

Power Grid Vulnerability to Geographically Correlated Failures – Analysis and Control Implications

Andrey Bernstein^{*†}, Daniel Bienstock[‡], David Hay[§], Meric Uzunoglu^{*}, and Gil Zussman^{*}

^{*}Department of Electrical Engineering, Columbia University, New York, NY 10027

[†]Department of Electrical Engineering, Technion, Haifa 32000, Israel,

[‡]Department of Applied Physics and Applied Math, Columbia University, New York, NY 10027

[§]School of Engineering and Computer Science, Hebrew University, Jerusalem 91904, Israel

Email: andreyb@techunix.technion.ac.il, dano@columbia.edu, dhay@cs.huji.ac.il,

meric.uzunoglu@gmail.com, gil@ee.columbia.edu

Abstract—We consider power line outages in the *transmission system* of the power grid, and specifically those caused by a natural disaster or a large scale physical attack. In the transmission system, an outage of a line may lead to overload on other lines, thereby eventually leading to their outage. While such cascading failures have been studied before, our focus is on cascading failures that follow an *outage of several lines in the same geographical area*. We provide an analytical model of such failures, investigate the model's properties, and show that it differs from other models used to analyze cascades in the power grid (e.g., epidemic/percolation-based models). We then show how to identify the most vulnerable locations in the grid and perform extensive numerical experiments with real grid data to investigate the various effects of geographically correlated outages and the resulting cascades. These results allow us to gain insights into the relationships between various parameters and performance metrics, such as the size of the original event, the final number of connected components, and the fraction of demand (load) satisfied after the cascade. In particular, we focus on the timing and nature of *optimal control* actions used to reduce the impact of a cascade, in real time. We also compare results obtained by our model to the results of a real cascade that occurred during a major blackout in the San Diego area on Sept. 2011. The analysis and results presented in this paper will have implications both on the design of new power grids and on identifying the locations for shielding, strengthening, and monitoring efforts in grid upgrades.

Index Terms—Power Grid, Geographically-Correlated Failures, Cascading Failures, Resilience, Survivability.

I. INTRODUCTION

Recent colossal failures of the power grid (such as the Aug. 2003 blackout in the Northeastern United States and Canada [2], [37]) demonstrated that large-scale and/or long-term failures will have devastating effects on almost every aspect in modern life, as well as on interdependent systems (e.g., telecommunications, gas and water supply, and transportation). Therefore, there is a need the study the vulnerability of the existing power transmission system and to identify ways to mitigate large-scale blackouts.

The power grid is vulnerable to natural disasters, such as earthquakes, hurricanes, floods, and solar flares [38] as well as

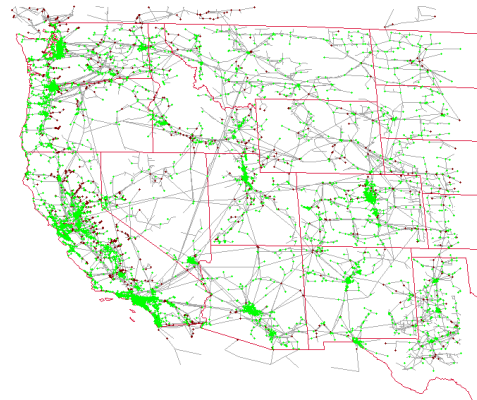


Fig. 1. The graph that represents part of the Western Interconnect and sections of the Texas, Oklahoma, Kansas, Nebraska, and the Dakotas' grids. Green dots are demands (loads), red dots are supplies (generators), and neutral nodes are not shown. This graph was used to derive numerical results.

to physical attacks, such as an Electromagnetic Pulse (EMP) attack [17], [18], [38]. Thus, we focus the vulnerability of the power grid to an outage of several lines in the same geographical area (i.e., to *geographically-correlated failures*). Recent works focused on identifying a *few vulnerable lines* throughout the entire network [7], [8], [32], on designing line or node interdiction strategies [9], [34], and on characterizing the network graph and studying probabilistic failure propagation models [12], [16], [23], [25], [39], [41]. However, our objective is to identify *the most vulnerable areas* in the power grid, and examine appropriate real-time control countermeasures to minimize the impact of an event of this type. Detection of the most vulnerable areas has various practical applications, since the system in these areas can be either shielded (e.g., against EMP attacks or solar flares), strengthened (e.g., by increasing the capacities of some relevant lines), or monitored (e.g., as part of smart grid upgrade projects). Real-time control will likely still be needed, in case the pre-strengthening of the grid has overlooked a particular “attack” pattern.

Unlike graph-theoretical network flows, the flow in the

power grid is governed by the laws of physics and there are *no strict capacity bounds on the lines* [6], [16], [27]. On the other hand, there is a *rating threshold* associated with each line, such that when the flow through a line exceeds the threshold, the line heats up and eventually faults. Such an outage, in turn, causes another change in the power grid, that can eventually lead to a *cascading failure*. We describe the *linearized (i.e., DC) power flow model* and the *cascading failure model* (originated from [13] and extended in [7], [8]) that allow us to obtain analytic and numerical results despite the problem’s complexity. We note that severe cascading failures are *hard to control in real time* [4], [7], [8], [16], [31], since the power grid optimization and control problems are of enormous size. Thus, in our numerical experiments, we apply basic control mechanisms that shed demands at the end of each round.

We initially use the linearized power flow model to study some simple motivating examples and to derive analytical results regarding the cascade propagation in simple ring-based topologies. We note that several previous works (e.g., [12], [25], [30], [35], [41] and references therein) assumed that a line or node failure leads with some probability to a failure of nearby nodes or lines. Such epidemic-based modeling allows using percolation-based tools to analyze the effects of a cascade. However, we show that using the more realistic power flow model leads to failure propagation characteristics that are significantly different. Specifically, we show that a failure of a line can lead to a failure of a line M hops away (with M arbitrarily large). This result is of particular importance, since it has been observed in real cascades (as the recent one in the San Diego area, discussed in Section VIII). Moreover, we prove that cascading failures can last arbitrarily long time and that a network whose topology is a subgraph of another topology can be more resilient to failures.

In this paper, we focus on contingency events in a grid that are initiated by geographically correlated failures. We represent the area affected by a contingency by a disk. Since such a disk can theoretically be placed in an infinite number of locations, we briefly discuss an efficient computational geometric method (which builds on results from [1]) that allows identifying a finite set of locations that includes all possible failure events.

In our numerical experiments, we use the WECC (Western-Interconnect) real power grid data taken from the Platts Geographic Information System (GIS) [33] (the resulting graph appears in Figure 1). We present extensive numerical results obtained by simulating the cascading failures for each of the possible disk centers (the results have been obtained by repeatedly and efficiently solving very large systems of equations). When only few lines initially fault, cascading failures usually start slowly and intensify over time [7], [8]. However, when the failures are geographically correlated, we notice that in many cases this *slow start* phenomenon does not exist. We illustrate the effects of the most (and the least) devastating failures and show the yield (the overall reduction in power generation) for all different failure locations in the

Western US. Moreover, we identify various relations between parameters and performance metrics (such as yield, number of components into which the network partitions, and number of faulted lines which corresponds to the length of the repair process). We also study the sensitivity of the results to different failure models (stochastic vs. deterministic model), to the value of the Factor of Safety (the ratio between line capacity and normal flow on the line), and to whether or not the line capacities are derived based on $N - 1$ contingency plan.

While large scale cascades are quite rare [13], during our work on this report, a cascade took place in the San Diego area (on Sept. 8, 2011). The causes of this cascade are still under investigation. Yet, we have been able to use some preliminary information [11] to assess the accuracy of our method and parameters. We conclude by discussing some of the numerical results that we obtained for a similar scenario and demonstrating that our methods have the potential to identify vulnerable parts of the network.

Finally, we consider optimal control actions to be taken in the event of such a failure. We show, experimentally, that appropriate action taken at the appropriate time (and not necessarily at the start of the cascade) can rapidly stop the process, while losing a minimum quantity of demand.

Below are the main contributions of this paper. First, we obtain analytical results regarding network vulnerability and resilience under the power flow model which is significantly different from the classical network flow models. These results provide *insights that significantly differ from insights obtained from epidemic/percolation-based models*. Second, we combine techniques from optimization, computational geometry, and communications network vulnerability analysis to develop a method that allows obtaining extensive numerical results regarding the effects of geographically correlated failures on a real grid. To the best of our knowledge, *this is the first attempt to obtain such results using real geographical data*. Third, we briefly illustrate that control actions have the potential to mitigate the effects of even the worst case failure. The results obtained in this paper will provide insight into the design of control algorithms and network architectures.

The rest of the paper is organized as follows. In Section II we present the related work and in Section III we describe the power flow and cascade models. Section IV provides analytical results regarding simple network topologies. Section V presents the method used to set some of the parameters and the algorithm used to identify the most vulnerable locations within the grid. Section VI describes the power grid data used in this paper. In Section VII we present our numerical findings regarding large scale failure and in Section VIII we discuss the recent cascade in the San Diego area. Section IX describes optimal control methods and their impact in the particular case of one of the simulated worst-case events that we compute. Section X provides concluding remarks and directions for future work.

II. RELATED WORK

The power grid and its robustness have drawn a lot of attention recently, as efforts are being made to create a smarter, more efficient, and more sustainable power grid (see, e.g., [2], [19]). The power grid is traditionally modeled as a complex system, made up of many components, whose interactions (e.g., power flows and control mechanisms) are not effectively computable (e.g., [3], [6] and references therein).

When investigating the robustness of the power grid, cascading failures are a major concern [2], [13], [16], [22]. This phenomenon, along with other vulnerabilities of the power grid, was studied from a few different viewpoints. First, several papers studied common topological properties of power grid networks and probabilistic failure propagation models, so that one can evaluate the behavior of a generic grid as a self-organized critical system using, for example, *percolation theory* (see [5], [12], [23], [25], [26], [30], [35], [39], [41] and references therein). These works are closely related to a long line of research in the power community which uses monte carlo techniques to analyze system reliability (e.g., [10]).

Another major line of research focused on specific (microscopic) power flow models and used them to identify a *few vulnerable lines* throughout the entire network [7], [8], [32]. In particular, [9], [34] focus on designing line or node interdiction strategies that will lead to an effective attack on the grid. Since the problem is computationally intractable, most of these papers use a linearized direct-current (DC) model, which is a tractable relaxation of the exact alternating-current (AC) power flow model. In addition, the initial failure events (causing eventually the cascading failures) are assumed to be sporadic link outages (and in most cases, a single link outage), with no correlation between them. On the contrary, we focus on events that cause *a large number of failures in a specific geographical region* (e.g., [18], [38]). To the best of our knowledge, geographically-correlated failures in the power grid have not been considered before and have been studied only recently in the context of communication networks [1], [14], [21], [29], [36].

Moreover, since cascading failures are highly dependent on the network topology, we use the real topology of the western U.S. power grid. While building on the linearized model, we manage to obtain numerical results for a very large scale real network (results for large networks have been derived in the past using mostly probabilistic models [12], [23], [30]). Perhaps, closest to us in its approach is [31] that analyzes the effects of single line failures on the Polish power grid. Yet, while [31] considers a real topology, it does not take into account geographical effects. Moreover, it applies control mechanisms that are more sophisticated than the ones considered here. Control mechanisms are also introduced in [4] that develops a method to trade off load shedding and cascade propagation risk. Recently, [27], [28] proposed efficient optimization algorithms for solving the classical power flow problems. These algorithms use efficient mathematical programming methods and can support *offline*

control decisions. However, they are not applicable when rapid *online* control of the grid is required.

III. BASIC MODELS

We adopt the *linearized* (or DC) *power flow model*, which is widely used as an approximation for more realistic non-linear AC power model (see [6] for a survey on the power flow models). In particular, we follow [7], [8] and represent the power grid by a directed graph \mathcal{G} , whose set of nodes is \mathcal{N} . Each of these nodes is classified either as a *supply node* (“generator”), a *demand node* (“load”), or a *neutral node*. Let $\mathcal{D} \subseteq \mathcal{N}$ be the set of the demand nodes, and for each node $i \in \mathcal{D}$, let D_i be its demand. Also, $\mathcal{C} \subseteq \mathcal{N}$ denotes the set of the supply nodes and for each node $i \in \mathcal{C}$, P_i is the active power generated at i . The edges of the graph \mathcal{G} represent the transmission lines. The orientation of the lines is arbitrarily and is simply used for notational convenience. We also assume *pure reactive* lines, implying that each line (i, j) is characterized by its *reactance* x_{ij} .

Given supply and demand vectors (P, D) , a *power flow* is a solution (f, θ) of the following system of equations:

$$\sum_{(i,j) \in \delta^+(i)} f_{ij} - \sum_{(j,i) \in \delta^-(i)} f_{ji} = \begin{cases} P_i, & i \in \mathcal{C} \\ -D_i, & i \in \mathcal{D} \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

$$\theta_i - \theta_j - x_{ij} f_{ij} = 0, \quad \forall (i, j) \quad (2)$$

where $\delta^+(i)$ ($\delta^-(i)$) is the set of lines oriented out of (into) node i , f_{ij} is the (real) power flow along line (i, j) , and θ_i is the phase angle of node i . These equations guarantee power flow conservation in each neutral node, and take into account the reactance and power capacity of each line. In addition, since the orientation of lines is arbitrary, a negative flow value simply means a flow in the opposite direction.

When \mathcal{G} is fully connected and $\sum_{i \in \mathcal{C}} P_i = \sum_{i \in \mathcal{D}} D_i$, (1)–(2) has a unique solution [8, Lemma 1.1]. This holds even when \mathcal{G} is not connected but the total supply and demand within each of the connected components are equal.

We note that the DC power flow model resembles an *electrical circuit model*, where phase angles are analogous to voltages, reactance is analogous to resistance, and the power flow is analogous to the current. The following observation, which is analogous to Kirchoff’s law, captures the essence of the model and provides easier way to look at (1)–(2) analytically. It uses the notion of a path between nodes, which is an alternating sequence of lines and adjacent nodes. (Proof is in the appendix).

Observation 3.1: Consider two nodes a and b and two paths $\pi_1 = (a, e_0, v_0, e_1, v_1, \dots, b)$ and $\pi_2 = (a, e'_0, v'_0, e'_1, v'_1, \dots, b)$. The sum of the flow-reactance product $f_{e_i} x_{e_i}$ along the lines of path π_1 is equal to the sum of the flow-reactance product along the lines of path π_2 . Specifically, the flows along parallel lines with the same reactance is the same.

Next we describe the *Cascading Failure Model*. We assume that each line (i, j) has a predetermined *power capacity* u_{ij} , which bounds its power flow in a normal operation of the

Cascading Failure Model (Deterministic Case)

Input: Connected network graph \mathcal{G} .

Initialization: Before time step $t = 0$, we have that $\sum_{i \in \mathcal{C}} P_i = \sum_{i \in \mathcal{D}} D_i$ (i.e., the power is balanced), (1)–(2) are satisfied for \mathcal{G} , and all flows along all lines are within the corresponding power capacity.

Failure event: At time step $t = 0$, a failure of some subset of links of \mathcal{G} occurs. Let $\mathcal{G}.\text{changed} = \text{true}$.

While $\mathcal{G}.\text{changed}$ is true do:

- 1) Adjust the total demand to the total supply *within each component* of \mathcal{G} .
- 2) Use the system (1)–(2) to recalculate the power flow in \mathcal{G} .
- 3) For all lines compute a moving average

$$\tilde{f}_{ij}^t = \alpha |f_{ij}| + (1 - \alpha) \tilde{f}_{ij}^{t-1}$$

- 4) Remove from \mathcal{G} all lines with flow *moving average* above power capacity ($\tilde{f}_{ij}^t > u_{ij}$). If at least one line was removed at this step, let $\mathcal{G}.\text{changed} = \text{true}$; otherwise, let $\mathcal{G}.\text{changed} = \text{false}$.

system (that is, $|f_{ij}| \leq u_{ij}$). We assume that before a failure event, \mathcal{G} is fully connected, the total supply and demand are equal, the power flows satisfy (1)–(2), and the power flow of each line is at most its power capacity. Upon a failure, some lines are removed from the graph, implying that it may become disconnected. Thus, within each component, we adjust the total demand to equal the total supply, by decreasing the demand (supply) by the same factor at all loads (generators). This process is sometimes called *demand shedding* and naturally it causes a decrease in demand/supply. Then, we use (1)–(2) to recalculate the power flows in the new graph. The new flows *may exceed the capacity* and as a result, the corresponding lines will become overheated. Thermal effects cause overloaded lines to become more sensitive to a large number of effects each of which could cause failure. We model outages using a *moving average* of the power flow, using a value $\tilde{f}_{ij}^t = \alpha |f_{ij}| + (1 - \alpha) \tilde{f}_{ij}^{t-1}$ (in this paper, we mostly use either $\alpha = 0.5$ or $\alpha = 1$). To first order, this approximates thermal effects, including heating and cooling from prior states. A similar moving average model was considered in [4], [8]. A general *outage rule* gives the fault probability of line (i, j) , given its moving average \tilde{f}_{ij}^t . In this paper, we consider the following rule:

$$\mathbb{P}\{\text{Line } (i, j) \text{ faults at round } t\} = \begin{cases} 1, & \tilde{f}_{ij}^t > (1 + \varepsilon)u_{ij} \\ 0, & \tilde{f}_{ij}^t \leq (1 - \varepsilon)u_{ij} \\ p, & \text{otherwise.} \end{cases} \quad (3)$$

where $0 \leq \varepsilon < 1$ and $0 \leq p \leq 1$ are parameters. When $\varepsilon = 0$, we obtain a *deterministic* version of this rule. In this case, lines (i, j) whose \tilde{f}_{ij}^t is above the power capacity u_{ij} are removed from the graph.

The process is repeated in rounds *until the system reaches stability*, namely until there is an iteration in which no lines are removed. We note that our model does not have a notion of exact time, however the relation between the elapsed time and the corresponding time can be adjusted by using different

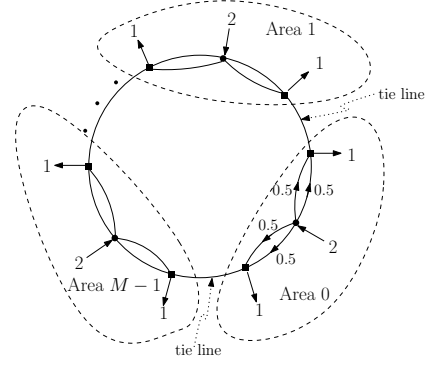


Fig. 2. The M -ring \mathcal{R}_M . Each generation node and its two adjacent demand node are a self-sustained area. No power flow is transmitted between these areas (that is, along the tie lines).

values of α ; in a sense, up to a certain degree, smaller value of α implies that we take a more microscopic look at the cascade.

Our major metric to assess the severity of a cascading failure is the system post-failure *yield* which is defined as follows:

$$Y \triangleq \frac{\text{The actual demand at the stability}}{\text{The original demand}}. \quad (4)$$

In addition to the yield, other performance metrics will be considered, such as the number of faulted lines, the number of connected components, and the maximum line overload. While the yield naturally gives an assessment of the severity of the cascade after the process has already finished, the other metrics may also shed light on the cascade properties after a fixed (given) number of rounds.

Note that our model contains a very simple control mechanism, namely, round-by-round demand shedding. In Section IX, we consider more elaborated control mechanisms and compare the results.

IV. CASCADING FAILURES PROPERTIES IN SIMPLE GRAPHS

In this section we describe important properties of the power flow and cascading failure models in a simple graph. Specifically, we show that unlike other flow models, *cascading failures in a power flow network are harder to predict and are different from epidemic-like failure models*, from the following four reasons:

- 1) Consecutive failures in a cascade may happen within an arbitrarily long distance of each other.
- 2) Cascading failures can last arbitrarily long time.
- 3) A network \mathcal{G}_1 whose topology is a subgraph of another network \mathcal{G}_2 can be more resilient to failures than \mathcal{G}_2 .
- 4) A failure event which results in initial failure of some set of lines A can cause more damage than a failure event, whose initial set of faulted lines is a superset of A .

In order to prove the results, we use the following simple power flow topology, which is depicted in Figure 2

Definition 1: An M -ring $\mathcal{R}_M = \langle \mathcal{C} \cup \mathcal{D}, \mathcal{E} \rangle$ is a directed graph with M supply nodes $\mathcal{C} = \{0, \dots, M-1\}$, $2M$ demand

nodes $\mathcal{D} = \{M, \dots, 3M - 1\}$, two parallel transmission lines connecting each generator $i \in \mathcal{C}$ to demand node $M + 2i$, two parallel transmission lines connecting generator $i \in \mathcal{C}$ to demand node $M + 2i + 1$, and a single transmission line connecting demand node $M + 2i + 1$ ($i \in \mathcal{C}$) to demand node $M + (2i + 2 \bmod 2M)$ ¹. For each $i \in \mathcal{C}$, the generation value is $P_i = 2$, and the demand values are $D_{2i} = D_{2i+1} = 1$. The reactance value x_{ij} of all transmission lines is 1.

Clearly, one can view the M -ring as a collection of M self-sustained areas: each generator i supplies the demands of $M + 2i$ and $M + 2i + 1$. For brevity, we call the lines connecting i and $M + 2i$ *even lines*, the lines connecting $M + 2i + 1$ *odd lines*, and the lines connecting two demands (that is, connecting two self-sustained areas) *tie lines*. Moreover, we refer to odd and even lines collectively as *internal lines*.

By symmetry, we have the same power flow on all even lines and odd lines. The solution of (1)–(2) verifies that this is indeed the case: a power flow of 0.5 is transmitted from each generator along its 4 adjacent internal lines, as shown in Figure 2, the phase angle of all generators is the same, and the phase angle of all demand nodes is the same. This also implies that there is no flow on the tie lines. Notice that if all lines has a power capacity of 0.5, it suffices to sustain that power flow.

In the rest of the section, we will consider the following types of failure events:

- **Area failure:** The four internal lines within Area i , as well as the two lines connecting Area i to the adjacent areas, fault. Namely, lines $(i, M + 2i)$, $(i, M + 2i)'$, $(i, M + 2i + 1)$, $(i, M + 2i + 1)'$, $(M + (2i - 1 \bmod 2M), M + 2i)$, $(M + 2i + 1, M + (2i + 2 \bmod 2M))$ fault.
- **Parallel lines failure:** Two parallel internal lines within Area i fault. Without loss of generality, we consider only even parallel lines, that is, $(i, M + 2i)$ and $(i, M + 2i)'$ fault.
- **Odd line and even line failure:** One odd line and one even line within Area i fault. Namely, lines $(i, M + 2i)$ and $(i, M + 2i + 1)$ fault.
- **Single internal line failure:** Single internal line in Area i faults. Without loss of generality, we consider only even line, that is, line $(i, M + 2i)$ faults.
- **Tie line and internal line failure:** Internal line in Area i , as well as the tie line connected to the corresponding demand, fault. Without loss of generality, we consider only an even line and its corresponding tie line, that is, $(i, M + 2i)$ and $(M + (2i - 1 \bmod 2M), M + 2i)$ fault.

It is important to notice that these failures cover the possible types of geographical failures over the ring (some combinations of these failures can also occur if the failure radius is large enough). In this section, we will not consider sporadic failure events (that is, in which the outaged lines are not geographically-correlated), or initial events that partition the

graph such that there is a component whose total demand is not equal to its total generation.

We next compare between an area failure and parallel lines failure. As for area failure, it is easy to see that the same power flow solution still holds, and therefore, the power flow along all operating lines does not change. The loss of demand is only 2 units and the resulting yield is $\frac{M-1}{M}$. On the other hand, upon parallel lines failure, the entire generation of node i must be transmitted on the only operating lines connected to node i : $(i, M + 2i + 1)$ and $(i, M + 2i + 1)'$. Since the lines are parallel, by symmetry, each of these lines transmits a power of 1 unit (cf. Observation 3.1). In addition, by (1), the tie line connecting area i to area $i + 1$ carries power of 1 unit since node $2i + 1$'s demand is only 1. Now focus on Area $i + 1$. The total amount of incoming generation to this area is 3 units (2 units generated in node $i + 1$ and one along the tie line in Area i), while the total demand is 2 units. This immediately implies that the tie line to Area $i + 2$ carries one unit of power. As for the flows on the lines within that area, one can verify that a valid (and therefore the only) solution of (1)–(2) has no power flow on lines $(i + 1, M + 2i + 2)$ and $(i + 1, M + 2i + 2)'$ and one unit of power flow along $(i + 1, M + 2i + 3)$ and $(i + 1, M + 2i + 3)'$. This flow assignment is identical to the assignment of Area i , and therefore, by inductive arguments, it is valid for all the areas in the ring. Moreover, there are three phase values in the solution: one for all generation nodes, one for all even demand nodes, and one for all odd demand nodes.

Based on the power capacity of the lines u_{ij} , the moving average parameter α of the cascading failure model, and the FoS of the entire system, we can derive the following results (all proofs are in the appendix):

Lemma 4.1: Consecutive failures in a cascade may happen within an arbitrarily long distance of each other.

We note that Lemma 4.1 captures an important difference between our model and previously-suggested models, that assume that power grid failures propagate in an epidemic-like manner. While, under these models, a line failure causes only adjacent node/line (or a line with small distance) to fault, our model captures situations in which the cascade “*skips*” large distances within a single iteration. As we will discuss in Section VIII, this was indeed the case in a recent real-life cascade causing a major blackout in California.

Lemma 4.2: A failure of $o(1)$ of the lines may cause an outage of a constant fraction of the lines, within one iteration.

The following two lemmas show that the failures do not always behave monotonically:

Lemma 4.3: An initial failure event of some set of lines A may result in a lower yield than a failure event, whose initial set of faulted lines is a superset of A .

Lemma 4.4: A network \mathcal{G}_1 , whose topology is a subgraph of the topology of another network \mathcal{G}_2 , may obtain higher yield.

We note that in practice, if the failures are geographically correlated, non-monotone situations, as described in Lemmas 4.3 and 4.4, rarely happen. Thus, in the rest of the paper (and specifically, in the algorithm described in Section V-B),

¹Parallel transmission lines are denoted (i, j) and $(i, j)'$. The orientation of the lines is only for notational convenience.

we assume only a monotone behavior of the failures.

Up until now, we dealt with very specific type of failures, which intuitively *breaks the ring into a chain*. In general, such failures are easier to analyze since (2) does not form constraints in a cyclic manner (that is, one can assign the phase value of the nodes based only on the flows and the phase value of the nodes that precedes it). More formally, this extra difficulty is captured by Observation 3.1, and looking at the different paths between each pair of nodes. When the ring breaks into a chain, all these paths follow one “direction”. On the other hand, when the ring does not break, there are both clockwise and counter-clockwise paths that need to be considered. We demonstrate this difference by comparing a single internal line failure of line $(0, M)$ (that does not break the ring into chain) with a tie line and internal line failure of lines $(0, M)$ and $(3M - 1, M)$ (which breaks the ring into chain).

In the tie line and internal line failure, one can simply assign a power flow of 1 unit along the parallel line $(0, M)'$, where the phase difference between node 0 and M is 1 (to meet the constraint of (2)); the rest of the flows and phases remain unchanged. However, this solution is not valid in case of a single internal line failure, since there is a phase difference between node M and $3M - 1$, and therefore, it is impossible that no flow traverses the operating line between them. We note also that in the odd and even lines failure (which does not break the chain either), the flows on the lines parallel to the failures increase to 1 by changing solely the generator phase. Since the demand nodes' phases do not change, this failure is localized within that single area.

We next show that the power flow values induced by the single line failure change across the entire graph and depends on the value of M :

Lemma 4.5: Consider an M -ring, in which line $(0, M)$ failed. Then, after the first iteration:

- (i) The flow along line $(0, M)'$ is $2M/(2M + 0.5)$.
- (ii) The flow along all other even lines is $M/(2M + 0.5)$.
- (iii) The flow along all odd lines is $1 - M/(2M + 0.5)$.
- (iv) The flow along all tie lines is $1 - 2M/(2M + 0.5)$.

Corollary 4.6: An M -ring requires power capacity of $1/9$ on its tie lines and a power capacity of 1 on its internal lines to withstand any failure of one line.

Interestingly, one can see that as M gets larger, a single internal line failure has the same effect as the corresponding tie line and internal line failure. This stems from the fact that the closed-loop effects, initially distinguishing between the failures, are fading away as the ring gets larger.

Finally, the next lemma shows that, unlike the examples we presented on the ring, cascades can be made arbitrarily long (in time). The lemma uses another topology which is depicted in Figure 15.

Lemma 4.7: The length of the cascade (the number of iterations until the system stabilizes) can be arbitrarily large.

To conclude, using simple examples we highlight the difficulties and differences between prior models used to analyze the power grid and the models we use. In the rest of the paper, we will investigate how these models behave in real-life power

grid and geographically-correlated failures.

V. POWER GRID RESILIENCE

A. Parameters Set-up

We note that in the cascading failure model, the power capacities u_{ij} of the lines are given a-priori. In practice, however, these capacities are hard to obtain and are usually estimated based on the actual operation of the power grid. In this paper, we take the $N-k$ contingency analysis approach [8] in order to estimate the power capacities. Namely, we set the capacities so that the network is resilient to failure of any set of k out of the N lines. In addition, we consider over-provisioning of lines capacity by a constant fraction of the required capacity of each line. This over-provisioning parameter, denoted by K , is often referred to as the *Factor of Safety* (FoS) of the grid.

Specifically, we focus on the following two cases.

- **N -resilient grids** (that is, $k = 0$). In this case, we solve (1)–(2) for the original grid graph (without failures) and set the power capacity to $u_{ij} = K \cdot f_{ij}$, where $K \geq 1$.
- **$(N - 1)$ -resilient grids** (that is, $k = 1$). In this case, we solve (1)–(2) for N graphs, each resulting from a single line failure event. The power capacity is set to $u_{ij} = K \cdot \max_r f_{ij}^r$, where f_{ij}^r is the flow assigned to line (i, j) when considering the r^{th} failure event.

It is worth mentioning that the real power grid is usually assumed to be at least N -resilient with $K \approx 1.2$ [15]. On the other hand, some data show that certain lines (or, more generally, *paths*) are more resilient than others. For example, a historical transmission paths data found in [40] shows that some transmission paths have power capacities which are 1.1 times their normal flow, while others have an FoS larger than 2. Nevertheless, the average FoS is indeed around 1.2. In addition, utility companies usually guarantee that their grid is at least $(N-1)$ -resilient [8]². Therefore, by setting the power capacities parameters, we examine in this paper both N - and $(N-1)$ -resilient grids with different FoS values K . Most of our numerical results are presented for a grid with FoS $K = 1.2$.

B. Identification of Vulnerable Locations

We consider a circular and deterministic *failure model*, where all lines and nodes within a radius r of the failure's epicenter are removed from the graph (this includes lines that pass through the affected area). In addition, we assume *monotonicity* of failures: if the initial set of faulted lines due to event A is a subset of the faulted lines due to event B , then the yield of A is at least that of B . We note that in the general case, this property does not hold (see Lemma 4.3). However, we observed that such events rarely happen in real power grid systems, and when they do, it is only when both events have a marginal effect. Since our goal is to identify the

²We note that early reports on the recent San Diego blackout indicate that this was not the case.

most vulnerable locations in a real power grid, this is a valid assumption for practical purposes.

To identify the candidates for the most vulnerable locations, we use computational geometric methods developed in [1] for identifying the vulnerable locations in fiber-optic networks. For each line, we define an *r*-hippodrome, which captures all points in the plane \mathbb{R}^2 whose distance from the line is at most the failure radius *r*. We focus on the *arrangement of hippodromes*, which is the subdivision of the plane into vertices, arcs, and faces. The vertices are the intersection points of the hippodromes, the arcs are either maximally connected circular arcs or straight line segments of the boundaries of hippodromes that occur between the vertices, and faces are maximally connected regions bounded by arcs.

Once the vertices of the arrangements are identified, we treat each vertex *v* as a *candidate* for a failure epicenter and denote by $\mathcal{L}(v)$ the set of lines within radius *r* of *v* ($\mathcal{L}(v)$ can be easily found). We then use the Cascading Failure Model, described in Section III, with $\mathcal{L}(v)$ as the set of lines that initially fault. Naturally, the process of checking all candidates (each with a different initial failure event) can be easily parallelized.

Arrangements are well-established concept in computational geometry, and it can be easily shown that in order to find the vulnerable locations, it is sufficient to consider only the vertices of the arrangements. In particular, for any point $p \in \mathbb{R}^2$, there is a vertex *v* such that $\mathcal{L}(p) \subseteq \mathcal{L}(v)$. Notice that computing arrangements is quadratic in the number of lines. Thus, we parallelized this computation as well by partitioning the graph into several sections (with small number of lines) and finding vertices of the arrangements in each section. To ensure that no vertices are lost in the border between two sections, the sections have a $2r$ overlap.

VI. POWER GRID DATA

We use *real power grid data* of the western US taken from the Platts Geographic Information System (GIS) [33]. This includes the information about the transmission lines, power substations, power plants, and the population at each geographic location. Since, in GIS each transmission line is defined as a link between two power substation, substations are used as nodes in our graph. In order not to expose the vulnerability of the real grid, we used a *part* of the Western Interconnect system which does not include the Canada and Mexico sections. On the other hand, we attached to the grid the Texas, Oklahoma, Kansas, Nebraska, and the Dakotas' grids, which are not part of the Western Interconnect. The resulting graph (see Figure 1) has 14,968 nodes (substations) and 19,513 lines. Moreover, it has 1,920 power stations which were merged with the substations as described below. We note that there is a small number of very dense areas (e.g., the Los Angeles area), while the rest of the grid is very sparse. This structure can be seen in many other typical power grids, such as the US Eastern Interconnect as well as European systems. Furthermore, recent research on topological models for power grid systems show similar results [23]. Thus, our results will probably carry over to other grids.

We performed the following processing steps for this graph.

Coordinate transformation: The coordinates of the substations in the GIS system are given by their *longitude and latitude*. We transformed these to *planar* (*x, y*) coordinates, using the *great-circle distance* method.

Connectivity check: We identified the connected components of the raw GIS graph, which consists of one large connected component and few small ones. Moreover, we identified some substations that appear in the transmission lines data but are absent in the substations data. The number of these substations is small, and therefore, after manual inspection, they were either eliminated or merged with other nearby substations (see also the next step).

Nodes merging and lines elimination: For each node outside the large component, we found the closest node within the large component. If the distance between them was below a given threshold (10 Km), we merged the two nodes. Then, the remaining disconnected nodes were inspected manually and were either eliminated or merged. At the end of this process, we obtained a fully-connected graph (that is, a single connected component) with 13,992 nodes and 18,681 lines.

Identifying demands and supplies: Demands were associated with the closest node to each populated area (i.e., zip-code) and were set to be proportional to the population size (the normalization factor is computed by comparing the total population and total generation output of the entire grid). Supplies were associated with closest node to each power plant (the generation capacity of the power plants is taken from the GIS). Then, for each node, we computed the difference between its total corresponding supplies and total corresponding demands. Thus, all nodes were characterized by a real number: positive for a supply node, negative for a demand node, and zero for a neutral node. The resulting categorization appears in Figure 1. Overall, 1,117 nodes were classified as generators (supplies), 5,591 as loads (demands), and 7,284 as neutral. Most of the neutral nodes are closely connected to each other and to one of the non-neutral nodes, thus drawing the power/demand from them.

Determining the system parameters: The GIS does not provide the power capacities of the transmission lines, nor their reactance. However, these parameters are needed for the power flow and cascading failure models. The reactance of a line depends on its physical properties (such as its material) and there is a linear relation between its length and reactance: the longer the line is, the larger its reactance. Thus, we assumed that all lines have the same physical properties (other than length) and used the length to determine the reactance. It is important to notice that the flow part of the solution of (1)–(2) is scale invariant to the reactance (that is, multiplying the reactance of all lines by the same factor does not change the values of the flows). Thus, we simply use the length of the line as its reactance. Regarding the power capacities, we take the approach described in Section V-A. In particular, we use $N - k$ contingency analysis, with $k \in \{0, 1\}$ and different FoS values *K*.

VII. NUMERICAL RESULTS

We identified the potential failure locations using the algorithm described in Section V-B implemented in MATLAB. We present results for $r = 50$ kilometers, which is small enough to capture realistic scenarios [18], [38], while it is large enough to generate a cascading failure in most cases (the results for other values of r are omitted, due to space constraints). For $r = 50$, the algorithm identified 61,327 potential failure locations. The identification of these locations was done on an eight-core server in less than 24 hours.

For each failure location v , we performed the simulation of the Cascading Failure Model, presented in Section III, assuming that all lines in $\mathcal{L}(v)$ fail. The simulation was performed using a program that efficiently solves very large systems of linear equations, using CPLEX [24] and Gurobi [20] optimization tools.

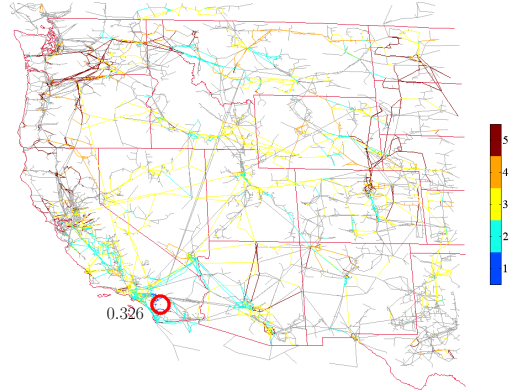
To assess the severity of a cascading failure, we use the following four metrics, which are measured in the end of the cascade: **The yield**, as defined in (4); **the total number of outaged lines**, which indicates the time it takes to recover the grid after the cascade: the larger the number of outaged lines, the longer is the actual time of the corresponding blackout; **the number of connected components**; and **the number of rounds until stability**.

A. N-Resilience Experiments

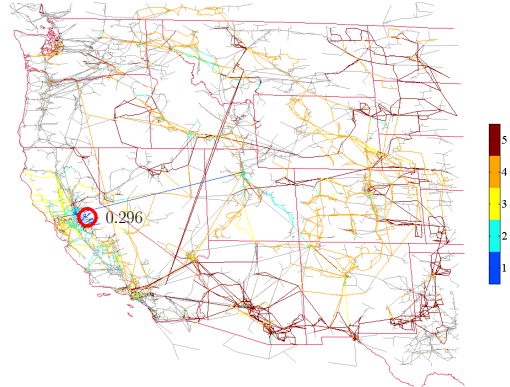
The first set of experiments was performed after using the N contingency analysis to set the capacities of the network, with FoS $K = 1.2$ and $K = 2$.

First, we plot specific failures to show how they evolved during the first *five* rounds of the cascade. Figure 3 shows three failure events for FoS $K = 1.2$: Two in California, both leading to severe blackouts, and another one around the Idaho-Montana-Wyoming border, which had a less severe effect. The round-by-round maximum overload (that is, $\max_{ij} f_{ij}/u_{ij}$) for these failures and $K = 1.2$ is shown in Table I. The same failures for $K = 2$ are shown in Figure 17 in the appendix. As expected, higher FoS usually leads to less severe blackout effect. Interestingly, the Idaho-Montana-Wyoming border failure with FoS $K = 1.2$ leads to low yield (0.39), although the development of the failure is very slow—after 5 rounds only few lines were faulted. However, the same event with $K = 2$ leads to near-unity yield. We note that this suggests that the assumption that $K = 1.2$ for all lines is quite pessimistic, as also can be seen from the actual data (see Section V-A for more details).

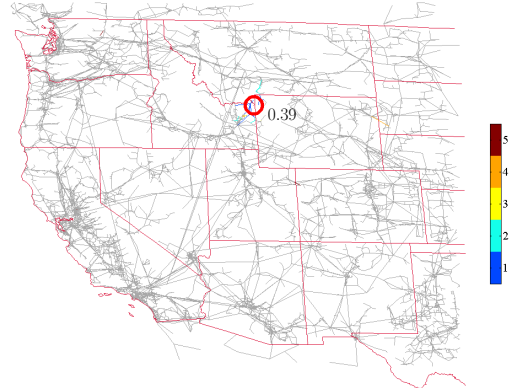
Scatter graphs for different metrics after 5 rounds and with FoS $K = 1.2$ are shown in Figure 4. It can be seen an increase in the initial number of faulted lines leads to an increase in the total number of faulted lines at the end of the fifth round: if 400, 800, 1,200 lines initially faulted, at least 2,847, 3,600, 4,669 are faulted at the end, respectively. Furthermore, an increase in the initial number of faulted lines leads also to an increase in the number of connected components: if 400, 800, 1,200 lines initially faulted, the number of components is at least 696, 1,382, 1,973, respectively. Similar results for



(a) San Diego area failure event.



(b) San Francisco area failure event.



(c) Idaho-Montana-Wyoming border failure event.

Fig. 3. Illustration of cascading failures over 5 rounds for N -resilient grid with FoS $K = 1.2$, where the initial failure locations are in the (a) Los Angeles area, (b) San Francisco area, and (c) Idaho-Montana-Wyoming border. The final yields are 0.326, 0.296, and 0.39, respectively. The colors represent the rounds in which the lines faulted.

$K = 2$ are shown in Figures 17 and 16 in the appendix. These results clearly show that in this case the grid is much more resilient to cascading failure.

Next, we analyze the severity of cascading failures once *stability* is reached. Namely, when no more line failures occur. The results for FoS $K = 1.2$ are shown in Figure 5. In this case, the vast majority of the failures resulted in yield in the

Round	San Diego Maximum Overload	San Francisco Maximum Overload	ID-MT-WY Maximum Overload
1	774.821	902.014	6.665
2	1724.525	24.923	1.193
3	16.527	57.87	1.836
4	14.99	2124	1.229

TABLE I
ROUND-BY-ROUND MAXIMUM LINE OVERLOAD $\max_{ij} f_{ij}/u_{ij}$ FOR THE
CASCADING FAILURES SHOWN IN FIGURE 3.

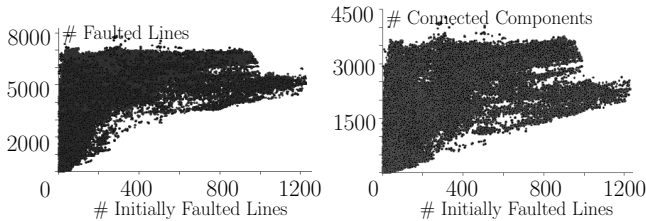


Fig. 4. The effects of the number of initially faulted lines on the total number of faulted lines (left) and the number of components (right), after 5 rounds of cascade (FoS $K = 1.2$).

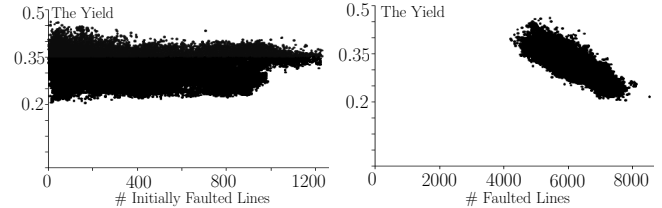
range of 0.2–0.46.

Figure 5(a) focuses on points whose yield is in $[0.2, 0.46]$. The variance of the yield is larger when the initial number of faulted lines is smaller. Moreover, as expected, there is an inverse correlation between the yield and the total number of faulted lines. Figure 5(b) shows the relation between the number of initially faulted lines and the number of rounds until stability, and in turn, the relation between the latter and the yield. We can see a *threshold effect* in the yield: when the number of rounds is small, the yield is around 1, while when more than 10 rounds are required, the yield drops to 0.5 and below. The correlation between the number of rounds and the number of initially faulted lines suggests (somewhat surprisingly) that usually a smaller number of initially faulted lines leads to a larger number of rounds until stability.

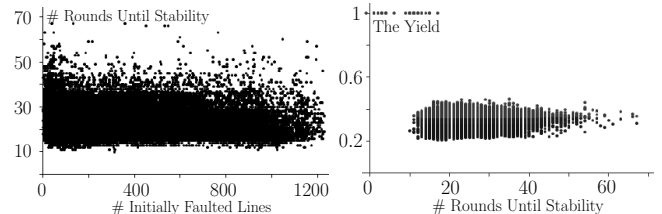
Finally, Figure 6 illustrates the yield, the number of rounds until stability, and the number of faulted lines at stability (respectively) by failure location.

B. $(N-1)$ -Resilience Experiments

The second set of experiments was performed after using the $N - 1$ contingency analysis to set the capacities of the network, with FoS $K = 1.2$. The results are presented in Figures 7 and 8. Regarding the failure events indicated in Figure 3, the corresponding yield values at stability in this case are 0.352, 0.333, and 0.999, respectively. The comparison between the results of N - and $(N-1)$ -resilience with the same FoS ($K = 1.2$) suggests, as expected, that $(N-1)$ -resilience helps when the initial event is not significant (such as the Idaho-Montana-Wyoming border event). However, it makes little difference when the initial event is significant (such as San Diego or San Francisco events). In particular, note that the failures in the artificially attached part of Texas do not lead to cascades when the network is $(N-1)$ -resilient. This is due to the fact that this part is connected to the whole network



(a) The yield as a function of the number of initially faulted lines (left) and the yield as a function of total number of faulted lines (right).



(b) The number of rounds as a function of the number of initially faulted lines (left) and the yield as a function of the number of rounds (right).

Fig. 5. Comparison between different performance metric at stability (FoS $K = 1.2$).

using small number of lines (which carry no power in normal operation in practice). However, when the network is only N resilient, these failures do propagate to the whole network³.

C. Stochastic Outage Rule

The third set of experiments was performed using a *stochastic* outage rule as defined in (3) with $\varepsilon > 0$ and $p = 0.5$.

In order to evaluate this outage rule, we performed two types of experiments on an N -resilient grid with FoS $K = 1.2$. First, for the *same* failure epicenter, we compared the yield of different values of ε : Figure 9 shows the average yield and its standard deviation for a representative failure epicenter (the results are based on 100 independent runs for each value of ε). Observe that $\varepsilon \in (0, 0.15)$ leads to a bit higher average yield than that of the deterministic rule. However, for $\varepsilon \geq 0.15$, the average yield obtained when using the stochastic rule is significantly lower.

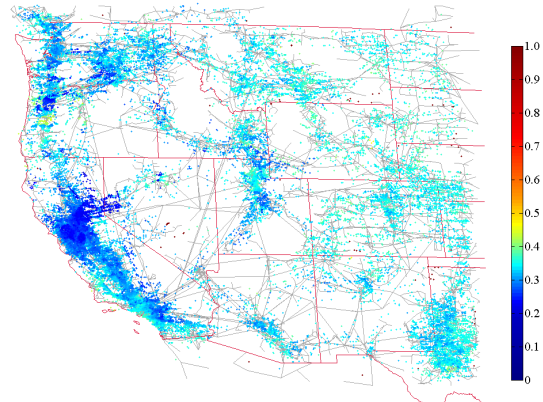
In the second type of experiments, we fixed $\varepsilon = 0.04$ and compared the results of *selected* failure epicenters with the results obtained for the deterministic outage rule. The failure epicenters were chosen such that the yield using deterministic rule grows approximately linearly with the failure index. The results, depicted in Figure 10, show that there is a certain yield range where the stochastic outage rule coincides with the deterministic outage rule. However, outside this range, the stochastic outage rule results in the yield values below 0.3, which are smaller than the yield obtained by a deterministic outage rule (even when this deterministic yield is almost 1).

VIII. SAN DIEGO BLACKOUT (SEPT. 2011)

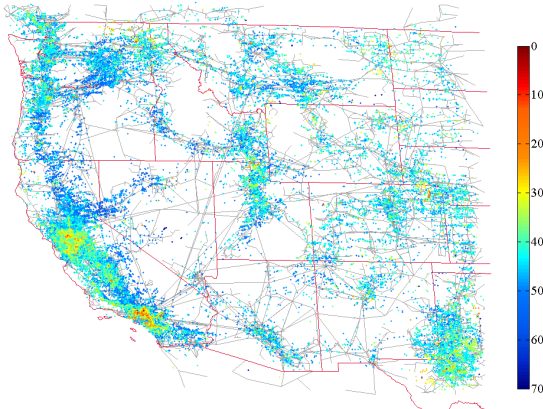
A. Description of the Blackout

On Sept. 8th, 2011, over 2.7 million people in southwestern United States experienced a massive power blackout. Al-

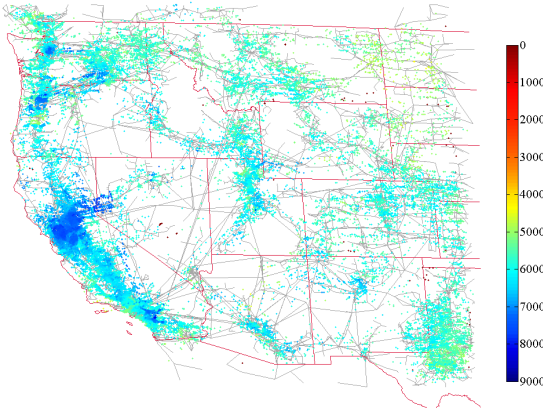
³This happens also even when the FoS is 2 (these results are not shown due to space constraints).



(a) The yield values at stability.



(b) The number of rounds until stability.



(c) The number of faulted lines at stability.

Fig. 6. Vulnerability analysis (at stability) of failure locations for N -resilient grid with FoS $K = 1.2$. The color of each point (which is a vertex of the arrangement) represents the value corresponding to a cascade whose epicenter is at that point (points that do not appear on the map cause outages that are a subset of the outages caused by a nearby vertex).

though the full details of this event are not known yet, several publicly available sources, such as [11], make it possible to reconstruct an approximate chain of events during this blackout. As a case study, we compare the reported chain of events to our simulation results. In particular, we use this

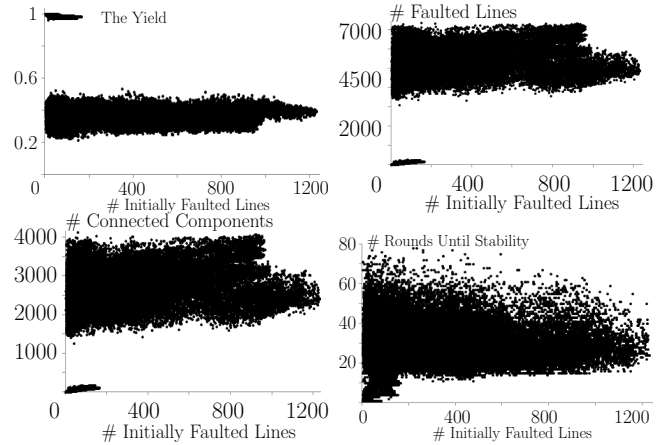


Fig. 7. Comparison between different performance metrics at stability; $(N-1)$ -resilient grid with FoS $K = 1.2$.

event in order to calibrate the various model parameters so that the simulation results match as closely as possible that cascade.

The blackout occurred around the San-Diego county area, and involved six utility companies: San-Diego Gas and Electric Co. (SDG&E), Southern California Edison (SCE), Comision Federal de Electricidad (CFE), Imperial Irrigation District (IID), Arizona Public Service (APS), and Western Area Power Association (WAPA). The power grid map of that area (using our data) is shown in Figure 11⁴. Specifically, there are two import generation paths into this area:

- 1) *SWPL*, which is represented by the 500KV Hassayampa-North Gila-Imperial Valley-Miguel transmission line. This path transmits the power generated in Palo Verde Nuclear Generating Station in Arizona.
- 2) *Path 44*, which is represented by the three 500KV transmission lines that connect SCE and SDG&E through the San Onofre Nuclear Generating Station (SONGS).

In addition, there are several SDG&E local power plants, and there is (relatively small) import of power from CFE.

Prior to the event, *SWPL* delivered 1370MW, *Path 44* delivered 1287MW, and the local generation was 2229MW [11] (this includes the generation of both SDG&E and CFE power plants). The cascade started at 15:27:39, when the 500KV Hassayampa-North Gila transmission line tripped at the North Gila substation. Several sources indicate that this failure was caused by maintenance works performed at this substation at that time. Initial investigation suggested that this single line failure caused the blackout⁵. The actual cascade development is shown in Figure 11.

⁴In these experiments, the *actual* Western Interconnect map was used, which includes the relevant parts of Northern Baja California in Mexico.

⁵Recently, some of the media publications mentioned the fact that this was not the only fault in that area. However, since these facts are still under investigation, we prefer to reconstruct the cascading failure according to the chain of event described in [11].

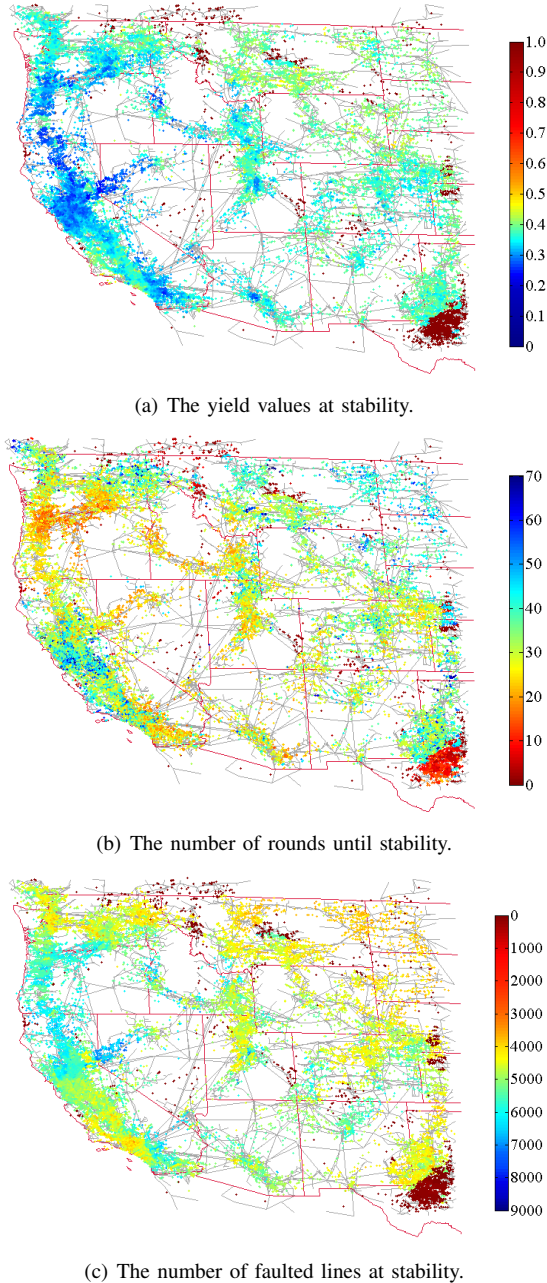


Fig. 8. Vulnerability analysis (at stability) of failure locations for an $(N-1)$ -resilient grid with FoS $K = 1.2$. The color of each point (which is a vertex of the arrangement) represents the value corresponding to a cascade whose epicenter is at that point.

B. Simulation Results

We performed two sets of experiments. In the first set, instead of performing simulation on the entire Western Interconnect, we chose to use a part of the grid which includes only the affected area. The initial conditions were set to match as close as possible the actual conditions prior to the event. In particular, we set the generation of the Palo Verde nuclear plant (which is the main contributing import generation unit) to 3,600 MW out of its nominal 4,300MW. This resulted

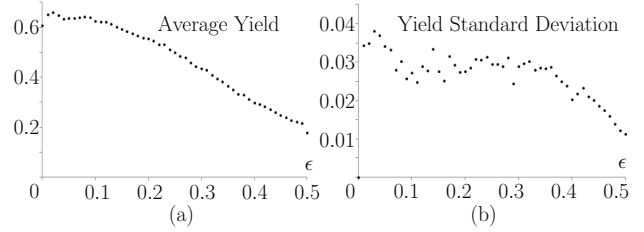


Fig. 9. The results for a representative failure epicenter, using stochastic outage rule, based on 100 independent runs. (a) presents the average yield, while (b) presents its standard deviation.

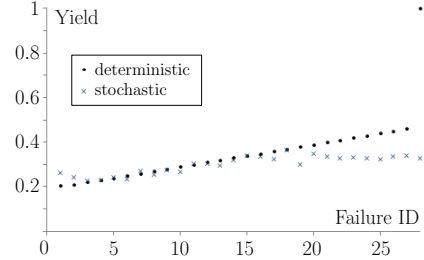


Fig. 10. A comparison of the deterministic and stochastic outage rules for selected failure events.

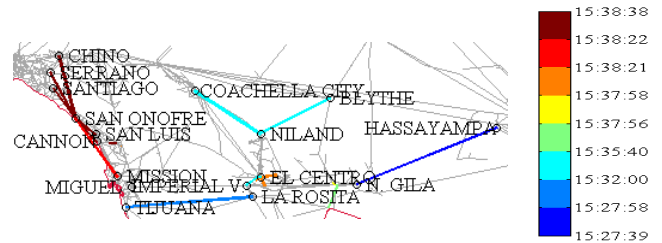


Fig. 11. The development of the San Diego blackout according to [11].

in the following initial conditions of the import generation: SWPL = 1,386MW, Path 44 = 1,284MW.

Moreover, since in the actual event there was no $(N-1)$ -resilience with respect to the faulted North Gila–Hassayampa line, we used an N -resilient grid with different values of FoS K (recall Section V-A). In addition, by [11], the actual capacity of Path 44 is almost 2.7 times the flow in normal operation. This information also correlates with other sources (e.g., [40]) which indicate that the power capacities are not based on a uniform FoS parameter. Since Path 44 was a major factor of the cascade development, as it carried most of the lost SWPL power, we decided to adjust its FoS accordingly. In particular, its FoS was set to 2.5. After experimenting with the value of K for other lines, we found out that $K = 1.5$ leads to a behavior that most resembles that of the actual event. The resulting cascade behavior is shown in Figure 12.

Table II presents a brief comparison of the simulation results and the known details of the actual event. The description of the actual event is presented exactly as in [11], without any interpretation. It can be observed that although the simulated cascade does not follow exactly the actual one, both of them developed in a similar way. This suggests that our model and data can be used to identify the vulnerable locations and design

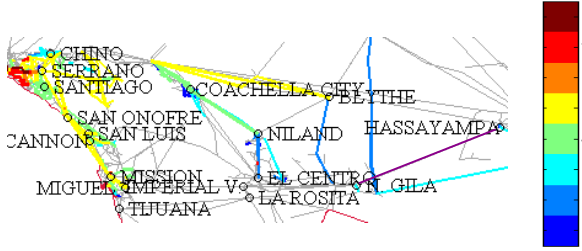


Fig. 12. The development of the San Diego blackout in the first eight rounds using our simulation.

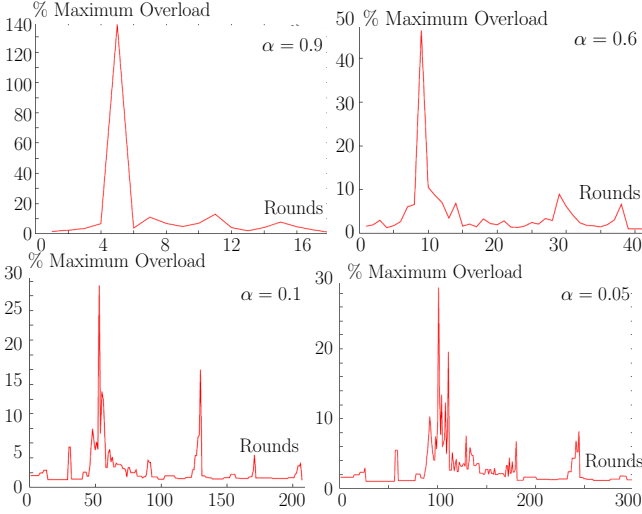


Fig. 13. Maximum line overload for different values of α .

corresponding control mechanisms that will allow to stop the cascades in the early stages.

The second set of experiments was performed on the whole Western Interconnect, and the goal was to examine the effect of the moving average parameter α on both the maximum line overloads and the length of the cascade. The results (see Figure 13) show that the larger α is, the higher is the maximum load and the shorter is the cascade. Moreover, when α is small (i.e., less than 0.5), there is a period of time when the maximum overload is smaller than that at the initial round. This suggests that a control mechanism that is applied at that time will stop the cascade with relatively high yield in early rounds.

IX. CONTROL

This section describes experiments with control algorithms in the specific case of the San Diego event, illustrated in Figure 3(a). The general goal of such algorithms is to stop the cascade (and, if possible, in a short time frame) without losing much demand.

In particular, we consider an algorithm that will shed a minimum amount of demand so as to yield a stable grid – the cascade has been stopped. In this paper (for lack of space), we focus on algorithms that will operate within a single round of the cascade. The critical question is, then, *at which round* control should be applied. Assuming that a given round t is

Actual		Simulation	
		1	Path 44: 1415MW; El Centro substation internal line overload of 100MW.
15:27	Path 44: 2407MW; Problems with Imperial Valley-El Centro line resulting in 100MW swing.	2	Path 44: 1438MW; El Centro internal line trip.
15:32	Path 44: 2616MW; Two lines trip at Niland-WAPA and Niland-Coachella Valley.	3	Path 44: 1992MW;
15:35	Path 44: 2959MW; IID and WAPA are separated.	4	Path 44: 3043MW; Niland-Coachella Valley line overload.
15:37	Path 44: 3006MW; IID tie line to WAPA trips.	5	Path 44: 2991MW; Niland-WAPA and Niland-Coachella Valley lines trip.
15:38	Path 44 trip; SONGS trips.	6	Path 44 trip; 4 out of 7 lines from SONGS to San Diego trip.
		7	SONGS stabilizes with total generation of 1350MW out of 2253MW.

TABLE II
COMPARISON OF THE ACTUAL EVENT AND THE SIMULATION RESULTS.

under consideration, our control is constrained as follows:

- At each demand point $i \in \mathcal{D}$, we reduce the demand by a certain quantity, s_i .
- We adjust generator output, within each component (a.k.a. island) so as to maintain overall balance.
- However, generators are furthermore constrained in that the amount of change in a generator must be proportional to its current output.
- After the demand shedding and generator adjustments, the moving-average power flow, on each operating line (i, j) , cannot exceed its capacity u_{ij} .

Rule (c) approximates generator “ramp-up” and “ramp-down” constraints (broadly speaking, generators cannot modify their output arbitrarily fast). Rule (d) states that, according to our thermal model, the cascade will stop. Rules (a)-(d) describe our constraints; the goal is to pick the round t and the quantities s_i so as to maximize the remaining demand.

Note that, at a given round t , this optimization problem can be written as a linear program. Specifically, denote by $\tilde{f}_{ij}^t, \tilde{D}_i^t, \tilde{P}_i^t$ the value, just before round t , of the flow on line (i, j) , the demand at demand point $i \in \mathcal{D}$, and the generation at supply point $i \in \mathcal{C}$, respectively. Furthermore, denote by C_1, \dots, C_n the connected components in the power grid graph before round t , and let $\text{comp}(i)$ denote the connected component that contain node i . The linear program is as follow:



Fig. 14. Illustration of cascading failures over 5 rounds using stochastic outage rule with $\varepsilon = 0.05$, $p = 0.5$, and $\alpha = 0.1$. The colors represent the rounds in which the lines faulted.

minimize $\sum_{i \in \mathcal{D}} s_i$ subject to

$$\begin{aligned}
 0 \leq s_i \leq \tilde{D}_i^t & \quad \forall i \in \mathcal{D} \\
 \alpha |f_{ij}| + (1 - \alpha) \tilde{f}_{ij}^t \leq (1 - \varepsilon) u_{ij} & \quad \forall \text{ line } (i, j) \\
 \sum_{(i,j) \in \delta^+(i)} f_{ij} - \sum_{(j,i) \in \delta^-(i)} f_{ji} = P_i, & \quad \forall i \in \mathcal{C} \\
 \sum_{(i,j) \in \delta^+(i)} f_{ij} - \sum_{(j,i) \in \delta^-(i)} f_{ji} = -(\tilde{D}_i^t - s_i) & \quad \forall i \in \mathcal{D} \\
 \sum_{(i,j) \in \delta^+(i)} f_{ij} - \sum_{(j,i) \in \delta^-(i)} f_{ji} = 0, & \quad \forall i \in \mathcal{N} \setminus (\mathcal{C} \cup \mathcal{D}) \\
 \theta_i - \theta_j - x_{ij} f_{ij} = 0 & \quad \forall \text{ line } (i, j) \\
 0 \leq \lambda^{C_m} \leq 1 & \quad \forall \text{ component } C_m \\
 P_i = \tilde{P}_i^t (1 - \lambda^{\text{comp}(i)}) & \quad \forall i \in \mathcal{C} \\
 \sum_{i \in C_m \cap \mathcal{C}} P_i = \sum_{i \in C_m \cap \mathcal{D}} D_i & \quad \forall \text{ component } C_m
 \end{aligned}$$

where, as in Section III, $\delta^+(i)$ ($\delta^-(i)$) is the set of lines oriented out of (into) node i . Notice that third-sixth equations in the linear program above are identical to Equations (1)-(2) in Section III.

As mentioned before, we demonstrate our control mechanism by considering a failure event in San Diego area. Figure 14 presents the development of this event over the first five rounds of a particular run obtained using the stochastic line failure model with $\varepsilon = 0.05$, $p = 0.5$, and $\alpha = 0.1$ (that is, a simulation with small time increments).

Table III outlines the performance of the optimal control mechanism, where ‘‘Round’’ refers to the round on which the optimal control is applied, while ‘‘Yield’’ refers to the outcome. We see from the table that applying control at the outset of the cascade is *not* optimal (this is typical, in our experience). On the other hand, waiting too long is not optimal either. Rather, there is a critical frame of time where effective control is possible; the precise time frame can be discovered by running our simulation upon the failure event, and applying the control only when we reach the round with optimal outcome. We also note that without control, the cascade stops at round 74 with the yield of 0.34. Currently, we are developing robust versions of this algorithm with respect to errors in data, timing, and delays in implementation.

TABLE III
OPTIMAL CONTROL OUTCOME.

Round	1	5	10	20	30	40	50	74
Yield	0.22	0.55	0.49	0.41	0.39	0.38	0.36	0.34

X. CONCLUSION AND FUTURE WORK

In this paper, we considered a DC power flow and an accompanied cascading failure model. We showed analytically that these models differ from previously-studied model based on an epidemic-like failures (which are often analyzed using percolation theory). Then, we used techniques from optimization and computational geometry along with detailed GIS data to develop a method for identifying the power grid locations that are vulnerable to geographically correlated failures. We performed extensive numerical experiments that show the relations between the various parameters and performance metrics. Specifically, we used a recent major blackout event in San Diego area as a case study to calibrate different parameters of the simulations. We also demonstrated that the use of control at the right point in the cascade can mitigate the effects of a large scale failure. While the presented results are for an intentionally modified version of the US Western Interconnect, they demonstrate the strength of our tools and provide insights into the issues affecting the resilience of the grid. These results can be used when designing new power grids, when making decisions regarding shielding or strengthening existing grids, and when determining the locations for deploying metering equipment.

This is one of the first steps towards an understanding of the grid resilience to large scale failures. Hence, there are still many open problems. For example, we plan to study the results’ sensitivity to the failure model (e.g., to consider a rule under which the probability of a line failure is a function of the overload). Moreover, we plan to study the effectiveness of some of the current control algorithms and their capability to cope with geographically correlated failures. Finally, we plan to develop control algorithms that will mitigate the effects of such failures and network design tools that would enable to construct resilient grids.

ACKNOWLEDGEMENTS

This work was supported in part by DOE award DE-SC000267, the Legacy Heritage Fund program of the Israel Science Foundation (Grant No. 1816/10), DTRA grant HDTRA1-09-1-0057, a grant from the from the U.S.-Israel Binational Science Foundation, and NSF grants CNS-10-18379 and CNS-10-54856. We thank Eric Glass for help with GIS data processing.

REFERENCES

- [1] P. Agarwal, A. Efrat, A. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, ‘‘The resilience of WDM networks to probabilistic geographical failures,’’ in *Proc. IEEE INFOCOM’11*, Apr. 2011.
- [2] M. Amin and P. F. Schewe, ‘‘Preventing blackouts: Building a smarter power grid,’’ *Scientific American*, May 2007.

- [3] G. Andersson, "Modelling and analysis of electric power systems," Lecture 227-0526-00, Power Systems Laboratory, ETH Zürich, March 2004, <http://www.eeh.ee.ethz.ch/downloads/academics/courses/227-0526-00.pdf>.
- [4] M. Anghel, K. A. Werley, and A. E. Motter, "Stochastic model for power grid dynamics," in *Proc. HICSS'07*, Jan. 2007.
- [5] K. Atkins, J. Chen, V. S. A. Kumar, and A. Marathe, "The structure of electrical networks: a graph theory-based analysis," *Int. J. Critical Infrastructures*, vol. 5, no. 3, pp. 265–284, 2009.
- [6] A. R. Bergen and V. Vittal, *Power Systems Analysis*. Prentice-Hall, 1999.
- [7] D. Bienstock, "Optimal control of cascading power grid failures," in *PES General Meeting*, July 2011.
- [8] D. Bienstock and A. Verma, "The $N - k$ problem in power grids: New models, formulations, and numerical experiments," *SIAM J. Optim.*, vol. 20, no. 5, pp. 2352–2380, 2010.
- [9] V. M. Bier, E. R. Gratz, N. Haphuriwat, W. Magua, and K. Wierzbicki, "Methodology for identifying near-optimal interdiction strategies for a power transmission system," *Reliab. Eng. Syst. Safety*, vol. 92, pp. 1155–1161, 2007.
- [10] R. Billinton and W. Li, *Reliability Assessment of Electrical Power Systems Using Monte Carlo Methods*. Plenum Press, 1994.
- [11] California Public Utilities Commission (CPUC), "CPUC briefing on San Diego blackout," http://media.signonsandiego.com/news/documents/2011/09/23/CPUC_briefing_on_San_Diego_blackout.pdf.
- [12] D. P. Chassin and C. Posse, "Evaluating north american electric grid reliability using the Barabási–Albert network model," *Physica A*, vol. 355, no. 2–4, pp. 667 – 677, 2005.
- [13] J. Chen, J. S. Thorp, and I. Dobson, "Cascading dynamics and mitigation assessment in power system disturbances via a hidden failure model," *Int. J. Elec. Power and Ener. Sys.*, vol. 27, no. 4, pp. 318 – 326, 2005.
- [14] T. N. Dinh, Y. Xuan, M. T. Thai, E. K. Park, and T. Znati, "On approximation of new optimization methods for assessing network vulnerability," in *Proc. IEEE INFOCOM'10*, 2010.
- [15] I. Dobson, "personal communication," 2010.
- [16] I. Dobson, B. Carreras, V. Lynch, and D. Newman, "Complex systems analysis of series of blackouts: cascading failure, critical points, and self-organization," *Chaos*, vol. 17, no. 2, p. 026103, June 2007.
- [17] W. R. Forstchen, *One Second After*. Forge Books, 2009.
- [18] J. S. Foster, E. Gjeldre, W. R. Graham, R. J. Hermann, H. M. Kluepfel, R. L. Lawson, G. K. Soper, L. L. Wood, and J. B. Woodard, "Report of the commission to assess the threat to the United States from electromagnetic pulse (EMP) attack, critical national infrastructures," Apr. 2008.
- [19] H. Gharavi and R. Ghafurian, "Special issue on smart grid: The electric energy system of the future," *Proc. IEEE.*, vol. 99, no. 6, 2011.
- [20] Gurobi, "Optimizer," <http://www.gurobi.com/>.
- [21] A. F. Hansen, A. Kvalbein, T. Cacic, and S. Gjessing, "Resilient routing layers for network disaster planning," in *Proc. ICN'05*, ser. LNCS, P. Lorenz and P. Dini, Eds. Springer, 2005, vol. 3421, pp. 1097–1105.
- [22] P. Hines, K. Balasubramaniam, and E. Cotilla-Sanchez, "Cascading failures in power grids," *IEEE Potentials*, vol. 28, no. 5, pp. 24–30, 2009.
- [23] P. Hines, E. Cotilla-Sanchez, and S. Blumsack, "Do topological models provide good information about electricity infrastructure vulnerability?" *Chaos*, vol. 20, no. 3, p. 033122, Sept. 2010.
- [24] IBM, "ILOG CPLEX Optimizer," <http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>.
- [25] Z. Kong and E. M. Yeh, "Resilience to degree-dependent and cascading node failures in random geometric networks," *IEEE Trans. Inf. Theory*, vol. 56, no. 11, pp. 5533–5546, Nov. 2010.
- [26] Y.-C. Lai, A. Motter, and T. Nishikawa, "Attacks and cascades in complex networks," in *Complex Networks*, ser. Lecture Notes in Phys., E. Ben-Naim, H. Frauenfelder, and Z. Toroczkai, Eds. Springer, 2004, vol. 650, pp. 299–310.
- [27] J. Lavaei and S. Low, "Zero duality gap in optimal power flow problem," to appear *IEEE Trans. Power Syst.*, 2011.
- [28] J. Lavaei, A. Rantzer, and S. H. Low, "Power flow optimization using positive quadratic programming," in *Proc. 18th IFAC World Congress*, Aug. 2011.
- [29] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the vulnerability of the fiber infrastructure to disasters," in *Proc. IEEE INFOCOM'09*, Apr. 2009.
- [30] G. A. Pagani and M. Aiello, "The power grid as a complex network: a survey," *ArXiv e-prints*, May 2011.
- [31] R. Pfitzner, K. Turitsyn, and M. Chertkov, "Controlled tripping of overheated lines mitigates power outages," *ArXiv e-prints*, Oct. 2011.
- [32] A. Pinar, J. Meza, V. Donde, and B. Lesieutre, "Optimization strategies for the vulnerability analysis of the electric power grid," *SIAM J. Optim.*, vol. 20, no. 4, pp. 1786–1810, Feb. 2010.
- [33] Platts, "GIS Data," <http://www.platts.com/Products/gisdata>.
- [34] J. Salmeron, K. Wood, and R. Baldick, "Worst-case interdiction analysis of large-scale electric power grids," *IEEE Trans. Power Syst.*, vol. 24, no. 1, pp. 96 –104, feb. 2009.
- [35] C. M. Schneider, A. A. Moreirab, J. J. S. Andrade, S. Havlin, and H. J. Herrmann, "Mitigation of malicious attacks on networks," in *Proceedings of the National Academy of Sciences of the United States of America*, vol. 108, no. 10, Mar. 2011, pp. 3838 – 3841.
- [36] A. Sen, S. Murthy, and S. Banerjee, "Region-based connectivity: a new paradigm for design of fault-tolerant networks," in *Proc. IEEE HPSR'09*, Jun. 2009.
- [37] U.S.-Canada Power System Outage Task Force, "Final report on the August 14, 2003 blackout in the United States and Canada: Causes and recommendations," Apr. 2004, <https://reports.energy.gov>.
- [38] U.S. Federal Energy Regulatory Commission, Dept. of Homeland Security, and Dept. of Energy, "Detailed technical report on EMP and severe solar flare threats to the U.S. power grid," Oct. 2010, <http://www.ornl.gov/sci/ees/etsd/pes/>.
- [39] Z. Wang, A. Scaglione, and R. Thomas, "Generating statistically correct random topologies for testing smart grid communication and control networks," *IEEE Trans. Smart Grid*, vol. 1, no. 1, pp. 28 –39, June 2010.
- [40] Western Electricity Coordinating Council (WECC), "Historical transmission paths database," <http://www.wecc.biz>.
- [41] H. Xiao and E. M. Yeh, "Cascading link failure in the power grid: A percolation-based analysis," in *Proc. IEEE Int. Work. on Smart Grid Commun.*, June 2011.

APPENDIX

Proof of Observation 3.1: Let f_{e_i} be the flow along line e_i , let x_{e_i} be its reactance and let θ_{v_i} be the phase angle of node v_i . By summing over the equalities of (2) for the lines of path p_1 we get $\theta_a - \theta_b = \sum_{e_i \in p_1} f_{e_i} x_{e_i}$. Similarly, for p_2 , $\theta_a - \theta_b = \sum_{e'_i \in p_2} f_{e'_i} x_{e'_i}$, and the claim follows.

Proof of Lemma 4.1: Consider an M -ring, and suppose that M is even (similar arguments can be used when M is odd). Assume that the capacity of all lines is 1, except lines $(M/2, M + M/2 + 1)$ and $(M/2, M + M/2 + 1)'$ whose capacity is 0.5. Assume also that $\alpha > 0$. Initially, the ring operates flawlessly, and all power flows are within the capacity of the corresponding lines.

Suppose that an initial failure event occur in lines $(0, M)$ and $(0, M)'$ (that is, a parallel lines failure). As described above, all lines will either carry flow of 0 or of 1, and therefore all lines except $(M/2, M + M/2 + 1)$ and $(M/2, M + M/2 + 1)'$ will continue operating normally.

As for lines $(M/2, M + M/2 + 1)$ and $(M/2, M + M/2 + 1)'$, their post-failure power flow is 1 while their pre-failure flow is 0.5. Since for each $\alpha > 0$, $\alpha \cdot 1 + (1 - \alpha) \cdot 0.5 > 0.5 = u_{M/2, M + M/2 + 1}$, these lines fault in the next round. Thus, the distance between consecutive failures in this cascade is $\Theta(M)$. As one can choose an arbitrarily large M -ring, the distance between there two consecutive failures can be made arbitrarily large and the claim follows.

Proof of Lemma 4.2: When the power capacity u of each line is 0.5, the same failure described in the proof of Lemma 4.1, which starts with 2 lines failure, causes an outage

of $3/5$ of the lines. Notice that in this case after the first iteration there is a demand shedding of half the total demand. Only odd lines still operate and each one of them carries a power flow of 0.5 unit (which is below its power capacity). Thus, this specific failure event stops after one iteration.

Proof of Lemma 4.3: Consider an M -ring and assume that the capacity of all lines is 0.5 , and, for ease of presentation, that $\alpha = 1$. Initially, the ring operates flawlessly, and all power flows are within the capacity of the corresponding lines.

Let $A = \{(0, M), (0, M)'\}$ (parallel lines failure). As described above, after a single iteration, all tie lines and odd lines will carry power flow of 1 unit, exceeding their capacity and therefore faulting. This will cause a demand shedding of 2 unit within Area 1 (no remaining active power lines); in the rest of the areas there will be a demand shedding of 1 unit (odd demand node will be disconnected), resulting in a total demand shedding of $M + 1$, and a total yield of $(M - 1)/2M < 0.5$.

On the other hand, assume an area failure of Area 0. Clearly, the area failure is a superset of A . However, the area failure event causes only demand shedding of 2 and therefore a yield of $(M - 1)/M > 0.5$ (for every $M > 2$). The reason is that in that case the cascade does not propagate outside Area 0, and all other power flows remain the same.

Proof of Lemma 4.4: Consider an M -ring and an M -ring with no tie-lines (which is a subgraph of the M -ring). Assume that the capacity of all lines is 0.5 , and, for ease of presentation, that $\alpha = 1$.

As we have shown before that a failure in $\{(0, M), (0, M)'\}$ causes a yield less than 0.5 in the M -ring. On the other hand, when there are no tie-line, this failure is contained within Area 1, resulting in a yield of at least $(2M - 2)/2M > 0.5$ (for $M > 2$). The same localization property holds for all other types of failures.

Proof of Lemma 4.5: Assume that the flow on line $(0, M)'$ is y . Thus, (1) and Observation 3.1 imply that the flow of each of the lines $(0, M + 1)$ and $(0, M + 1)'$ is $\frac{1}{2}(2 - y) = 1 - \frac{y}{2}$. In addition, (1) implies that the flow on the tie line connecting Area 0 and Area 1 is $1 - y$.

Focus on Area 1. Similar arguments yield that the flow on $(1, M + 2)$ and $(1, M + 2)'$ is $y/2$, the flow on $(1, M + 2)$ and $(1, M + 2)'$ is $1 - y/2$ and the flow on the tie line is $1 - y$. By symmetry, the same flow assignment holds for all other areas.

Now consider two paths between node 0 and node M . One traverses the single line $(0, M)'$ that carries y unit of power flow. The other traverses M tie lines, M odd lines, and $M - 1$ even lines; the power flow along even lines is in the opposite direction. Thus, the total flow along this path is $M(1 - y) + M(1 - y/2) - (M - 1)(y/2) = 2M - 2My + y/2$. Thus, by Observation 3.1, $y = 2M - 2My + y/2$ implying that $y = \frac{2M}{(0.5+2M)}$.

Proof of Corollary 4.6: The proof follows by case analysis. First, consider a failure event of a tie line. Since there is no power flow along the tie line, such a failure does not change the operations of the power grid and a capacity of 0.5 along internal lines and 0 along tie lines suffices to

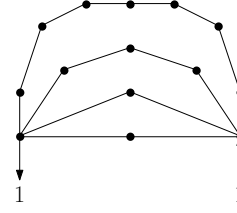


Fig. 15. Illustration of the proof of Lemma 4.7. The graph \mathcal{Q}_4 has 4 paths between its generation node and its demand node, two of length 2, one of length 4, and one of length 8.

withstand such a failure.

On the other hand, Lemma 4.5 analyzes a single internal line failures. In that case, as M goes to infinity, the post-failure flow on an internal lines approaches 1 . On the other hand, the maximum post-failure flow on the tie lines is $1/9$ when $M = 2$.

Proof of Lemma 4.7: For any $m > 2$, let $\mathcal{Q}_m = \langle \mathcal{CUD}, \mathcal{E} \rangle$ be an undirected graph with a single supply node and a single demand node, that are connected through m disjoint path in the following manner: the first two paths are of length 2 (implying that along each of these path there is an additional intermediate neutral node, connected to the generator on one side and to the demand node on the other side). For any $i > 2$, path i is of length 2^{i+1} . See Figure 15 depicting \mathcal{Q}_4 . Assume also that all lines have the same capacity $u = 1/2$ and the same reactance $x = 1$. We note that the total number of lines in \mathcal{Q}_m is 2^m and the total number of nodes is $2^m - m + 2$.

By Observation 3.1, the sum of flows along each of the m path is the same, denoted by y . In addition, since there is no generation or demand along each path and since each intermediate node has degree 2, (1) immediately implies that all lines along the same path carry the same amount of flow. Namely, the flow along the lines of the first two path is $y/2$ and for each line of path i ($i > 2$) the flow is $y/2^{i-1}$. Solving (1) for the demand node, implies that $y/2 + \sum_{i=1}^{m-1} y/2^i = 1$ and therefore $y = \frac{1}{1.5 - \frac{1}{2^{m-1}}} < 1$ (for any $m > 2$). Since the largest amount of flow on any line in the graph is $y/2 < 1/2 = u$, there power initially flow flawlessly.

Suppose now that there is a failure event in one of the lines along the first path. Applying (1) on the intermediate node implies that there is no flow along the other line of that path. We next show that this failure event causes a cascade that last m iterations, implying that cascades can be made arbitrarily large (by choosing larger graph \mathcal{Q}_m).

After each iteration ℓ , denote by y_ℓ the total amount of flow on each of the remaining path. We next show by induction that the flow $y_{\ell+1}/2^\ell$ along each of the edges of the $(\ell+1)$ -th path exceeds $1/2$, implying that these lines fault in the next iteration. In addition, for each $j > \ell + 1$ the flow along all lines of the j -th path is less than $1/2$, implying that they still operate at the end of the iteration.

In the base case, (1) for the demand node implies that $\sum_{i=1}^{m-1} y_1/2^i = 1$, and therefore $y_1 = \frac{1}{1 - \frac{1}{2^{m-1}}} > 1$, thus the lines along the second path have power flow of $y_1/2 > 1/2$

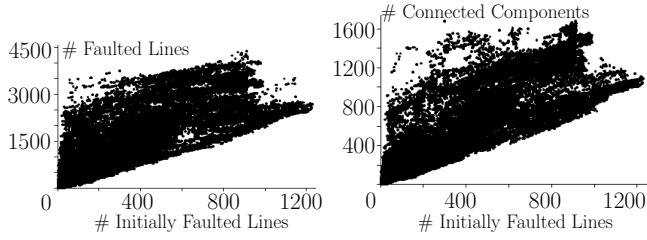


Fig. 16. The effects of the number of initially faulted lines on the number of faulted lines (left) and the number of components (right), after 5 rounds of cascade for FoS $K = 2$.

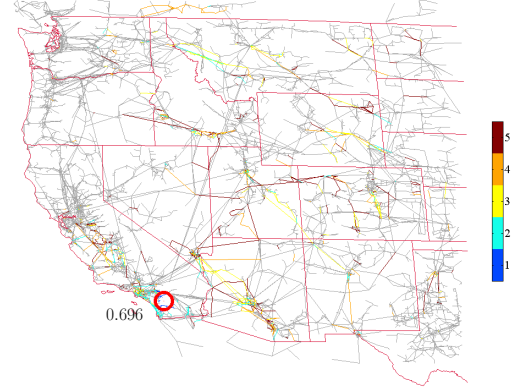
and they fail. On the other hand, for any $m > 2$, $y_1 < 2$, thus all lines of path $j > 2$ has capacity of at most $y_1/4 < 1/2$.

Suppose now that in iteration ℓ all lines of the first ℓ paths fault. We show that iteration $\ell + 1$ the lines of path $\ell + 1$, and the lines of the rest of the paths survive. The proof follows by solving (1) for the demand node with the surviving lines; namely, $\sum_{i=\ell}^{m-1} \frac{y_{\ell+1}}{2^i} = 1$, implying that $y_{\ell+1} = \frac{1}{\frac{1}{2^{\ell-1}} - \frac{1}{2^{m-1}}} = \frac{2^{m-1}}{2^{m-\ell}-1}$. Thus, the flow on the lines of the $(\ell + 1)$ -th path is $\frac{y_{\ell+1}}{2^\ell} = \frac{1}{2} \frac{2^{m-\ell}}{2^{m-\ell}-1} > \frac{1}{2}$, while the flow on the lines of any path of index larger than $\ell + 1$ is at most $\frac{y_{\ell+1}}{2^{\ell+1}} = \frac{1}{4} \frac{2^{m-\ell}}{2^{m-\ell}-1} < \frac{1}{2}$ (for any $\ell < m$).

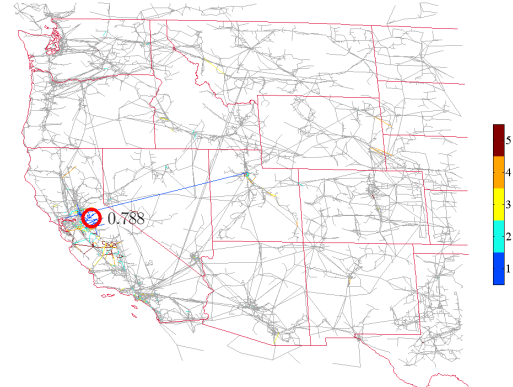
Hence, the above cascade over \mathcal{Q}_m lasts for m iterations, such that at each iteration the lines of one path fail. At the end of the cascade, the demand and generation nodes are disconnected and the yield is 0.

Cascades in an N -resilient grid with FoS $K = 2$:

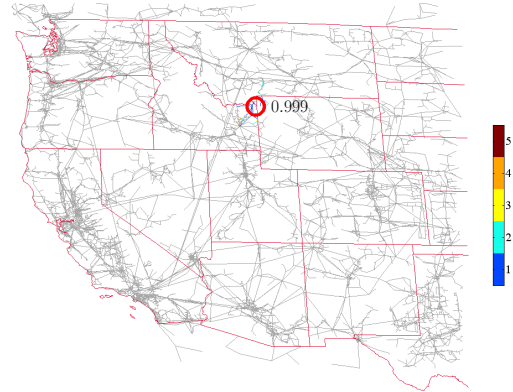
Figure 17 shows the first five rounds of three failures, initiated in three different locations. By comparing these results to the cascade resulting from the same failures, but when the grid has FoS $K = 1.2$ (Figure 3), one can clearly see than increasing the FoS significantly slows down the cascade. These results are further verified in Figure 16 which considers many failures on that grid and shows the effects of the number of initially faulted lines on the number of faulted lines after 5 rounds, as well as the number of components.



(a) San Diego area failure event



(b) San Francisco area failure event



(c) Idaho-Montana-Wyoming border failure event

Fig. 17. Illustration of the cascading failures from Figure 3 for N -resilient grid with FoS $K = 2$. The final yields are 0.696, 0.788, and 0.999.