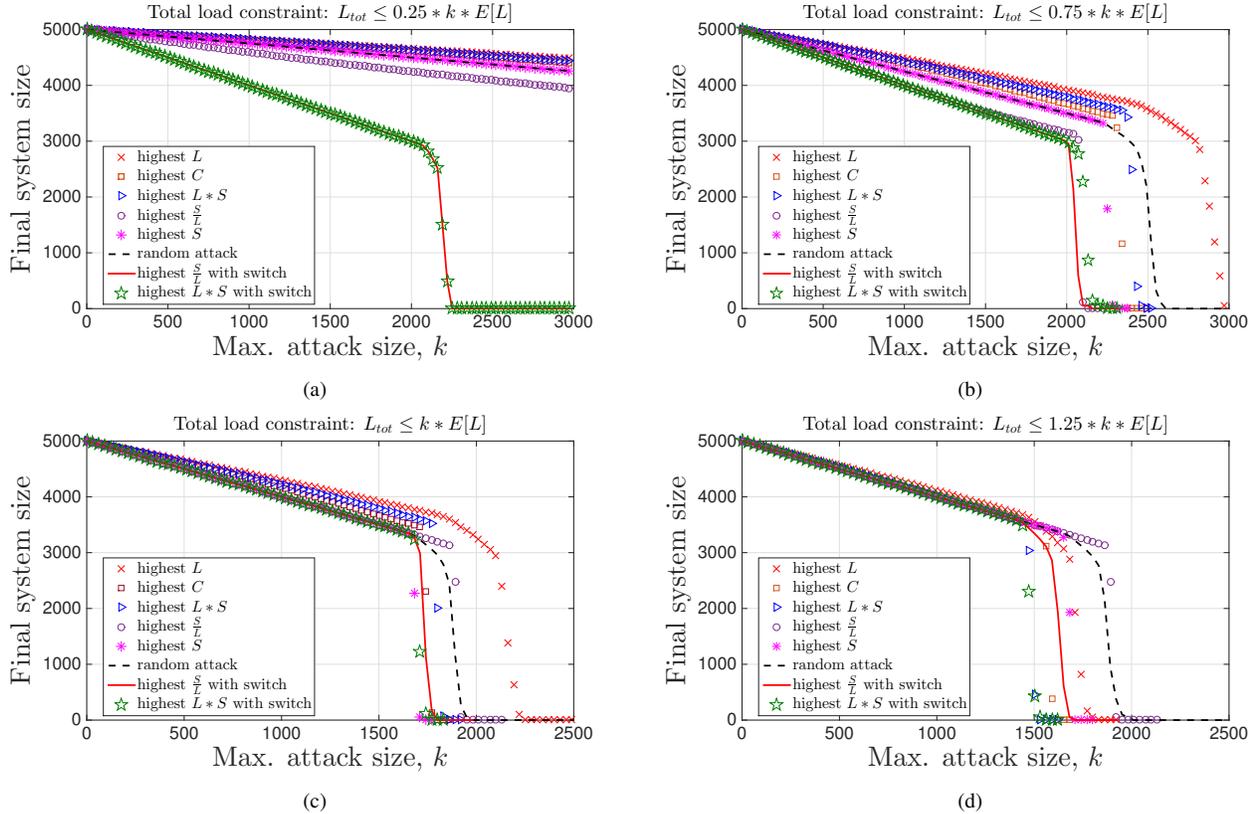


Fig. 8. The performance comparison of different heuristic algorithms for $L \sim U[10, 30]$, $S \sim U[10, 60]$, $N = 5000$, when the attack is constrained to k lines such that their total load satisfies a) $L_{tot} \leq 0.25 * k * \mathbb{E}[L]$; b) $L_{tot} \leq 0.75 * k * \mathbb{E}[L]$; c) $L_{tot} \leq 1.0 * k * \mathbb{E}[L]$; d) $L_{tot} \leq 1.25 * k * \mathbb{E}[L]$.

$$\sum_{j=1}^{k-k'} L^{(j)} + \sum_{i \in A'} L_i \leq Q$$

If the answer is YES, we continue with the second condition for the switch, while if the answer is NO, the switch is activated and algorithm is finished by appending A' with $k - k' - 1$ lines with the smallest load-carrying lines available; i.e., with $\mathcal{L}^{(1)}, \dots, \mathcal{L}^{(k-k'-1)}$. Alternatively, one can release the latest added member of A' and append it with $k - k'$ lines $\mathcal{L}^{(1)}, \dots, \mathcal{L}^{(k-k')}$; we found no major performance difference between these two approaches.

- ii) Next, check whether it is feasible to select all remaining $k - k'$ lines from the largest load carrying lines available. Namely, with the remaining lines' loads sorted in ascending order $L^{(1)} \leq L^{(2)} \leq \dots \leq L^{(N-k')}$, we check if

$$\sum_{N-k+1}^{N-k'} L^{(j)} + \sum_{i \in A'} L_i \leq Q$$

If the answer is NO, we continue the algorithm with the next selection, while if the answer is YES, the switch is activated and algorithm is finished by appending A' with $k - k'$ lines with the largest load-carrying lines available; i.e., with $\mathcal{L}^{(1)}, \dots, \mathcal{L}^{(k-k'-1)}$.

The two conditions of the switch described above ensure that while the initial selections are made in line with the original objectives of picking lines with high load and high free-space, care is also given so as to be able to pick k (or, $k-1$) lines whose total load is close to the limit Q . Of course, any algorithm including the benchmarks can be modified using

the switch idea to better accommodate total load constraints. In particular, when the total load limit Q is extremely stringent, it would be tempting to pick lines with small load so as to not exhaust the total load limit quickly, while aiming to choose lines with high free-space. This prompts us to consider the max- S/L heuristic as well, including its modification with the *switch* idea described above. To keep the discussion brief we do not present results for the max- $L * S^\beta$ attack (with or without switch) and consider only the case where $\beta = 1$; this is in part due to the fact that when the switch is added, the performance of the max- $L * S^\beta$ attack becomes much less sensitive to variations in β over small ranges.

We now present numerical results to evaluate the performance of the heuristic attacks developed here and compare them against benchmark algorithms; as before, we will use max- C , S , L and random attacks as benchmarks. Different from the experiments conducted in Section VII-A, here we have to vary not only the maximum number k of lines that can be attacked, but also the limit Q on the total load of the attacked lines. In particular, we would expect the performance of the algorithms to depend heavily on Q . To this end, we find it meaningful to let Q vary with k and to set it in reference to the *mean* total load of k randomly selected lines; i.e., we set

$$Q = Q(k) = c * k * \mathbb{E}[L] \quad (10)$$

for some constant $c > 0$. This choice enables us to tune c to different levels and check performance in cases where (i) the total load is extremely limited (i.e., $c \ll 1$); ii) total load limit is such that heuristics that do not take load into account (such as max- S or random attacks) will likely be able to select

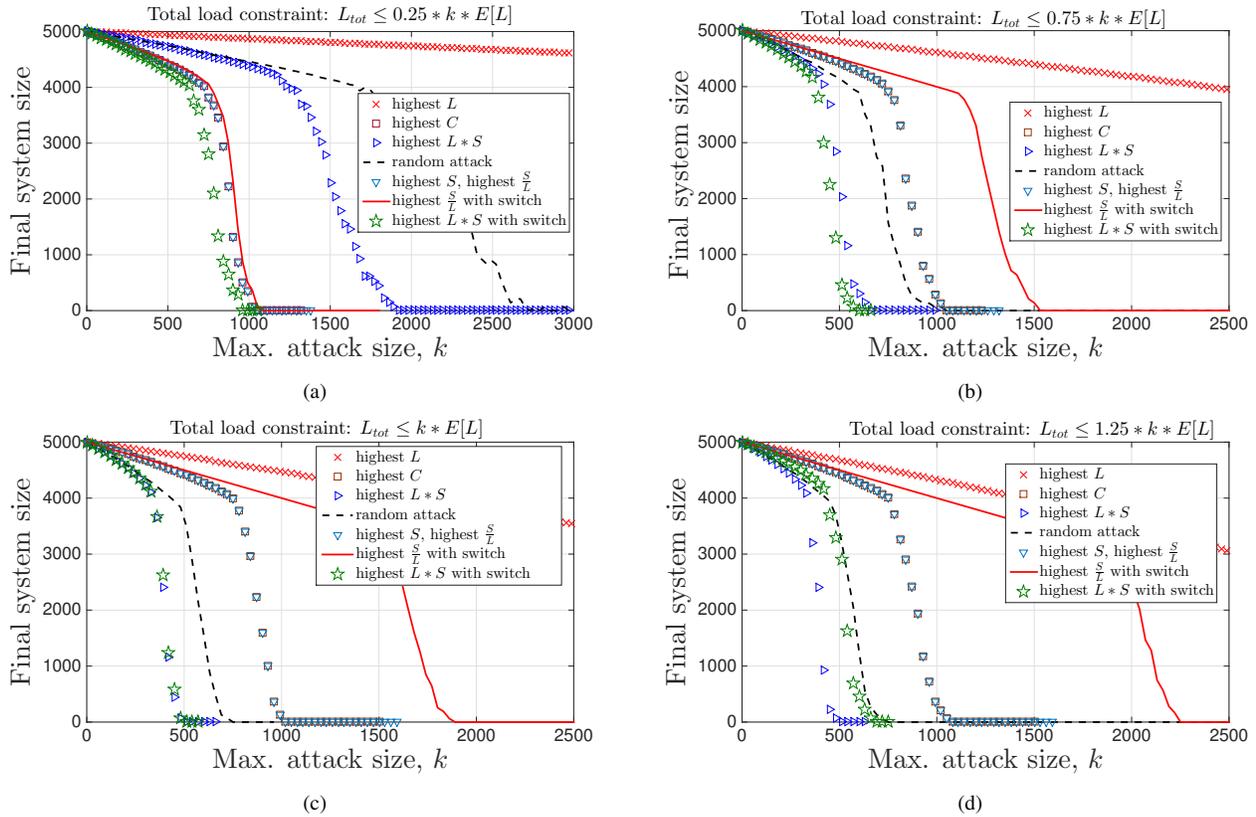


Fig. 9. The performance comparison of different heuristic algorithms for $L \sim U[0.4, 100]$, $S \sim U[0.05, 150]$, with L and S sorted in reverse order, $N = 5000$, when the attack is constrained to k lines such that their total load satisfies a) $L_{tot} \leq 0.25 * k * \mathbb{E}[L]$; b) $L_{tot} \leq 0.75 * k * \mathbb{E}[L]$; c) $L_{tot} \leq 1.0 * k * \mathbb{E}[L]$; d) $L_{tot} \leq 1.25 * k * \mathbb{E}[L]$. Each data point is obtained by averaging over 100 independent runs.

(close to) k lines with total load very close to Q (i.e., $c \approx 1$); or iii) total load limit is not stringent at all (i.e., $c \gg 1$) and the problem is similar to the unconstrained load case.

In the first set of experiments, we set $N = 5000$ and generate load-free space values independently from the distributions $L \sim U[10, 30]$ and $S \sim U[10, 60]$. For brevity we consider four values of c given at (10): $c = 0.25$, $c = 0.75$, $c = 1.0$, $c = 1.25$. The results are presented in Figure 8 from which a number of interesting observations can be made. When $c = 0.25$, i.e., when total load is extremely constrained, we see that all heuristics without a switch perform poorly and are not able to fail the whole system even until $k = 3000$. This can be attributed to their inability to attack the maximum allowed number k of lines as they quickly exhaust the total load limit. On the other hand, we see that both max- $L * S$ -with-switch and max- S/L -with-switch attacks perform much better, and despite the stringent limit on the total load are able to fail the system by attacking about 50% more lines than required in the case where the total load is unlimited. When c is increased to 0.75, we see that the performance of the benchmarks improve but still are significantly worse than the two heuristics that use the switch; in this case we also see that the max- S/L -with-switch attack slightly outperforms max- $L * S$ -with-switch.

With $c = 1$, we see that algorithms that ignore the loads such as max- S and random attacks perform as well as they do in the unconstrained case; this is expected by virtue of the law of large numbers. In particular, when $c = 1$, we would expect max- S to perform well since it picks the most robust lines in the system and is likely to reach the limits k and Q

simultaneously given that L is independent from S . Figure 8(c) confirms this intuition where we also see that both heuristics with switch match the performance of the max- S attack. Finally, with $c = 1.25$, we get close to the unconstrained load case, and as expected see that the performance of the max- $L * S$ algorithm becomes the best. What is interesting here is that the max- $L * S$ -attack-with-switch is able to match this performance, showing its versatile performance across very different cases considered here. Overall, these experiments demonstrate that incorporating the *switch* significantly improves the performance and the max- $L * S$ -attack performs well across different ranges of the total load limit.

As in Section VII-A, it is of interest to check the performance of these algorithms in difficult cases where load and free-space values are sorted in reverse orders. To this end, we consider one of the settings used in Figure 7, and generate load and free-space independently from $L \sim U[0.4, 100]$ and $S \sim U[0.05, 150]$, and then sort them in reverse orders so that the line with maximum load gets the smallest free-space and so on and so forth. In this setting, max- S and max- S/L heuristics become equivalent. Also, since large S values are around 150 while L is limited to 100, the lines with maximum- C will be those with large S (and small L due to reverse ordering). As seen in Figure 9, this leads to three benchmarks (max- S , C , S/L) performing almost equally in this setting.

We see from Figure 9 that the max- $L * S$ -attack-with-switch once again performs well. It leads to the best performance among all heuristics considered for $c = 0.25$, $c = 0.75$, and $c = 1.0$ (equal with max- $L * S$), while coming second for

$c = 1.25$ after $\max-L * S$; this is expected since without a stringent limit on total load, the problem gets closer to the unconstrained case where a switch is not needed. Also, in this case we see that the performance of the $\max-S$ attack (along with $\max-S/L$ and $\max-C$ attacks) is rather unaffected by the load constraint. We attribute this to the fact that since load and free-space are reverse ordered, targeting $\max-S$ lines is equal to targeting $\min-L$ lines and even when $c = 0.25$ the total load limit is not likely to be exhausted easily; i.e., the algorithm is able to choose k lines without exceeding the total load limit. We think this is also the reason why the *switch* is not helping (and, actually hurting) the $\max-S/L$ algorithm in this case.

VIII. SIMULATIONS WITH UK NATIONAL GRID DATA

In this section, we provide simulation results illustrating how the attack strategies covered here performs when the load and free space distribution are based on real data. We have used National Grid Electricity Ten Year Statement 2016 Model of Great Britain [15] to generate load-free space pairs. To be more precise, the load distribution is chosen from the winter peak power flow diagram presented in [15, Appendix C]. For the free-space distribution, the transmission line ratings given in [15, Appendix B] has been used. As in the case of previous examples, the number of lines is taken as $N = 5000$.

In Figure 10, we show the performance of the heuristic algorithms for the unconstrained case, i.e., for the ER- k problem. The results are very similar to those obtained under synthetic load-free space distributions and demonstrate that heuristics developed here perform well under real-world distributions as well. In particular, we see that the proposed $\max-L * S$ heuristic performs better than all benchmarks considered, and its performance can further be improved by the $\max-L * S^\beta$ attack. For the UK National Grid data, our results indicated that the best performance is obtained when $\beta = 1.5$; see also the last row in Table I.

Next, we consider the ER- $k-k'-Q$ problem where the total load of the attack set is bounded by Q . As in Figures 8 and 9, we set Q according to (10) for several c values, and compare the performance of our attack strategies with benchmarks; this time the load-free space distribution is set according to the aforementioned UK National Grid data. The results are depicted in Figure 11. Once again we see that the $\max-L * S$ attack *with switch* leads to the best overall performance among

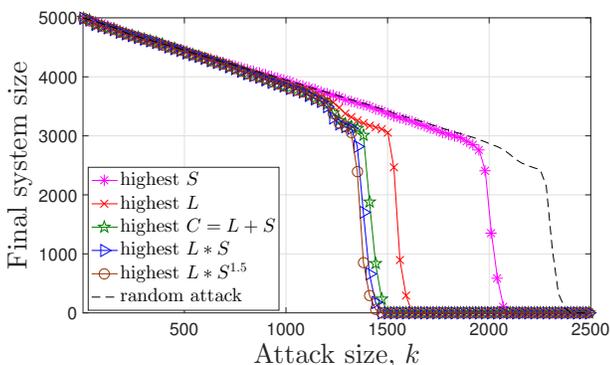


Fig. 10. The performance comparison of different heuristic algorithms when L, S follow the UK National Grid data [15].

all heuristics considered; it leads to the best performance when $c = 0.25$ (tied with $\max-S/L$ with switch), $c = 0.75$, and $c = 1.25$, while coming as second for $c = 1$ after $\max-S/L$. This shows that the heuristic attacks proposed here have versatile performance also under distributions observed in real-world power systems.

IX. CONCLUSION

In this paper, we consider a flow-redistribution based model to understand cascading failures in load-carrying networks, e.g., electrical power systems, transportation systems, etc. We focus on assessing the vulnerability of such systems against adversarial attacks. In particular, in a system with N lines with initial loads L_1, \dots, L_N and capacities C_1, \dots, C_N , we study the constrained optimization problem of finding k initial lines to be attacked to minimize the final number of alive lines in the system. We give optimal greedy algorithms in several special cases, and prove that a variation of the problem (with a bound on the total load of the initial attack set) is NP-Hard. Several heuristics are developed and their performance is compared with benchmark attacks under various settings. Overall, it is seen that the system is most vulnerable against attacks that target lines with maximum load free-space product $L * (C - L)$.

There are many interesting directions to consider for future work. First, the complexity of the optimal k -attack problem (without a bound on the total load) is not known. Also, with the results of this paper revealing *good* attack strategies, one might now seek optimal design strategies (e.g., in the form of load-capacity distributions) that lead to maximum robustness against such attacks. It would also be interesting to consider new cascading failure models 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 last problem, we provide preliminary simulation results in Supplementary Material. They suggest that the mean-field (i.e., equal) redistribution model captures the qualitative behavior of system robustness well. Finally, it might be interesting to study *information cascades* in social networks [27], [30], [31], [33] using the models considered here; the optimal attack problem studied here will then amount to *influence maximization* problem [10].

REFERENCES

- [1] J. V. Andersen, D. Sornette, and K.-t. Leung. Tricritical behavior in rupture induced by disorder. *Phys. Rev. Lett.*, 78:2140–2143, Mar 1997.
- [2] G. Andersson, P. Donalek, R. Farmer, N. Hatziairgyriou, I. Kamwa, et al. Causes of the 2003 major grid blackouts in north america and europe, and recommended means to improve system dynamic performance. *Power Systems, IEEE Transactions on*, 20(4):1922–1928, 2005.
- [3] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin. Catastrophic cascade of failures in interdependent networks. *Nature*, 464:1025–1028, 2010.
- [4] E. Chatzifratris, Y. Zhang, and O. Yağan. On the robustness of power systems: optimal load-capacity distributions and hardness of attacking. In *Information Theory and Applications Workshop (ITA)*, 2016.
- [5] P. Crucitti, V. Latora, and M. Marchiori. Model for cascading failures in complex networks. *Phys. Rev. E*, 69:045104, Apr 2004.
- [6] H. Daniels. The statistical theory of the strength of bundles of threads. i. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 183(995):405–435, 1945.

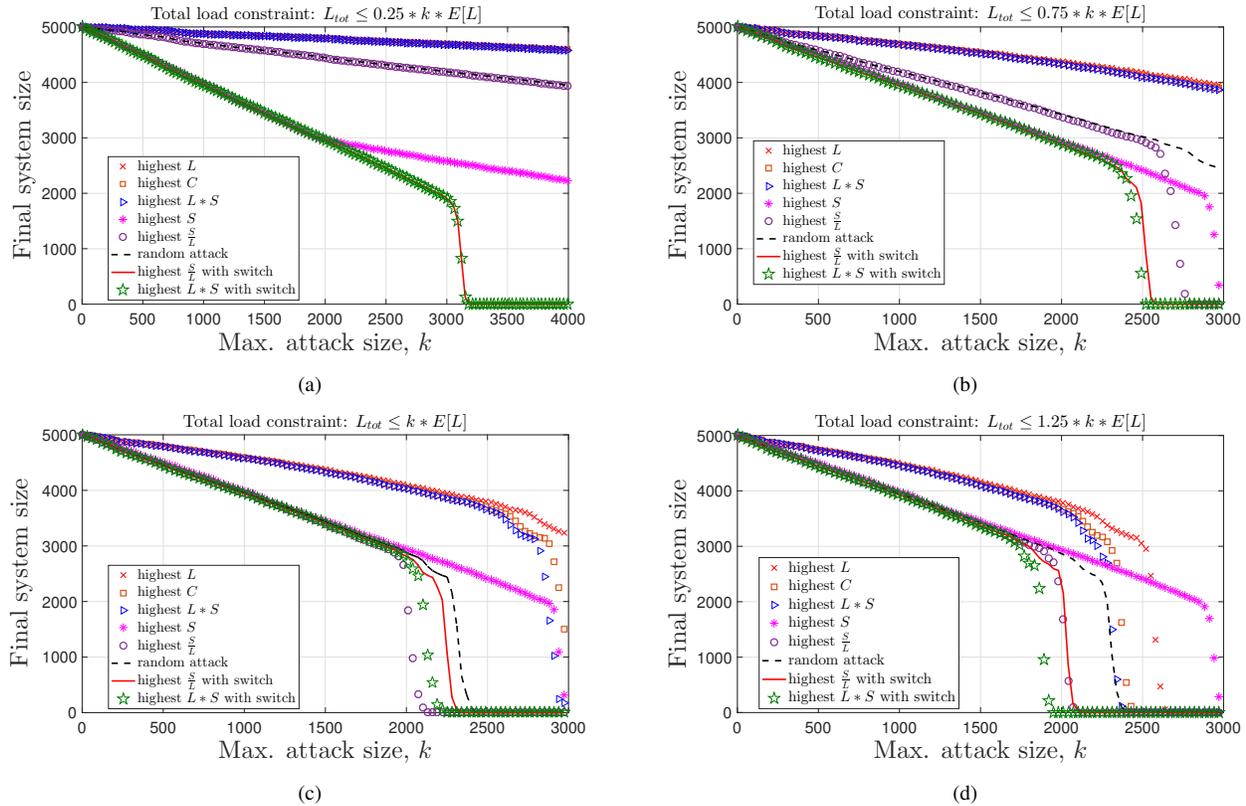


Fig. 11. The performance comparison of different heuristic algorithms when L, S pairs are distributed according to the UK National Grid data [15]. The attack is constrained to k lines such that their total load satisfies a) $L_{tot} \leq 0.25 * k * \mathbb{E}[L]$; b) $L_{tot} \leq 0.75 * k * \mathbb{E}[L]$; c) $L_{tot} \leq 1.0 * k * \mathbb{E}[L]$; d) $L_{tot} \leq 1.25 * k * \mathbb{E}[L]$.

- [7] L. Daqing, J. Yanan, K. Rui, and S. Havlin. Spatial correlation analysis of cascading failures: congestions and blackouts. *Sci. Reps.*, 4, 2014.
- [8] I. Dobson, B. A. Carreras, V. E. Lynch, and D. E. Newman. Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization. *Chaos*, 17(2):026103, 2007.
- [9] I. Dobson, J. Chen, J. Thorp, B. A. Carreras, and D. E. Newman. Examining criticality of blackouts in power system models with cascading events. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*, 2002.
- [10] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proc. of ACM SIGKDD Conference*, pages 137–146, 2003.
- [11] R. Kinney, P. Crucitti, R. Albert, and V. Latora. Modeling cascading failures in the north american power grid. *The European Physical Journal B*, 46(1):101–107, 2005.
- [12] R. Loulou and E. Michaelides. New greedy-like heuristics for the multidimensional 0-1 knapsack problem. *Operations Research*, 27(6):1101–1114, 1979.
- [13] B. Mirzasoleiman, M. Babaei, M. Jalili, and M. Safari. Cascaded failures in weighted networks. *Physical Review E*, 84(4):046114, 2011.
- [14] A. E. Motter and Y.-C. Lai. Cascade-based attacks on complex networks. *Phys. Rev. E*, 66:065102, Dec 2002.
- [15] NationalGrid. <http://www2.nationalgrid.com/uk/industry-information/future-of-energy/electricity-ten-year-statement/>. Technical report, 2016.
- [16] T. D. O'Rourke. Critical infrastructure, interdependencies, and resilience. *BRIDGE-WASHINGTON-NATIONAL ACADEMY OF ENGINEERING-*, 37(1):22, 2007.
- [17] T. J. Overbye, X. Cheng, and Y. Sun. A comparison of the ac and dc power flow models for Imp calculations. In *Proceedings, 37th Hawaii International Conference on System Sciences*, 2004.
- [18] O. Ozel, B. Sinopoli, and O. Yağan. Robustness of flow networks against cascading failures under partial load redistribution. *arXiv:1802.07664*.
- [19] S. Pahwa, C. Scoglio, and A. Scala. Abruptness of cascade failures in power grids. *Scientific reports*, 4, 2014.
- [20] S. Pradhan and B. K. Chakrabarti. Failure properties of fiber bundle models. *Int'l. Journal of Modern Physics B*, 17(29):5565–5581, 2003.
- [21] M. Rosas-Casals and R. Solé. Analysis of major failures in european power grid. *International Journal of Electrical Power & Energy Systems*, 33(3):805–808, 2011.
- [22] C. Roy, S. Kundu, and S. Manna. Fiber bundle model with highly disordered breaking thresholds. *Physical Review E*, 91(3):032103, 2015.
- [23] S. Senju and Y. Toyoda. An approach to linear programming with 0-1 variables. *Management Science*, pages B196–B207, 1968.
- [24] B. Stott, J. Jardim, and O. Alsac. Dc power flow revisited. *Power Systems, IEEE Transactions on*, 24(3):1290–1300, Aug 2009.
- [25] Y. Toyoda. A simplified algorithm for obtaining approximate solutions to zero-one programming problems. *Management Science*, 21(12):1417–1427, 1975.
- [26] W.-X. Wang and G. Chen. Universal robustness characteristic of weighted networks against cascading failure. *Phys. Rev. E*, 77, 2008.
- [27] D. J. Watts. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99:5766–5771, 2002.
- [28] Y. Yang, T. Nishikawa, and A. Motter. Small vulnerable sets determine large network cascades in power grids. *Science*, 358(6365), 2017.
- [29] O. Yağan. Robustness of power systems under a democratic-fiber-bundle-like model. *Phys. Rev. E*, 91:062811, Jun 2015.
- [30] O. Yağan and V. Gligor. Analysis of complex contagions in random multiplex networks. *Physical Review E*, 86(3):036103, 2012.
- [31] O. Yağan, D. Qian, J. Zhang, and D. Cochran. Information diffusion in overlaying social-physical networks. In *Proc. of the 46th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6, March 2012.
- [32] O. Yağan, D. Qian, J. Zhang, and D. Cochran. Optimal allocation of interconnecting links in cyber-physical systems: Interdependence, cascading failures, and robustness. *IEEE Transactions on Parallel and Distributed Systems*, 23(9):1708–1720, Sept 2012.
- [33] O. Yağan, D. Qian, J. Zhang, and D. Cochran. Conjoining speeds up information diffusion in overlaying social-physical networks. *IEEE Journal on Selected Areas in Communications*, 31(6):1038–1048, 2013.
- [34] G. Zhang, Z. Li, B. Zhang, and W. A. Halang. Understanding the cascading failures in indian power grids with complex networks theory. *Physica A*, 392(15):3273–3280, 2013.
- [35] Y. Zhang, A. Arenas, and O. Yağan. Cascading failures in interdependent systems under a flow redistribution model. *Phys. Rev. E*, 97:022307, Feb 2018.
- [36] Y. Zhang and O. Yağan. Optimizing the robustness of electrical power systems against cascading failures. *Scientific reports*, 6:27625, 2016.

Attack vulnerability of power systems under an equal load redistribution model (Supplementary Material)

Talha Cihad Gulcu, Vaggos Chatziafratis, Yingrui Zhang, and Osman Yağan

APPENDIX

A. A proof of Theorem 5.5

A key observation is that since the failed load is always redistributed *equally* among *alive* lines, the system will preserve the “equal load” property through the cascading failures. Namely, at any stage the load of a line that is functioning will be given by $L(1 + \frac{M}{N-M})$ where M is the number of lines (out of N) failed thus far. In addition the sequence of attacks does not affect the final state of the system as discussed before. Therefore, the claim would follow for general attack sizes k , if we establish it for $k = 1$. This is because after a single line is attacked, the system will again be one with equal loads and the optimization problem will repeat itself with $k-1$ additional lines to be attacked. Continuing in this manner, we see that the optimal attack strategy in this case would be a combination of *optimal* single-line attacks launched sequentially.

Assume now that $k = 1$, i.e., the goal is to attack the line that will lead to the maximum number of failed lines. Since all loads are equal, the lines that fail *initially* as a result of this attack will be (with A_0 denoting the line chosen)

$$proj(A_0) := \{A_0\} \cup \left\{ \ell \in \{1, 2, \dots, N\} / A_0 : C_\ell \leq LN/(N-1) \right\}$$

With \mathcal{L}_i and \mathcal{L}_j denoting arbitrary distinct lines, we have

$$|proj(\mathcal{L}_i)| - |proj(\mathcal{L}_j)| = \mathbf{1} [C_j \leq LN/(N-1)] - \mathbf{1} [C_i \leq LN/(N-1)]$$

which automatically gives $|proj(\mathcal{L}_i)| \geq |proj(\mathcal{L}_j)|$ if $C_i \geq C_j$. Since i and j are arbitrary, this shows that

$$\arg \max_{i=1, \dots, N} proj(\mathcal{L}_i) = \arg \max_{i=1, \dots, N} C_i,$$

meaning that attacking the max- C line maximizes the number of lines failed initially. Also, with all loads equal, attacking the max- C line is equivalent to attacking the max- S line so this

This research was supported by the National Science Foundation through grant CCF #1422165 and a generous gift from *Persistent Systems, Inc.*

Preliminary version of some of the material was presented at the Information Theory and Applications Workshop (ITA 2016), San Diego, CA [1].

T. C. Gulcu is with TUBITAK Advanced Technologies Research Institute, 06800, Umitkoy, Ankara, Turkey. E-mail: tcgulcu@gmail.com

V. Chatziafratis is with the Computer Science Department, Stanford University, Palo Alto, CA 94305. Email: vaggos@stanford.edu

Y. Zhang and O. Yağan are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213. E-mail: {yingruiz, oyagan}@andrew.cmu.edu

attack also collectively minimizes the free space of remaining lines. In other words, with $i^* = \arg \max_{i=1, \dots, N} C_i$, we have

$$\{S_\ell\}_{\ell \notin proj(\mathcal{L}_{i^*})} \leq \{S_\ell\}_{\ell \notin proj(\mathcal{L}_j)}$$

for any $j = 1, \dots, N$ where the comparison is made pairwise. Combining, we see that the max- C attack maximizes the number of initially failed lines and minimizes the free-space of those that remain. Since all lines have equal load, this then shows that the final number of failed at the end of a potential cascade will be minimized by attacking the max- C line.

B. A proof of Theorem 5.6

As before, a key observation is that the optimization problem can be reduced to finding the optimal single-line attack, and repeating this recursively. The reason is that since failed load is equally redistributed, the system will maintain to have the same free space (among all alive lines) throughout the cascade process. Given also that the order of the attack does not matter, it remains to find the optimal single-line attack, i.e., the case where $k = 1$. Here, we have

$$proj(A_0) = \{A_0\} \cup \{ \ell \in \{1, 2, \dots, N\} / A_0 : S_\ell \leq L_{A_0}/N \}$$

where A_0 is the line that is attacked. Since all lines have equal free space S , this gives

$$|proj(A_0)| = N \mathbf{1} [S \leq L_{A_0}/N]$$

showing that there can be no cascades in this case; after an attack either all lines fail simultaneously or no single line other than those directly attacked fails. It is clear that $|proj(A_0)|$ is monotone increasing in the load L_{A_0} of A_0 so that it is maximized by attacking the line with the maximum load. Repeating this argument recursively, we see that the max- L greedy strategy is optimal for the ER- k -Same-Space problem.

C. A proof of Theorem 5.7

The key observation about the ER- k -($C \propto L$) Problem is that, given $C_i = (1 + \alpha)L_i$, the load, capacity, and free spaces of lines all follow the same order. Namely, the lines with the largest loads, who shall intuitively be attacked to shed more load on others, are also the ones with the largest free spaces, who are also good to attack given the difficulty of failing them by load redistribution. This eliminates the trade-off faced in the optimal attack problem and simplifies it greatly.

In this setting, the problem does not repeat itself since after load redistribution, it may no longer be the case that all

lines have the same tolerance factor (i.e., S/L). However, the aforementioned key property will be maintained throughout. For instance, assume without loss of generality that initial loads are ordered as $L_1 \leq L_2 \leq \dots \leq L_N$. Then, at any stage of the cascading failures, L_i , C_i , and S_i will all be in increasing order for all $i = 1, 2, \dots, N$ that are still alive.

With these in mind, we can show the optimality of *max-L, C, S-Greedy Algorithms* for single line attacks in any system whose loads L_1, L_2, \dots, L_N and free spaces S_1, \dots, S_N follow the same order. Since this property will be preserved throughout the cascades and the sequence of attacks doesn't affect the final state of the system, the proof of Theorem 5.7 will be completed. The rest of the proof is similar to that of Theorem 5.5 and omitted here for brevity.

D. Simulation results under a topology-based load redistribution model

The results presented in this paper are obtained under the *mean field* assumption that when a line fails, its load gets redistributed globally and equally among all functional lines in the network. In doing so, our motivation was to get insights on the vulnerability of power systems (in terms of the set of lines that are most likely to be attacked by a knowledgeable adversary) through a simplified model that captures the qualitative behavior of system robustness.

One way to check whether the mean-field assumption leads to useful results is to run simulations under different redistribution models and compare them against the equal redistribution model. A popular redistribution model [2], [3], [4] from the literature considers a graph topology over which the power system or the flow network is defined, and assumes that whenever a line fails, its load gets redistributed (oftentimes in an equal manner, but in some cases based on some weights) only in the neighborhood of the failure point; e.g., when line i fails, all neighboring lines receive an equal portion of the failed load L_i . As mentioned in the paper, such local redistribution models fail to capture the long-range nature of Kirchhoff's Law and thus it would be more appropriate to consider a model where the failed load is redistributed both locally and globally.

With these in mind, we consider a model that combines the equal redistribution model described in Section II and the purely topology-based redistribution models (e.g., similar to studied in [3]). More precisely, we assume that upon failures, a γ -fraction of the failed flow is redistributed solely in the local neighborhood of the failed line (according to the network topology), while the rest gets redistributed among *all* functional lines. With this approach, we recover the case studied originally in our paper when $\gamma = 0$, while setting $\gamma = 1$ gives the merely topology-based model.

Under this redistribution model, we have run simulations for a specific case of network topology to see how the results presented in Figure 5 are affected under different γ values. In particular, we generate graph topologies randomly according to the Erdős-Rényi $\mathbb{G}(n; M)$ model, where n denotes the number of nodes and M denotes the number of edges. According to the $\mathbb{G}(n; M)$ model, the M edges are selected uniformly at random from among all possible $\binom{n}{2}$ edges. We

set the number of edges (i.e., lines) at $M = 5000$ so that a fair comparison can be made with the setting of Figure 5. For simplicity, we set $n = 230$ which leads to each edge having a mean number 85 of neighbors; two edges are deemed to be neighbors if they have a common end node. Also, the load and free-space values of these 5000 lines are drawn independently from the distributions $L \sim U[10, 30]$, $S \sim U[10, 60]$, as in the set-up of Figure 5.

The results are depicted in Figure 1, where each data point represents the average obtained over 100 independent experiments. We see that the mean-field approach (i.e., the case with $\gamma = 0$) is able to capture very well the *qualitative* behavior of the final system size for all γ values and for the attack strategies considered. As would be expected, larger values of γ make the system more vulnerable against attacks by virtue of (a portion of) failed flow getting redistributed among smaller number of (neighboring) lines, thereby increasing the chance of further overloading and cascading failures. However, the qualitative behavior of the final system size as a function of the attack size is seen to be unaffected by γ . In particular, we see that the conclusions drawn under the mean-field model (i.e., $\gamma = 0$) concerning how different attack strategies compare with each other remains the same, even under the fully topology-based redistribution model (i.e., $\gamma = 1$). For instance, irrespective of γ , the highest $L * S$ attack performs the best among other strategies considered.

An important remark here is that, even though the alternative redistribution model discussed above is based on the network topology, the highest- $L * S$ attack strategy (or, other strategies considered in this paper) does not take the network topology into account. We believe that it is this aspect that leads to the mean-field assumption still being reasonably accurate under different γ values. If the attacker picks the set of lines taking both L, S values and network topology into account, we would expect the mean-field assumption to fail (at least for certain attack strategies). This is because of the fact that we can no longer state that all non-attacked lines are indistinguishable from each other and are equally likely to share the failed flow.

Finally, it is worth noting that an analogous version of Theorem 6.8 can be easily established under the topological redistribution model discussed above. This is because when the underlying network topology is one where every pair of lines are neighbors with each other, the failed load will be redistributed equally across all non-failed lines, irrespective of the value of γ . In other words, the topological redistribution model will be equivalent to the equal redistribution (i.e., mean-field) model for such networks. This means that the ER- k - Q problem (shown to be NP-Complete in Theorem 6.8) constitutes a *special* case of the analogous problem defined under the topological redistribution model; thus, the analogous problem (for general network topologies) is also NP-Complete.

REFERENCES

- [1] E. Chatziafratis, Y. Zhang, and O. Yağan. On the robustness of power systems: optimal load-capacity distributions and hardness of attacking. In *Information Theory and Applications Workshop (ITA)*, 2016.
- [2] B. Mirzasoleiman, M. Babaei, M. Jalili, and M. Safari. Cascaded failures in weighted networks. *Physical Review E*, 84(4):046114, 2011.

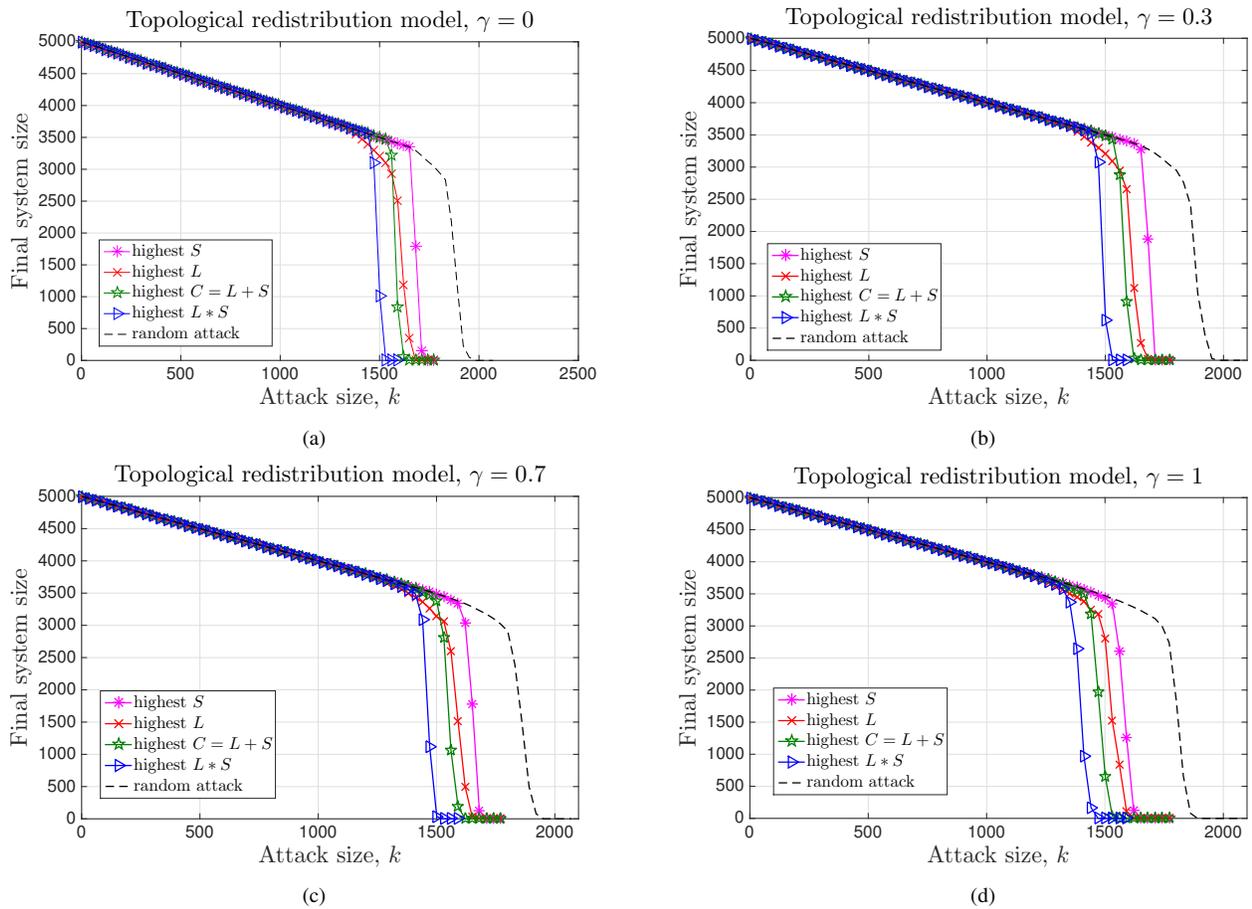


Fig. 1. The comparison of attack strategies under different flow redistribution models. Here, γ controls the fraction of flow that will be redistributed (upon failure) locally according to the network topology (an Erdős-Rényi network, $\mathbb{G}(n; M)$ with $M = 5000$ lines, each connected on average to 85 other lines). The remaining $1 - \gamma$ fraction of the failed flow gets redistributed globally among all functioning lines. As in Figure 5, we set $L \sim U[10, 30]$ and $S \sim U[10, 60]$.

- [3] A. E. Motter and Y.-C. Lai. Cascade-based attacks on complex networks. *Phys. Rev. E*, 66:065102, Dec 2002.
- [4] W.-X. Wang and G. Chen. Universal robustness characteristic of weighted networks against cascading failure. *Phys. Rev. E*, 77, 2008.