Computational Analysis of Cascading Failures in Power Networks

Dorian Mazauric
Laboratoire d'Informatique Fondamentale de Marseille
Marseille, France
dorian.mazauric@lif.univ-mrs.fr

Saleh Soltan, Gil Zussman
Electrical Engineering
Columbia University
New York, NY
{saleh,gil}@ee.columbia.edu

ABSTRACT

We focus on *cascading line failures in the transmission system of the power grid*. Recent large-scale power outages demonstrated the limitations of epidemic- and percolation-based tools in modeling the cascade evolution. Hence, based on a linearized power flow model, we obtain results regarding the various properties of a cascade. Specifically, we consider performance metrics such as the distance between failures, the length of the cascade, and the fraction of demand (load) satisfied after the cascade. We show, for example, that due to the unique properties of the model: (i) a set of initial line failures may have a smaller effect than a failure of one of the lines in the set, (ii) the distance between subsequent failures can be arbitrarily large and the cascade may be arbitrarily long, and (iii) minor changes to the network parameters may have a significant impact. Moreover, we show that finding the set of lines whose removal has the most significant impact (under different metrics) is NP-hard. Finally, for specific graphs, we develop a fast algorithm to determine if a set of line failures initiates a cascade. The results can provide insight into the design of smart grid measurement and control algorithms that can mitigate a cascade.

Categories and Subject Descriptors: C.4 [Performance of Systems]: Reliability, availability, and serviceability; G.2.2 [Discrete Mathematics]: Graph Theory–*Graph algorithms, Network problems.*

Keywords: Power Grid, Cascading Failures, Performance Metrics, Computational Complexity, Survivability.

1. CASCADING FAILURE MODEL

We adopt the linearized (or DC) *power flow model*, which is widely used as an approximation for the non-linear AC power flow model. In particular, we follow [2, 3] and represent the power grid by a graph $G = (V, E)$ where V and E are the set of nodes and edges representing the buses and transmission lines, respectively. P_v is the active power supply or demand at node $v \in V$. Each node is classified either as a *supply node* ($P_v > 0$), a *demand node* ($P_v < 0$), or a *neutral node* ($P_v = 0$). We assume *pure reactive* lines, implying that each edge $\{u, v\} \in E$ is characterized by its *reactance* $x_{uv} = x_{vu}$.

Given an active power vector P, a *power flow* is a solution

Cascading Failure Model

Input: A connected graph $G = (V, E)$ and an initial edge failures event $F_0 \subseteq E$.

Output: The length of the cascade $t \geq 0$, the sequence $(F_0, F_1, \ldots, F_t)$ of the sets of edge failures at each round, and the power flows $f_e(F_t^*)$ $\forall e \in E$, at stabilization.

$F_0^* \leftarrow F_0$ and $i \leftarrow 0$.
while $F_i \neq \emptyset$ **do**
  Adjust the total demand to total supply within each connected component of $G = (V, E \backslash F_i^*)$ by decreasing the demand (supply) by the same factor at all demand (supply) nodes.
  Compute the new flows $f_e(F_i^*)$ $\forall e \in E$.
  Find the set of new edge failures F_{i+1}.
  $F_{i+1}^* \leftarrow F_i^* \cup F_{i+1}$ and $i \leftarrow i + 1$.
return $t = i - 1$, $(F_0, \ldots, F_t)$, and $f_e(F_t^*)$ $\forall e \in E$.

(f, θ) of the following system of equations:

$$\sum_{v \in N(u)} f_{uv} = P_u, \ \forall \ u \in V \tag{1}$$

$$\theta_u - \theta_v - x_{uv} f_{uv} = 0, \ \forall \ \{u, v\} \in E \tag{2}$$

where $N(u)$ is the set of neighbors of node u, f_{uv} is the power flow from node u to node v, and θ_u is the phase angle of node u. Note that the edge capacities are not taken into account in determining the flows.

The *cascading failure model* described above is similar to the model used in [1, 3, 4]. We assume that each edge $e = \{u, v\} \in E$ has a predetermined power capacity $c_e = c_{uv} = c_{vu}$, which bounds its power flow (we define $f_e = |f_{uv}| = |f_{vu}|$) in a normal operation of the system (that is, $f_e \leq c_e$). The cascade proceeds in rounds. We denote by $F_i \subseteq E$ the set of edge failures in round i and by $F_i^* = F_{i-1}^* \cup F_i$ the set of edge failures until the end of round i ($i \geq 1$). $f_e(F)$ denotes the flows in $G \backslash F$. We assume that before the initial failure event $F_0 \subseteq E$, the power flows satisfy (1)-(2), and $f_e \leq c_e \ \forall \ e \in E$. We use a deterministic outage rule [1] assuming that edge $e \in E$ fails as soon as the flow exceeds its capacity ($f_e > c_e$). To avoid cases in which all edges of an induced path fail (due to path flow exceeding their capacities), we assume that only one edge in the path fails.

We present the metrics for evaluating the grid vulnerability following an initial failure:
Yield ($Y(G, F_0)$): The ratio between the demand supplied at stabilization and its original value. Accordingly,

SIGMETRICS'13, June 17-21, 2013, Pittsburgh, PA, USA.
ACM 978-1-4503-1900-3/13/06.

$Y(G,k) = \min_{F_0 \subseteq E, |F_0| \leq k} Y(G,F_0)$ is the minimum yield for any F_0 of size at most k.

Number of edge failures ($|F^*_+(G,F_0)|$): The number of edges that fail until the system stabilizes. Accordingly, $|F^*_+(G,k)| = \max_{F_0 \subseteq E, |F_0| \leq k} |F^*_+(G,F_0)|$.

Number of rounds ($L(G,F_0)$): The number of cascade rounds until the system stabilizes. Accordingly, $L(G,k) = \max_{F_0 \subseteq E, |F_0| \leq k} L(G,F_0)$, and $L(G) = L(G,|E|)$.

Distance between consecutive failures ($D(G,F_0)$): We define the *distance sequence* $(d_1, d_2, \ldots, d_t)$ associated with $(F_0, F_1, \ldots, F_t)$ as follows. For any i, $1 \leq i \leq t$, $d_i = d(F_{i-1}, F_i)$ where $d(F_{i-1}, F_i) = \min_{e \in F_{i-1}, e' \in F_i} d(e,e')$ and $d(e,e')$ is the distance between edges e and e' in G. The minimum distance between consecutive failures is $D(G,F_0) = \min_{i, 1 \leq i \leq t} d_i$. Accordingly,
$D(G,k) = \max_{F_0 \subseteq E, |F_0| \leq k} D(G,F_0)$, and $D(G) = D(G,|E|)$.

2. CASCADE PROPERTIES

The observations below demonstrate unique properties of the power flow and cascading failure models. Obs. 1 describes non-monotonicity effects of failures (i.e., a set of edge failures may have less significant impact than some of its subsets). It is shown in [1, Lemma 4.3] that a failure event F_0 may result in a lower yield than a failure event $F \supset F_0$. However, in the proof in [1], $G \backslash F$ is disconnected. Obs. 2 and 3 show that the metric values may be arbitrarily large or small even for a single edge failure event. In [1, Lemma 4.2] it was shown that cascading failures may happen within arbitrarily long distance of each other, and in [1, Lemma 4.7] it was shown that they can last arbitrarily long time. However, Obs. 3 shows that these two events can happen simultaneously. Obs. 2 and 3 are summarized in Table 1. Obs. 4 and 5 show that small changes of the parameters may have a large effect on the metric values.

Observation 1 (non-monotonicity). *There exists a graph $G = (V,E)$, an initial failure $F_0 = \{e\}$, $e \in E$, and $F'_0 \supset F_0$, such that $(V, E \setminus F'_0)$ is a connected graph, $Y(G,F_0) = 0$ and $Y(G,F'_0) = 1$.*

Observation 2 (Rounds, failures, and yield). *For any $m \geq 2$, there exists a graph $G = (V,E)$ with $|E| \geq m$, such that $L(G,1) = |E|-1$, $|F^*_+(G,1)| = |E|$, and $Y(G,1) = 0$.*

Observation 3 (Rounds and distance). *For any l, $d \geq 1$, there exists a graph $G = (V,E)$ such that $L(G,1) \geq l$ and for any i, $1 \leq i \leq l$, $d_i \geq d$. As a result $D(G,1) \geq d$.*

Define graphs $G^c_\pm$ and $G^x_\pm$ as modified versions of the graph $G = (V,E)$ with a small difference in the parameter of a particular edge $e \in E$ (in $G^c_\pm$, $c^\pm_e = c_e \pm \varepsilon$; and in $G^x_\pm$, $x^\pm_e = x_e \pm \varepsilon$).

Observation 4 (Parameters decrease). *For any $\varepsilon > 0$ and any $m \geq 2$, there exists a graph $G = (V,E)$ with $|E| \geq m$, an edge $e \in E$, and an initial failure $F_0 \subseteq E$ with $|F_0| = 1$, such that:*
$L(G,F_0) = 0$, $|F^*_+(G,F_0)| = |F_0| = 1$, $Y(G,F_0) = 1$; but
***(a)** $L(G^c_-,F_0) = |F^*_+(G^c_-,F_0)|-1 = |E|-1$, $Y(G^c_-,F_0) = 0$,*
***(b)** $L(G^x_-,F_0) = |F^*_+(G^x_-,F_0)|-1 = |E|-1$, $Y(G^x_-,F_0) = 0$.*

Observation 5 (Capacity increase). *There exists a graph $G = (V,E)$, an edge $e \in E$, and an initial failure $F_0 \subseteq E$ with $|F_0| = 1$, such that:*
$L(G,F_0) = 1$, $|F^*_+(G,F_0)| = |F_0| + 1 = 2$, $Y(G,F_0) = 1/3$; but $L(G^c_+,F_0) = 1$, $|F^*_+(G^c_+,F_0)| = 3$, $Y(G^c_+,F_0) = 0$.

Table 1: Worst case values of the metrics for cascades caused by a single edge failure.

Metric		Worst case	
Yield	$Y(G,1)$	0	Obs. 2
Number of edge failures	$\lvert F^*_+(G,1)\rvert$	$\lvert E\rvert$	Obs. 2
Number of rounds	$L(G,1)$	$\lvert E\rvert - 1$	Obs. 2
Distance between failures	$D(G,1)$	$O(\lvert E\rvert)$	Obs. 3

3. HARDNESS RESULTS

The following lemmas show that finding an initial set of failures F_0 that minimizes the *yield*, maximizes the *number of rounds*, or maximizes the *distance between consecutive failures* are all hard problems.

Lemma 1. *Given a graph G, a real number y, $0 \leq y \leq 1$, and an integer $k \geq 1$, the problem of deciding if $Y(G,k) \leq y$ is NP-complete.*

Lemma 2. *Given a graph G and an integer $t \geq 1$, the problem of deciding if $L(G) \geq t$ is NP-complete.*

Lemma 3. *Given a graph G, the problem of computing $D(G)$ is not in APX.*

4. ALGORITHMS

Inspired by circuit theory methods, we introduce the (r,h)-decomposition of the *single supply–single demand* graph G. It consists of replacing the network between arbitrary pairs of nodes by edges with equivalent reactance values such that the flow between the supply and demand nodes is preserved. We show below that using the pre-computed (r,h)-decomposition of the graph, it is possible to compute and recompute the flows with low complexity. When r and h are relatively small, this allows checking if a cascade has been initiated in a more efficient method than the classical one (which requires $O(|V|^3)$ time [5]).

Lemma 4. *Given a constant integer $r \geq 1$, the problem of deciding if $G = (V,E)$ admits an (r,h)-decomposition for some $h > 0$, can be solved in $O((|V||E|)^2)$ time.*

Lemma 5. *Given an (r,h)-decomposition of a given graph $G = (V,E)$, and F_0 of any size, $f_e(F_0)$ $\forall e \in E$, can be computed in $O(r^3|E|)$ time.*

Lemma 6. *Given an (r,h)-decomposition of a given graph $G = (V,E)$, and F_0 of any size, deciding if F_0 initiates a cascade can be done in $O(r^3 h|F_0|)$ time.*

5. ACKNOWLEDGMENTS

This work was supported by NSF grant CNS-1018379 and DTRA grants HDTRA1-09-1-0057 and HDTRA1-13-1-0021.

6. REFERENCES

[1] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman. Power grid vulnerability to geographically correlated failures - analysis and control implications. Technical Report CU-EE-2011-05-06, Columbia University, Nov. 2011.

[2] D. Bienstock. Optimal control of cascading power grid failures. *Proc. IEEE CDC-ECC*, Dec. 2011.

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

[4] 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, 2007.

[5] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins Studies in Mathematical Sciences, 4th edition, 2012.