# Sensitivity Analysis of the Power Grid Vulnerability to Large-Scale Cascading Failures

Andrey Bernstein*, Daniel Bienstock*, David Hay†, Meric Uzunoglu*, Gil Zussman*

*Columbia University
New York, NY 10027

†The Hebrew University
Jerusalem, Israel

{andreyb@ee, dano@ , meu2107@, gil@ee}.columbia.edu, dhay@cs.huji.ac.il

## ABSTRACT

This paper revisits models of cascading failures in the transmission system of the power grid. It has been recently shown that since power flows are governed by the laws of physics, these models significantly differ from epidemic/percolation-based models. Yet, while some numerical results have been recently obtained based on these models, there is a need to investigate the sensitivity of the results to various parameters and to evaluate the models' accuracy. In this paper, through numerical experiments with real grid data, we study the effects of geographically correlated outages and the resulting cascades. We consider a wide range of parameters, such as the power lines' Factor of Safety and the sensitivity of the lines to power flow spikes. Moreover, we compare our numerical results to the actual events in a recent blackout in the San Diego area (Sept. 2011), thereby demonstrating that the model's predictions are consistent with real events.

**Categories and Subject Descriptors:** C.4 [Computer Systems Organization]: Performance of Systems — *Reliability, availability, and serviceability*

**General Terms:** Design, Reliability, Performance

**Keywords:** Power Grid, Geographically-Correlated Failures, Cascading Failures, Resilience, Survivability.

## 1. INTRODUCTION

The power grid is vulnerable to natural disasters, such as earthquakes and solar flares [22] as well as to physical attacks, such as an Electromagnetic Pulse (EMP) attack [14, 22]. Failures following such large-scale events will have devastating effects on almost every aspect in modern life, as many systems (e.g., telecommunications, gas and water supply, and transportation) highly depend on the power grid. These adverse effects were recently demonstrated in several large-scale blackouts in North America (e.g., the Aug. 2003 blackout in the northeastern United States and in Canada [21], and the Sept. 2011 blackout in southwestern United States and in Mexico [7]), Europe (e.g., the Sept. 2003 blackout in Italy [10] and the Apr. 2012 blackout in Cyprus [16]), and Asia (India's blackout in July 2012 which left 670 million people without power [24]).

Hence, the power grid and its robustness have drawn a lot of attention recently [2], where cascading failures have been a major concern [2, 9, 13]. One approach to study these phenomena is by probabilistic failure propagation models (e.g., based on percolation theory) on graphs with specific topological properties, which are common to contemporary power grids [8, 15, 17].



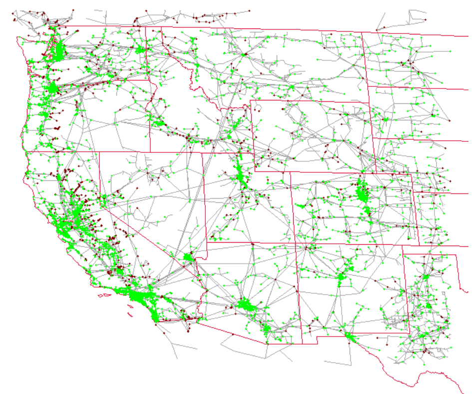

**Figure 1: The power grid graph of the Western Interconnect (including some sections of neighboring states' grids). Green dots represent demand nodes and red dots represent supply nodes.**

Alternatively, in this paper, as in [4], we focus on an approach which is based on (microscopic) power flow models. Specifically, the flow in the power grid is governed by the laws of physics and there are *no strict capacity bounds on the lines* [3]. Instead, 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 use a *linearized (i.e., DC) power flow model* and a *cascading failure model* (originated from [9]) to obtain results despite the problem's complexity. Previous works use these models to identify a *few vulnerable lines* throughout the network [5, 6, 19], where the initial failure events (causing eventually the cascading failures) are assumed to be sporadic link outages, with no correlation between them. On the contrary, we focus on *geographically-correlated failures*: events that cause *a large number of failures in a specific geographical region* (e.g., [14, 22]).

To the best of our knowledge, geographically-correlated failures have been studied only in the context of communication networks [1, 11]. Recently, a survivability analysis of power networks to geographically-correlated failures was performed in [4]. This paper complements the results of [4] by providing a sensitivity analysis of the results. Specifically, we present the results of numerical experiments performed on the real U.S. power grid data, taken from the Platts Geographic Information System (GIS) [20] (see Figure 1). The results provide a sensitivity analysis of the models to various parameters, such as the grid's a-priori resilience, the power lines' Factor of Safety, and the sensitivity to power flow spikes. We also compare the results obtained using our model to the events in a real cascade which took place in the San Diego area on Sept. 8, 2011 [7]. This allows us to assess the accuracy of our methods and parameters.

# 2. PRELIMINARIES

## 2.1 The Linearized Power Flow Model

We follow [6] and represent the power grid as a directed graph $\mathcal{G} = \langle \mathcal{N}, \mathcal{E} \rangle$, whose nodes are either *supply nodes* ("generators"), *demand nodes* ("loads"), or *neutral nodes*. Let $\mathcal{D} \subseteq \mathcal{N}$ denote 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 $(i,j) \in E$ is characterized by its *reactance* $x_{ij}$.

We adopt the *linearized* (or DC) *power flow model*, which is widely used as an approximation for more realistic nonlinear AC power model [3]. Given supply and demand vectors $(P, D)$, a *power flow* is a solution $(f, \theta)$ 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 $\theta_i$ is the phase angle of node $i$. The equations guarantee power flow conservation and consider the reactance of each line.

## 2.2 Cascading Failure Model

We use the cascading failure model proposed in [9] (also used in [4–6,18]). 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 system (that is, $|f_{ij}| \le u_{ij}$). Upon a failure, some lines in a specific geographic areas are removed from the graph. Specifically, we consider a circular and deterministic initial failure, 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). This, in turn, implies that the graph may become disconnected. Thus, within each component, the demand (supply) is decreased by the same factor at all loads (generators), so that the total demand is equal the total supply. Then, the power flows in the new graph are recalculated by (1)–(2). 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 factors, each of which could cause failure. Such outages are modeled by a *moving average* of the power flow $\tilde{f}^t_{ij} = \alpha|f_{ij}| + (1-\alpha)\tilde{f}^{t-1}_{ij}$. Lines $(i,j)$ whose $\tilde{f}^t_{ij}$ is above the power capacity $u_{ij}$ are removed from the graph. To first order, this approximates thermal effects, including heating and cooling from prior states. The process is repeated in rounds *until the system reaches stability*, namely until there are no overloaded lines in the graph. The 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 $\alpha$. Smaller value of $\alpha$ implies that we take a more microscopic look at the cascade.

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

$$Y \triangleq \frac{\text{The actual demand at the stability}}{\text{The original demand}}. \quad (3)$$

## 2.3 Parameters Set-up

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 grid. In this paper, we take the $N-k$ contingency analysis approach [6] 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 line capacities 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 two cases: (i) $N$**-resilient grids** (that is, $k = 0$): we solve (1)–(2) for the original grid graph (without failures) and set the power capacity to $u_{ij} = K \cdot f_{ij}$, where $K \ge 1$; and (ii) $(N-1)$**-resilient grids** (that is, $k = 1$): we solve (1)– (2) for $N$ graphs, each resulting from a single line failure event. The power capacity is set to $u_{ij} = K \cdot \max_r f^r_{ij}$, where $f^r_{ij}$ is the flow assigned to line $(i,j)$ when considering the $r$-th failure event.

The real power grid is usually assumed to be at least $N$-resilient with $K \approx 1.2$ [12]. On the other hand, some data shows that certain lines (or, more generally, *paths*) are more resilient than others. For example, a historical transmission paths data found in [23] shows that some transmission paths have power capacities which are 1.1 times their normal flows, 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 [6].[1] Therefore, we examine in this paper both $N$- and $(N-1)$-resilient grids with different FoS values $K$.

# 3. SENSITIVITY ANALYSIS

We use *real power grid data* of the western U.S. taken from the Platts Geographic Information System (GIS) [20]. This includes the information about the transmission lines, power substations, power plants, and the population at each geographic location. The nodes of the graph are substations, while the arcs are the transmission lines. In order not to expose the vulnerability of the real grid, we used a *part* of the Western Interconnect system which does not include most of 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. In order to obtain a connected graph of the network that can be used to simulate the cascading failure

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

model, we performed different processing steps of the raw data, whose details can be found in [4]. At the end of this process, we obtained a fully-connected graph with 13,992 nodes and 18,681 lines. Overall, 1,117 nodes were classified as generators (supply nodes), 5,591 as loads (demand nodes), and 7,284 as neutral. We assumed that all lines have the same physical properties (other than length) and used the length to determine the reactance.

We present results for an attack radius $r = 50$ kilometers, which is small enough to capture realistic scenarios [14, 22], while it is large enough to generate a cascading failure in most cases. A computational-geometry–based algorithm to discretisize the entire plane to finite search space was presented in [1, 4]. In our case, the algorithm identified 61,327 potential failure locations. For each failure location $v$, we performed the simulation of the cascading failure model (Section 2), assuming that all lines within radius $r$ of $v$ fail. Our program used CPLEX and Gurobi optimization tools in order to solve the power flow equations efficiently.

Most of our results using the $N$ contingency analysis to set the capacities of the network, with FoS $K = 1.2$, can be found in [4]. In this paper, we perform a sensitivity analysis for different FoS values and also use $N - 1$ contingency analysis for comparison.

The first set of experiments was performed using the $N$ contingency analysis with FoS $K = 2$. We plot specific failures to show the evolution during the first *five* rounds of the cascade. Figure 2 shows three failure: Two in California, both leading to severe blackouts, and another one around the Idaho-Montana-Wyoming border, which had a negligible effect. A comparison with the results in [4] shows that 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 faulted. However, the same event with $K = 2$ leads to near-unity yield. 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 2.3 for more details).

We now analyze the failures severity once *stability* is reached (namely, when no more line failures occur) using the $N - 1$ contingency analysis to set the capacities of the network, with FoS $K = 1.2$. Figure 3 shows the yield values upon cascades which start at the possible attack locations that we identified. It can be seen that the most vulnerable areas are the highly populated areas of the West Coast: San-Francisco, Los Angeles, and Seattle. We compare these results to those of $N$-resilience experiments from [4]. As expected, it can be seen 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, since this part is connected to the whole network using a small number of lines (which in practice carry no power in normal operation). However, when the network is only $N$ resilient, these failures do propagate to the whole network[2].

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

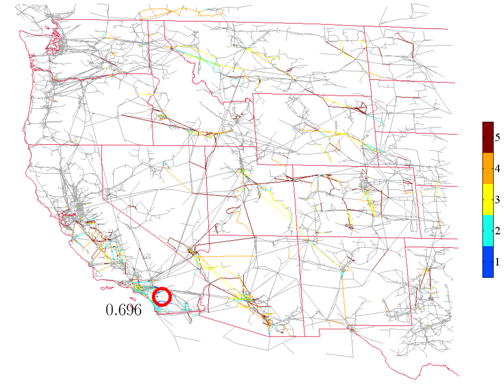


(a) San Diego area failure event

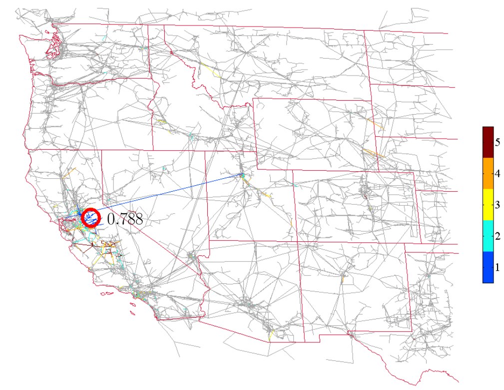


(b) San Francisco area failure event


(c) Idaho-Montana-Wyoming border failure event

**Figure 2: Illustration of cascading failures over 5 rounds for $N$-resilient grid with FoS $K = 2$. The colors represent the rounds in which the lines faulted. The final yields are (a) 0.696, (b) 0.788, and (c) 0.999.**

## 4. SAN DIEGO BLACKOUT (SEPT. 2011)

### 4.1 Description of the Blackout

On Sept. 8$^{\text{th}}$, 2011, over 2.7 million people in southwestern U.S. and in Mexico experienced a massive power blackout. Although the full details are not known yet, several publicly available sources, such as [7], 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, since such blackouts are rare, we use this blackout in order to calibrate the various model parameters so that the simulation results

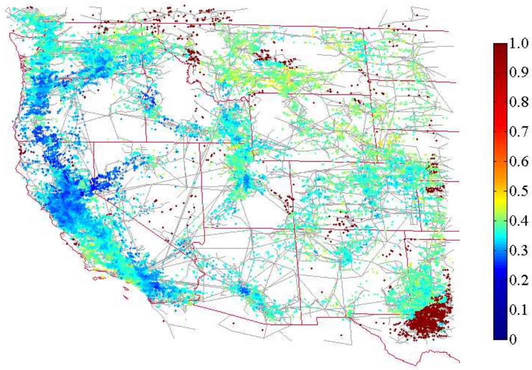


**Figure 3: Vulnerability analysis (at stability) of failure locations for $N-1$-resilient grid with FoS $K = 1.2$. The color of each point represents the yield value corresponding to a cascade whose epicenter is at that point.**

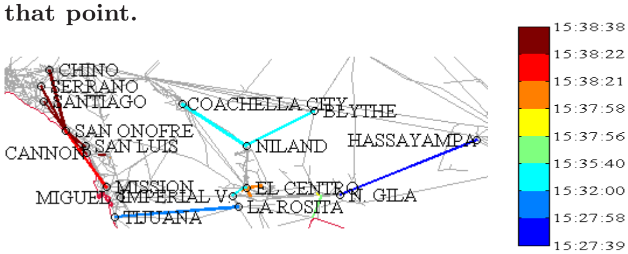


**Figure 4: The development of the San Diego blackout according to [7].**

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 4. There are two import generation paths into this area: (i) *SWPL*: represented by the 500KV Hassayampa-North Gila-Imperial Valley-Miguel transmission line. This path transfers power from the Palo Verdi Nuclear Generating Station in Arizona; and (ii) *Path 44*: 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 [7] (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 blackout. The actual cascade development is shown in Figure 4.

## 4.2 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

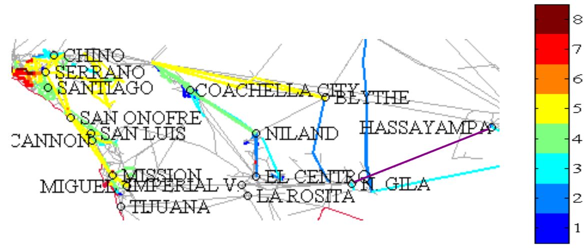


**Figure 5: The development of the San Diego blackout in the first eight rounds using our simulation.**

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

Moreover, since in the actual event there was no $(N-1)$-resilience with respect to the faulted North Gila–Hassyampa line, we used an $N$-resilient grid with different values of FoS $K$ (recall Section 2.3). In addition, according to [7], 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., [23]) 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 adjusted its FoS accordingly and set it to 2.5. After experimenting with the value of $K$ for other lines, we found that $K = 1.5$ leads to a behavior that most resembles that of the actual event. The resulting cascade development is shown in Figure 5.

Table 1 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 [7], 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 corresponding control mechanisms that will allow stopping 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 $\alpha$ on both the maximum line overloads and the length of the cascade. Large $\alpha$ corresponds to the case in which the faults are mostly determined by the very recent flow values, while small $\alpha$ corresponds to the case in which the historical flow values affect the faults. In other words, with large $\alpha$ values, lines are more sensitive to power flow spikes. The results (see Figure 6) show that the larger $\alpha$ is, the higher is the maximum load and the shorter is the cascade. Moreover, when $\alpha$ 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 applied at that time will stop the cascade with relatively high yield in early rounds.

# 5. CONCLUSION

We performed a sensitivity analysis of the power grid cascading failure model. We used detailed GIS data of the western U.S. power grid and performed numerical experiments that demonstrate the effects of different parameters. In addition, we used a recent major blackout event in the

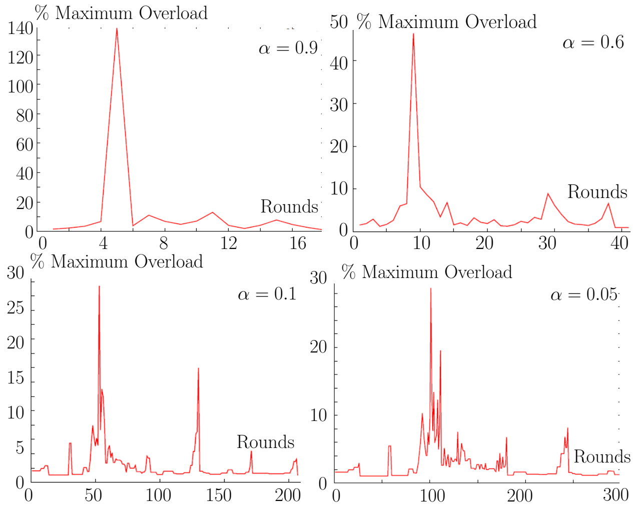


**Figure 6: Maximum line overload for different values of $\alpha$.**

| 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 1: Comparison of the actual event and the simulation results.**

San Diego area as a case study to demonstrate the consistency of the models and to calibrate different parameters. In general, the vulnerability 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.

## 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 U.S.-Israel Binational Science Foundation, NSF grant CNS-10-18379, the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), and the Israeli Smart Grid consortium. We thank Eric Glass for help with GIS data processing.

# 6. 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] A. R. Bergen and V. Vittal. *Power Systems Analysis*. Prentice-Hall, 1999.

[4] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman. Geographically correlated failures in power networks – survivability analysis. Submitted, 2012.

[5] D. Bienstock. Optimal control of cascading power grid failures. *in IEEE PES General Meeting*, July 2011.

[6] D. Bienstock and A. Verma. The $N-k$ problem in power grids: New models, formulations, and numerical experiments. *SIAM J. Optim.*, 20(5):2352–2380, 2010.

[7] 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`.

[8] D. P. Chassin and C. Posse. Evaluating north american electric grid reliability using the Barabsi–Albert network model. *Physica A*, 355(2-4):667 – 677, 2005.

[9] 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.*, 27(4):318 – 326, 2005.

[10] S. Corsi and C. Sabelli. General blackout in Italy Sunday September 28, 2003, h. 03:28:00. In *IEEE PES General Meeting*, June 2004.

[11] T. N. Dinh, Y. Xuan, M. T. Thai, P. M. Pardalos, and T. Znati. On new approaches of assessing network vulnerability: hardness and approximation. *IEEE/ACM Trans. Netw.*, 20(2):609–619, Apr. 2012.

[12] I. Dobson. personal communication, 2010.

[13] I. Dobson, B. Carreras, V. Lynch, and D. Newman. Complex systems analysis of series of blackouts: cascading failure, critical points, and self-organization. *Chaos*, 17(2):026103, June 2007.

[14] J. S. Foster, E. Gjelde, 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.

[15] P. Hines, E. Cotilla-Sanchez, and S. Blumsack. Do topological models provide good information about electricity infrastructure vulenrablity? *Chaos*, 20(3):033122, Sept. 2010.

[16] msnbc.com staff. Blackout nation: fault leaves cyprus without power, 2012. `http://worldnews.nbcnews.com/_news/2012/04/04/11016803-blackout-nation-fault-leaves-cyprus-without-power?lite`.

[17] G. A. Pagani and M. Aiello. The power grid as a complex network: a survey. *ArXiv e-prints*, May 2011.

[18] R. Pfitzner, K. Turitsyn, and M. Chertkov. Controlled tripping of overheated lines mitigates power outages. *ArXiv e-prints*, Oct. 2011.

[19] A. Pinar, J. Meza, V. Donde, and B. Lesieutre. Optimization strategies for the vulnerability analysis of the electric power grid. *SIAM J. Optim.*, 20(4):1786–1810, 2010.

[20] Platts. GIS Data. `http://www.platts.com/Products/gisdata`.

[21] 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`.

[22] 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/`.

[23] Western Electricity Coordinating Council (WECC). Historical transmission paths database. `http://www.wecc.biz`.

[24] J. Yardley and G. Harris. India staggered by power blackout; 670 million people in grip, 2012. `http://www.nytimes.com/2012/08/01/world/asia/power-outages-hit-600-million-in-india.html/`.