

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

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

[†]School of Computer and Communication Sciences, EPFL, Lausanne, Switzerland

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

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

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

andrey.bernstein@epfl.ch, {dano@, meu2107@, gil@ee.}columbia.edu, dhay@cs.huji.ac.il

Abstract—We consider line outages in the transmission network of the power grid, and specifically those caused by natural disasters or large-scale physical attacks. In such networks, an outage of a line may lead to overload on other lines, thereby leading to their outage. Such a cascade may have devastating effects not only on the power grid but also on the interconnected communication networks. We study a model of such failures and show that it differs from other models used to analyze cascades (e.g., epidemic/percolation-based models). Inspired by methods developed for network-survivability analysis, we show how to identify the most vulnerable locations in the network. We also perform extensive numerical experiments with real grid data to estimate the effects of geographically correlated outages and briefly discuss mitigation methods. The developed techniques can indicate potential locations for grid monitoring, and hence, will have impact on the deployment of the smart-grid networking infrastructure.

Keywords—Survivability, Geographically-Correlated Failures, Cascading Failures, Power Grid, Network Science.

I. INTRODUCTION

Recent outages of the power grid (such as the Aug. 2003 blackout in the Northeastern United States and Canada [33]) demonstrated that large-scale failures can have devastating effects on almost every aspect in modern life, as well as on interdependent systems such as telecommunications networks [14]. The power grid is vulnerable to natural disasters, such as earthquakes, hurricanes, floods, and solar flares as well as to physical attacks, such as an Electromagnetic Pulse (EMP) attack [17], [34]. Thus, we focus on the vulnerability of the power grid to an outage of several lines in the same geographical area (i.e., to *geographically-correlated failures* which were recently studied in the network survivability community, e.g., [1], [15], [25], [32], [35], [40]).

Recent works focused on identifying a *few vulnerable lines* throughout the entire network [7], [8], [29] and on designing line or node interdiction strategies [31]. On the other hand, our objective is to identify *the most vulnerable areas* in the power grid and to characterize the properties of the cascade. Detection of the most vulnerable areas has various practical applications, since the system in these areas




Fig. 1. The development of the Sept. 8, 2011 San Diego blackout according to [11] (the color of each line represents the time in which it tripped). The cascade development is not contiguous.

can be either shielded (e.g., against EMP attacks or solar flares), strengthened (e.g., by increasing the capabilities of some relevant lines), or monitored. Since within the framework of the *smart grid*, significant effort is dedicated to identifying locations in which monitoring and communication equipment should be deployed (e.g., [39]), identifying potential monitoring locations will affect the communication network topology. Moreover, as indicated in [10], [17], [26], [28] there is a strong interdependence between the communication networks and the power grid (the grid supplies power to critical network components and the networks are used in order to control the grid). Hence, improving the grid’s resilience will have a direct impact on the communication networks’ resilience.

Unlike graph-theoretical network flows, power flows are governed by the laws of physics and there are *no strict capacity bounds on the lines* [4]. On the other hand, there is a *rating threshold* associated with each line – if the flow through a line exceeds the threshold, the line will eventually experience thermal failure (and will likely be turned off, or “tripped” for protection). Such an outage alters network topology, giving rise to a different flow pattern which, in turn, could cause other line outages. The repetition of this process constitutes a *cascading failure*. In this work we employ the (linearized) *direct-current (DC) power flow model*, which is a tractable relaxation of the exact alternating-current (AC) model, and an extension of the *cascading failure model* in [13] (also see [7], [8]). We note that cascading failures are *difficult to control in real time* [7], [8] because of their inherent multi-stage nature, problem size, and underlying noise.

Previous work (e.g., [12], [27] and references therein)

assumed that a line or bus (node) failure in a grid leads with some probability to a failure of nearby nodes or lines. Such epidemic-based modeling allows using percolation-based tools to analyze the cascade's effects. Yet, in real large scale cascades, a failure of a specific line can affect a remote line and the cascade does not necessarily develop in a contiguous manner. For example, the evolution of the the cascade that took place in the San Diego area on Sept. 8, 2011 appears in Fig. 1. It can be seen that consecutive line outages took place a few hops from each other. Motivated by this observation, we show that using the DC power flow model along with the cascading failure model leads to failure propagation characteristics that are significantly different from the epidemic-based models. For example, there exist graph topologies where failure of a line can lead to failure of a non-adjacent line.

Then, we focus on contingency events that are initiated by geographically correlated failures. We use a disk to represent the area affected by a contingency. 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.

Next, we present extensive numerical results which are based on the WECC (Western-Interconnect) real power grid data taken from the Platts Geographic Information System (GIS) [30]. The results are obtained by simulating the cascading failures for each of the possible disk centers (i.e., for each disk, our tool repeatedly and efficiently solved very large systems of equations). For example, Fig. 2 illustrates the 5 first rounds of a simulated cascade.

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. Our simulations identify vulnerable locations which are not “the usual suspects” (namely, highly-populated areas or power plants locations). We note that the San Diego blackout illustrated in Fig. 1 was caused by a single failure in such seemingly-negligible location (Yuma, AZ). 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 study the sensitivity of the results to different failure models (namely, stochastic vs. deterministic) and to different attack radii (naturally, the smaller the radius required to cause devastating effect, the more vulnerable the location). We also observe that while cascading failures usually start slowly and intensify over time when only few lines initially fault [7], [8], this *slow start* phenomenon often does not exist if the failures are geographically correlated.

Finally, we briefly consider control actions to be taken following a geographically correlated failure. We show, experimentally, that appropriate action taken at the appropriate time (not necessarily at the start of the cascade) can rapidly stop the cascade, while losing minimum demand.

The main contributions of this paper are twofold. First, we combine techniques from network survivability, optimization, and computational geometry to develop a method that allows obtaining extensive numerical results regarding the effects of




Fig. 2. Simulation of a geographically correlated failure in the San Diego area. All power lines, power plants, and substations within the marked 50 km radius disk were removed at round 1. The colors represent the rounds in which additional lines faulted. The final yield (the fraction of the original demand that was satisfied at the end of the cascade) is 0.326.

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*. The numerical results provide insights into the various effects of a geographically correlated failure. Second, we obtain analytical results regarding the cascade propagation under the power flow model which significantly differ from the epidemic/percolation-based results. We note that several network science tools were recently successfully used to analyze power networks (e.g., [3], [18], [20]–[23], [26], [28], [39] and references therein). This paper continues this line of work and is mostly inspired by the area of network survivability.

The rest of the paper is organized as follows. Section II reviews related work and Section III describes the power flow and cascade models. Section IV provides analytical results regarding the cascade propagation. Section V presents the algorithm used to identify the most vulnerable locations. Section VI describes the power grid data. In Section VII we present our numerical findings. Section VIII describes optimal control methods and demonstrates their impact. Section IX provides concluding remarks and directions for future work.

II. RELATED WORK

Geographically correlated failures have been recently studied in the context of communication networks [1], [15], [24], [25], [32], [35], [40] but to the best of our knowledge, they have not been extensively studied in the context of power networks. On the other hand, analysis of the power grid and its robustness have drawn a lot of recent attention, as part of efforts to develop smarter and more sustainable power infrastructure. The power grid is traditionally modeled as a complex system, made up of many components, whose interactions are not effectively computable (e.g., [4] and references therein).

Cascading failures are of high interest in the context of robustness of the grid [13]. Much work has focused on common topological properties of power grids and probabilistic failure propagation models, with the goal of evaluating the behavior of a generic grid as a self-organized critical system using, for example, *percolation theory* (see [12], [19], [36], [38], and references therein). These works are closely related to a long line of research in the power community which uses Monte Carlo methods to analyze system reliability (e.g., [9]).

Another major line of research focused on specific power flow models and used them to identify *small sets* of lines or stations whose simultaneous faults would create a severe contingency [7], [8], [22], [29], [31]. Since the problem is computationally intractable, most of these works used the DC model.

In our studies we use the real topology of the western U.S. power grid, or “interconnect” (WECC). While relying on the linearized model, we obtain numerical results for a very large scale real networks. In the past, results for large networks have been derived using mostly probabilistic models [12], [19]. Control mechanisms were also considered in [2].

The *interdependence* between the power and the communication networks upon failures in either network was investigated in [10], [26], [28]. Specifically, Parandehgheibi and Modiano [28] recently introduced a model that captures this interdependence and shows how failures cascade alternately between the networks. Our work complements [28] by studying cascades in the power grid network, implying a greater effect on the two interconnected networks.

III. MODELS

We adopt the *linearized* (or DC) *power flow model* that is widely used as an approximation for the non-linear AC power model (see [4] for a survey on the power flow models). In particular, we represent the power grid by a directed graph \mathcal{G} , whose set of nodes is \mathcal{N} , some of which are *supply nodes* (“generators”), or *demand nodes* (“loads”). 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 transmission lines. The orientation of the lines is arbitrarily and is used for notational convenience. We also assume *pure reactive* lines, characterizing each line (i, j) 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 balance, and take into account the reactance of each line. In addition, since the orientation of lines is arbitrary, a negative flow value simply means a flow in the opposite direction.

It is known that when \mathcal{G} is connected and $\sum_{i \in \mathcal{C}} P_i = \sum_{i \in \mathcal{D}} D_i$, (1)–(2) has a unique solution. As a result, this holds even when \mathcal{G} is not connected but the total supply and demand within each of its connected components are equal.

Next we describe the *Cascading Failure Model* (as appears above), which is an extension of the model in [13] (also see [7], [8]). We assume that each line (i, j) has a predetermined *capacity* u_{ij} , which bounds its power flow in a normal operation of the system (that is, $|f_{ij}| \leq u_{ij}$). We assume that before a failure event, \mathcal{G} is connected, the total supply and demand

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 $\mathcal{G}.\text{isStable} = \text{false}$.

While $\mathcal{G}.\text{isStable}$ is false 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 no line was removed at this step and the actual flow in all lines is below the capacity ($|f_{ij}| \leq u_{ij}, \forall (i, j)$), let $\mathcal{G}.\text{isStable} = \text{true}$.
-

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. Then, 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). 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, denoted by \tilde{f}_{ij}^t , where $\tilde{f}_{ij}^t = \alpha |f_{ij}| + (1 - \alpha) \tilde{f}_{ij}^{t-1}$ (in this paper, we mostly use $\alpha = 0.5$). To first order, the moving average approximates thermal effects, including heating and cooling from prior states. A similar moving average model was considered in [2], [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. We also consider the specific case of the *stochastic* rule, where $\varepsilon > 0$ and $p = 0.5$.

The process is repeated in rounds *until the system reaches stability*, namely until there is an iteration in which no lines are removed and there are no overloads. 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 values of α : 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:

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

In addition to 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.

We next discuss geographically-correlated initial failures. For simplicity, we assume that the initial failures are represented by disks with radius r . Namely, given a failure epicenter p , all lines within distance r from p initially fail¹. Let $Y(r, p)$ be the post-cascade yield of an initial failure of radius r whose epicenter is p . To capture the effect of different failure radii, we define the *critical failure radius* $r(Y_T, p)$ with respect to a yield threshold Y_T such that $r(Y_T, p) = \min_r Y(r, p) < Y_T$.

As will be shown later (Property 4.3), in some cases, failures do not have a monotonic impact, implying that $Y(r, p) > Y(r', p)$ even though $r > r'$. Thus, it is useful to classify epicenter points by their *monotonicity*. A point p is referred to as *monotone*, if for every $r > r'$, $Y(r, p) \leq Y(r', p)$. A point p is referred to as *pseudo-monotone*, if for every $r > r'$, $Y(r, p) \leq Y(r', p) + \delta$ (namely, allowing fluctuations of up to δ in the yield; typically we set $\delta = 0.05$). Otherwise, a point is *non-monotone*.

Note that our model contains a very simple control mechanism, namely, round-by-round demand shedding. In Section VIII, we consider a more elaborate control mechanism.

IV. CASCADING FAILURE PROPERTIES

In this section we describe important properties of the power flow and cascading failure models. Our goal is to highlight the differences between prior models used to characterize cascades in the power grid and the models we use. Note that these prior models did not compute power flows directly and assume some kind of contiguity in link failures. The proofs of these properties are based on constructing succinct examples of power grids and corresponding failures. They are omitted due to space constraints and can be found in [5].

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

Property 4.1 captures an important difference between our model and previously-suggested models, which 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 within a small distance) to fault, our model captures situations in which the cascade “skips” large distances within a single iteration. This situation can actually be observed in practice, such as in the recent major blackout in the San Diego area [11] shown in Fig. 1.

The following property shows that a devastating effect can be caused by an outage of a small fraction of lines.




Fig. 3. Classification of 1,870 potential failure locations of the WECC power grid by their monotonicity (see Section VI for details on this power grid and its parameters). For each location, the yield values resulting from failures of radii 10 km, 20 km, 30 km, 40 km, and 50 km are obtained. A location is monotone (and marked in blue) if the yield never increases, when the attack radius increases. Pseudo-monotonicity (marked in light blue) allows fluctuations of up to $\delta = 5\%$ in the yield. All other locations are non-monotone and marked in red.

Property 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 properties show that the failures do not always behave monotonically.

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

Property 4.4: For the same initial failure event, a network \mathcal{G}_1 , whose topology is a subgraph of the topology of another network \mathcal{G}_2 , may obtain a higher yield.

In practice, if the failures are geographically correlated, such non-monotone behavior rarely occurs. Fig. 3 shows the *monotonicity* of different failure locations (recall Section III for the exact definitions) in the WECC power grid. Out of the 1,870 considered locations, only 14.65% were non-monotone. As expected, non-monotonicity usually occurs in dense parts of the graph. Thus, in Section V-B, we assume a monotone behavior of the failures and, additionally, make sure that dense areas are examined thoroughly.

Finally, the next property shows that cascades can be made arbitrarily long (in time).

Property 4.5: The length of the cascade (the number of rounds until stability) can be arbitrarily large.

V. POWER GRID RESILIENCE

A. Parameters Set-up

In the cascading failure model, the power capacities u_{ij} of the lines are given a-priori. In practice, 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 consider the following two cases.

¹Our results can be easily extended to other attack shapes.

- **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 have FoS of at least $K \approx 1.2$ [16]. On the other hand, some data shows that certain lines (or, more generally, *paths*) are more resilient than others. For example, [37] shows that some transmission paths have power capacities which are 1.1 times their normal flow, while others have an FoS larger than 2. In addition, utility companies usually guarantee that their grid is at least $(N-1)$ -resilient [8]. Our experiments showed that an $(N-1)$ -resilient grid with an FoS K is essentially equivalent to an N -resilient grid with a higher FoS. Therefore, due to space constraints, we present numerical results for N -resilient grids with FoS $K = 1.2$. For extensive sensitivity analysis and results using $(N-1)$ -resilient grids and different FoS values K , see [6].

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).

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. (see Fig. 4).

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 . 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.

It was shown in [1] 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.




Fig. 4. The arrangement of hippodromes for two lines. The dots are the vertices and the dashed lines are the arcs.

We note that the method implicitly assumes *monotonicity* of failures, which as we showed does not always hold (see Property 4.3). However, Fig. 3 evidently shows that non-monotone locations are relatively rare, and that they usually appear only in dense areas. On the other hand, dense areas induce many vertices in the arrangement (as they contain many hippodromes’ intersections), guaranteeing that they are checked thoroughly.

VI. POWER GRID DATA

We use *real power grid data* of the western US taken from the Platts Geographic Information System (GIS) [30]. This includes the information about the transmission lines, power substations, power plants, and population at each 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. Moreover, 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 has 14,968 nodes (substations) and 19,513 lines. Moreover, it has 1,920 power stations, each of which was merged with the nearest substation. Notice 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 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 [19]. Thus, our results will probably carry over to other grids.

In order to obtain a connected graph of the grid that can be used in the simulation of the cascading failure model, we performed different processing steps of the raw data, including *longitude-latitude to planar coordinate transformation, connectivity check and nodes merging/elimination, and identifying demands and supplies*. As a result, we obtained a fully-connected graph (that is, a single connected component) with 13,992 nodes and 18,681 lines. 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. The details of these processing steps can be found in [5].

In addition, the GIS does not provide the power capacities of most of the transmission lines², 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

²The capacities are given in the GIS only for some congested lines.

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). Hence, we simply use the length of the line as its reactance. Regarding the power capacities, we take the approach described in Section V-A.

VII. NUMERICAL RESULTS

First, we note that as part of our experimental work, we performed a simulation of a recent major blackout event that occurred on Sept. 8th, 2011, in the San-Diego area. The results, reported in [6], suggest that our model follows closely the development of the actual event.

We identified the potential failure locations using the algorithm described in Section V-B implemented in MATLAB. We present results for failures with radius $r = 50$ km, which captures realistic scenarios such as an EMP attack [17], [34] and is large enough to generate a cascading failure in most cases. Due to space constraints, most of the results for other values of r are omitted, but we briefly discuss the critical failure radius below. For $r = 50$ km, the algorithm identified 61,327 potential failure locations. The identification of these locations was done within 24 hours on an eight-core server.

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 and Gurobi optimization tools. We present the results of simulation experiments for the N -resilient grid with FoS $K = 1.2$. In [6] we present results for an $(N-1)$ -resilient grid. The comparison between the results of N - and $(N-1)$ -resilience with the same FoS ($K = 1.2$) suggests that $(N-1)$ -resilience helps when the initial event is insignificant. However, it makes little difference when the initial event is significant.

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**. While the results do not point to a specific vulnerability of the power grid (due to the use of the modified grid map), they provide insights into the possible values of these metrics and the relations between them.

A. Deterministic Outage Rule

We first examined the deterministic outage rule, i.e. the rule defined in (3) with $\varepsilon = 0$. We plot specific failures to show how they evolved during the first *five* rounds of the cascade. Figs. 2 and 5 show two failure events: One in California, leading to a severe blackout, and another one around the Idaho-Montana-Wyoming border, which had a less severe effect. In general, higher FoS usually leads to less severe blackout effect (see [6] for details). Interestingly, the Idaho-Montana-Wyoming border




Fig. 5. Illustration of cascading failures over 5 rounds, where the initial failure location is in the Idaho-Montana-Wyoming border. The final yield values is 0.39. The colors represent the rounds in which the lines faulted. See also Fig. 2.




Fig. 6. 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 $K = 1.2$).

failure 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.

Scatter graphs for different metrics after 5 rounds are shown in Fig. 6. It can be seen that 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, and 1,200 lines initially faulted, at least 2,847, 3,600, and 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, and 1,200 lines initially faulted, the number of components is at least 696, 1,382, and 1,973, respectively.

Next, we analyze the severity of cascading failures once *stability* is reached. The results are shown in Figs. 7 and 8. In this case, the vast majority of failures resulted in yield in the range of 0.2–0.46. Fig. 7 also shows that, as expected, there is an inverse correlation between the yield and the total number of faulted lines. The relation between the number of rounds and the number of initially faulted lines suggests that when the number of initially faulted lines is small, the time until stability can be either large or small. However, as the number of initially faulted lines gets larger, the distribution of the time until stability becomes more concentrated. Also, our experiments indicate that there is no clear correlation between the final yield and the number of initially faulted lines as well as the final yield and the number of rounds until stability. Finally, Fig. 8 illustrates the yield and the number of rounds until stability by failure location. It can be seen that there is a high variability in the effects of the attacks in nearby locations and that a location that leads to low yield does not necessarily






Fig. 7. The number of rounds as a function of the number of initially faulted lines (left), and the yield as a function of total number of faulted lines (right)



(a) The yield values at stability.



(b) The number of rounds until stability.

Fig. 8. Vulnerability analysis (at stability) of failure locations. 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).

lead to high number of rounds (and vice versa).

B. Stochastic Outage Rule

In order to demonstrate the sensitivity to the assumptions regarding the outage rule, we also present two types of experiments using a *stochastic* outage rule, as defined in (3) with $\varepsilon > 0$ and $p = 0.5$. First, for the *same* failure epicenter, we compared the yield of different values of ε : Fig. 9(a) 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




Fig. 9. Results for the stochastic outage rule with $p = 0.5$. (a) the average yield and standard deviation of a representative failure epicenter, based on 100 independent runs. (b) a comparison of the deterministic and stochastic outage rules for selected failure events

epicenters were chosen such that the yield using deterministic rule grows approximately linearly with the failure index. The results, depicted in Fig. 9(b), 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 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).

C. Critical Failure Radius

We now demonstrate the relation between the initial failure radius and the resulting yield. In particular, we identify the failure radius that is required in order to obtain yield value below a certain yield threshold. The results demonstrate that for some attack epicenters, a failure of radius 10 km is sufficient for a devastating effect, while for some epicenters even a failure of radius 50 km has only a mild effect.

For example, we considered 1,870 equally spaced apart points. For each point, we considered 5 initial failure radii: 10 km, 20 km, 30 km, 40 km, and 50 km. Fig 10 depicts for each such point p , the value of $r(0.3, p)$ and $r(0.9, p)$: the minimal radius (out of the five radii checked) that leads to a post-cascade yield below 0.3 and 0.9, respectively.

VIII. CONTROL

This section describes experiments with control algorithms in the specific case of the San Diego event, illustrated in Fig. 2. The general goal of such algorithms is to stop the cascade without losing much demand (if possible, in a short time frame). It is well known that the exact optimal control of cascading failures becomes infeasible for large practical networks, due to various reasons (see, e.g., [7]). The main challenge is that any formulation has to deal with the combination of combinatorics in the network dynamics, multistage behavior, stochastic behavior, and very large size. It seems that this combination places the problem outside the capabilities of current optimization methodology. Therefore, some approximate or heuristic algorithms should be considered.

Particularly, we are interested in algorithms that will shed a minimum amount of demand and lead to a stable grid – the cascade has been stopped. Specifically, in this paper (for lack of space), we focus on an algorithm that operates within a single round of the cascade. The question is, then, *at which round* control should be applied. Assuming that a given round t is under consideration, our control is constrained as follows:

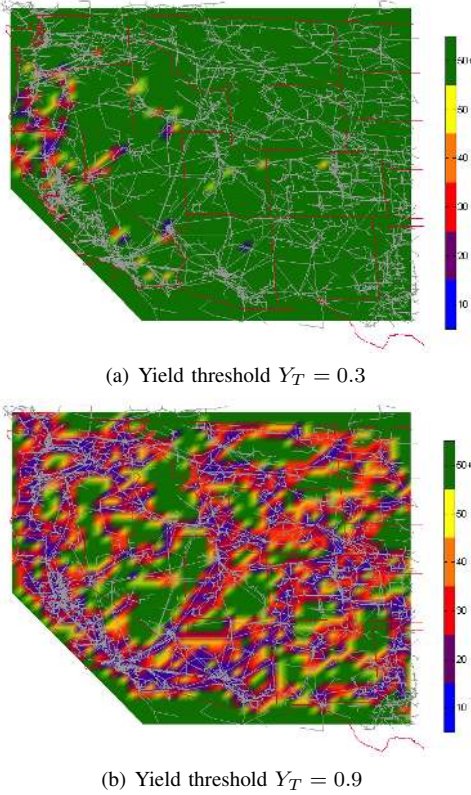


Fig. 10. The critical failure radius $r(Y_T, p)$ of 1,870 equally spaced apart points on the grid.

- (a) At each demand point $i \in \mathcal{D}$, we reduce the demand by a certain quantity, s_i .
- (b) We adjust generator output, within each component so as to maintain overall balance between supply and demand.
- (c) However, generators are furthermore constrained in that the amount of change in a generator must be proportional to its current output.
- (d) After the demand shedding and generator adjustments, the power flow on each operating line (i, j) , and its moving-average, 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 the cascading failure model, the cascade will stop. Rules (a)-(d) describe the constraints; the goal is to pick the round t and quantities s_i so as to maximize the remaining demand. At a given round t , this optimization problem can be written as a linear program. Denote by $\tilde{f}_{ij}^t, \tilde{D}_i^t, \tilde{P}_i^t$ the value, just before round t , of the flow moving average on line (i, j) , the demand at demand point $i \in \mathcal{D}$, and the generation at supply point $i \in \mathcal{C}$, respectively. Moreover, denote by C_1, \dots, C_n the connected components in the grid graph before round t , and let $\text{comp}(i)$ denote the connected component that contains node i . The

TABLE I. OPTIMAL CONTROL OUTCOME. “ROUND” REFERS TO THE ROUND ON WHICH THE CONTROL IS APPLIED, WHILE “YIELD” IS THE 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

linear program is as follows:

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

where, as in Section III, $\delta^+(i)$ ($\delta^-(i)$) is the set of lines oriented out of (into) node i . Notice that 3rd – 6th equations in the linear program above are identical to (1)-(2) in Section III.

We demonstrate the operation of the control mechanism by considering a failure event in San Diego area (i.e., in the location illustrated in Fig. 2) under the stochastic outage rule with $\varepsilon = 0.05$, $p = 0.5$, and $\alpha = 0.1$.

Table I outlines the performance of the control mechanism and shows that *neither* applying control at the outset of the cascade is optimal (this is typical, in our experience), nor waiting too long. 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 the 74th round with yield value of 0.34. Currently, we are developing robust versions of this algorithm with respect to errors in data, timing, and delays in implementation.

IX. 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 models based on epidemic-like failures. Then, we used techniques from network survivability analysis along with detailed GIS data to develop a method for analyzing the vulnerability of different grid locations to geographically correlated failures. We performed extensive numerical experiments to obtain insight into the relations between the various parameters and performance metrics. Finally, we demonstrated that the use of control at the right point can mitigate the effects of a large scale failure.

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. In particular, we plan to extend this work to study the effect of geographically correlated failures on the *interdependent grid and communication networks* (see e.g., [10], [26], [28]). Moreover, we will extend the cascade model to incorporate effects such as frequency collapse and generator

tripping. In addition, while due to its relative simplicity, most previous work in the area of grid vulnerability is based on the DC model, this model does not capture effects such as voltage collapse that may occur during a cascade. Hence, we plan to develop methods to analyze the effect of an attack using the more realistic AC model. Finally, we plan to study the effectiveness of some of the existing control algorithms in coping with geographically correlated failures.

ACKNOWLEDGEMENTS

This work was supported by NSF grant CNS-1018379, NSF CIAN ERC under grant EEC-0812072, DTRA grants HDTRA1-09-1-0057 and HDTRA1-13-1-0021, DOE award DE-SC0002676, a grant from the U.S.-Israel Binational Science Foundation, the Legacy Heritage Fund program of the Israel Science Foundation under Grant No. 1816/10, the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), and the Israeli Smart Grid (ISG) Consortium, administered by the Office of the Chief Scientist of the Israeli ministry of Industry and Trade and Labor.

REFERENCES

- [1] P. Agarwal, A. Efrat, A. Ganjugunte, D. Hay, S. Sankaraman, and G. Zussman, "The resilience of WDM networks to probabilistic geographical failures," *IEEE/ACM Trans. Netw.*, vol. 21, no. 5, pp. 1525–1538, 2013.
- [2] M. Anghel, K. A. Werley, and A. E. Motter, "Stochastic model for power grid dynamics," in *Proc. HICSS'07*, Jan. 2007.
- [3] O. Ardakanian, C. Rosenberg, and S. Keshav, "On the use of teletraffic theory in power distribution systems," in *Proc. e-Energy'12*, May 2012.
- [4] A. R. Bergen and V. Vittal, *Power Systems Analysis*. Prentice-Hall, 1999.
- [5] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman, "Power grid vulnerability to geographically correlated failures – analysis and control implications," *ArXiv e-prints*, 2012, available at <http://arxiv.org/abs/1206.1099v1>.
- [6] —, "Sensitivity analysis of the power grid vulnerability to large-scale cascading failures," *SIGMETRICS Perform. Eval. Rev.*, vol. 40, no. 3, pp. 33–37, Dec 2012.
- [7] D. Bienstock, "Optimal control of cascading power grid failures," in *IEEE 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] R. Billinton and W. Li, *Reliability Assessment of Electrical Power Systems Using Monte Carlo Methods*. Plenum Press, 1994.
- [10] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin, "Catastrophic cascade of failures in interdependent networks," *Nature*, vol. 464, pp. 1025–1028, 2010.
- [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] J. Cowie, A. Ogielski, B. Premore, E. Smith, and T. Underwood, "Impact of the 2003 blackouts on Internet communications," *Preliminary Report, Renesys Corporation*, 2004.
- [15] 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*, Mar. 2010.
- [16] I. Dobson, "personal communication," 2012.
- [17] J. S. Foster, E. Gjeldel, W. R. Graham, R. J. Hermann, H. M. Kluepfel, R. L. Lawson, G. K. Soper, L. L. Wood, and J. B. Woodard, "Report on the commission to assess the threat to the United States from electromagnetic pulse (EMP) attack, critical national infrastructures," Apr. 2008.
- [18] L. Gan, A. Wierman, U. Topcu, N. Chen, and S. H. Low, "Real-time deferrable load control: Handling the uncertainties of renewable generation," in *Proc. e-Energy'13*, 2013.
- [19] 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.
- [20] Y. Huang, S. Mao, and R. M. Nelms, "Adaptive electricity scheduling in microgrids," in *Proc. IEEE INFOCOM'13*, 2013.
- [21] N. Li, L. Chen, and S. H. Low, "Optimal demand response based on utility maximization in power networks," in *IEEE PES General Meeting*, July 2011.
- [22] X. Liu, K. Ren, Y. Yuan, Z. Li, and Q. Wang, "Optimal budget deployment strategy against power grid interdiction," in *Proc. IEEE INFOCOM'13*, 2013.
- [23] L. Lu, J. Tu, C.-K. Chau, M. Chen, and X. Lin, "Online energy generation scheduling for microgrids with intermittent energy sources and co-generation," in *Proc. ACM SIGMETRICS'13*, 2013.
- [24] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the vulnerability of the fiber infrastructure to disasters," *IEEE/ACM Trans. Netw.*, vol. 19, no. 6, pp. 1610–1623, 2011.
- [25] S. Neumayer and E. Modiano, "Network reliability with geographically correlated failures," in *Proc. IEEE INFOCOM'10*, Mar. 2010.
- [26] —, "Assessing the effects of a randomly located circular disaster on the power grid and dependent networks," in *Proc. IEEE SmartGridComm'13*, Oct. 2013.
- [27] G. A. Pagani and M. Aiello, "The power grid as a complex network: a survey," *Physica A: Statistical Mechanics and its Applications*, vol. 392, no. 11, pp. 2688 – 2700, 2013.
- [28] M. Parandehgheibi and E. Modiano, "Robustness of interdependent networks: The case of communication networks and the power grid," in *Proc. IEEE GLOBECOM'13*, Dec. 2013.
- [29] 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.
- [30] Platts, "GIS Data," <http://www.platts.com/Products/gisdata>.
- [31] 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.
- [32] A. Sen, B. Shen, L. Zhou, and B. Hao, "Fault-tolerance in sensor networks: a new evaluation metric," in *Proc. IEEE INFOCOM'06*, 2006.
- [33] 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>.
- [34] 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/>.
- [35] J. Wang, C. Qiao, and H. Yu, "On progressive network recovery after a major disruption," in *Proc. IEEE INFOCOM'11*, 2011.
- [36] 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.
- [37] Western Electricity Coordinating Council (WECC), "Historical transmission paths database," <http://www.wecc.biz>.
- [38] 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.
- [39] Y. Zhao, A. Goldsmith, and V. Poor, "On PMU location selection for line outage detection in wide-area transmission networks," in *IEEE PES General Meeting*, July 2012.
- [40] Q. Zheng, G. Cao, T. L. Porta, and A. Swami, "Optimal recovery from large-scale failures in IP networks," in *Proc. IEEE ICDCS'12*, 2012.