

LA-UR- 10-07971

Approved for public release;
distribution is unlimited.

Title: Statistical Analysis of Cascading Failures in Power Grids

Author(s): Michael Chertkov, T-4, LANL
Rene Pfitzner, T-4, LANL
Konstantin Turitsyn, T-4, LANL

Intended for: Proceedings: IEEE PES (Power Energy systems) Annual
Meeting
World-wide Web: arxiv



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Statistical Analysis of Cascading Failures in Power Grids

René Pfitzner, Konstantin Turitsyn and Michael Chertkov, *Member, IEEE*

Abstract—We introduce a new microscopic model of cascading failures in transmission power grids. This model accounts for automatic response of the grid to load fluctuations that take place on the scale of minutes, when optimum power flow adjustments and load shedding controls are unavailable. We describe extreme events, caused by load fluctuations, which cause cascading failures of loads, generators and lines. Our model is quasi-static in the causal, discrete time and sequential resolution of individual failures. The model, in its simplest realization based on the Directed Current description of the power flow problem, is tested on three standard IEEE systems consisting of 30, 39 and 118 buses. Our statistical analysis suggests a straightforward classification of cascading and islanding phases in terms of the ratios between average number of removed loads, generators and links. The analysis also demonstrates sensitivity to variations in line capacities. Future research challenges in modeling and control of cascading outages over real-world power networks are discussed.

Index Terms—Power grids, Power system dynamics, Power system faults, Power system reliability

I. INTRODUCTION

The power transmission system is one of the greatest engineering achievements of the past century. The power grid system provides electricity 24 hours a day, seven days a week. However, due to its complexity and spatial extent, it is also vulnerable to failures of various sizes and significance. Extreme events, like the infamous Aug 2003 blackout [1] which left a significant part of North Eastern US without electricity, are rare but their costs to the economy and society are enormous. The significance of this subject has stimulated research in this important area, summarized in a recent review paper by Dobson et al. [2]. Well known cascading models include random network models [3], initial power flow based models with gradual load increase, maintenance and random failures [4], [5], [6], phenomenological stochastic models [7], [8], hidden-failure embedded power flow models [9], [10], and recently power flow based models accounting for adaptive control [11].

In this manuscript, we consider a power flow model of cascading failures, applicable to temporal scales shorter than the time of operator-induced Optimal Power Flow (OPF) control. (OPF, redistributing power generation, shedding load and adjusting frequency, is typically executed on the scale

R.P. is a student at the Department of Physics, University of Jena, Germany as well as at New Mexico Consortium and CNLS, Los Alamos National Laboratory, Los Alamos, NM, USA. E-mail: rene.pfitzner@uni-jena.de.

K.T. is at MIT, Department of Mechanical Engineering and at Los Alamos National Laboratory, Los Alamos, NM, USA.

M.C. is at Los Alamos National Laboratory and at New Mexico Consortium, Los Alamos, NM, USA.

of minutes.) The demand on the transmission grid grows at a pace significantly exceeding maintenance and system reinforcement. Hence the stress on the grid gradually increases, forcing the system to operate close to its capacity. In a system far from saturation (call it the "grid of yesterday") fluctuations in demand rarely led to any significant outage, not to mention a devastating cascade. On the contrary, in the system under stress (the "grid of today or tomorrow") short-term fluctuations in demand can and unfortunately will lead to significant and possibly cascading outages. Such short-term fluctuations may be caused by local penetration of the grid with renewables or, for this case, the short-term *lack* of local generation due to weather conditions etc. This consideration motivates us to analyze outages and cascades generated solely by short-term fluctuations in demand.

Our model and approach extends the line of research started in [4], [5], [6], in which the first quasi-static and microscopic (i.e. based on power flows and not on an abstract model of stress redistribution) models of cascades were considered. As in [4], [5], [6], tripping of overloaded lines is a significant part of our power flow dynamics. However, our approach is different in what causes the tripping and how it occurs. Furthermore, in contrast to [4], [5] our analysis concentrates solely on the short-time scale. To allow for a more physical causal resolution of the actual dynamics of the grid, we do not trip multiple overloaded lines at once. We also consider a somewhat more basic model of cascade initiation, caused only by fluctuations of loads. In our model of a grid close to saturation lines are tripped solely in response to excessive demand. This differs from previous approaches, where pure random line outages [5], effects of sympathetic line trippings [6] or so called hidden failures [10] were taken into account. (In a more realistic scenario, initiations of multiple types may cause cascades. See Section IV.)

This study is inspired by the actual operational paradigm guiding the grid dynamics on the scale of minutes or even seconds, when an intelligent manual or semi-automatic control (typically including a human decision in the loop) is simply unavailable. We only account for standard, automatic and system inherent control, such as voltage control, automatic line tripping and droop control executed at the generators. The highlights of our method and techniques in this manuscript are:

- We study cascading behavior by solving the power flow equations. Unlike the phenomenological ("disease spread"-like) models, our model accounts for *non-local, non-uniform* power-flow based dynamic responses of the power system.
- We propose a realistic cascade algorithm executed over

non-uniform IEEE test beds, and not over abstract (tree-like or random) structures. This results in a realistic spatio-temporal redistribution of overloads. We point out the importance of this course of action since power systems are not arbitrary but designed intelligently (e.g. capable to withstand any $N - 1$ contingency).

- Our cascade algorithm relies heavily on inhomogeneous line overload and tripping, thus resulting in a dynamical change of the network structure and (in the case of significant damage) leading to the emergence of islands. To study these cascades, we then solve the power flow problem on every island independently. Islanding may also result in an extreme form of load shedding, i.e. blackout of the entire island if the cumulative load exceeds the generation capacity.
- The only source of randomness in our model is associated with fluctuations in demand. We show that these demand fluctuations, without other random external influences, are sufficient for causing cascades.
- The outage growth is quantified in terms of the average number of tripped elements (loads, generators and lines) and their dependence on the magnitude of demand fluctuations.

This manuscript is organized in four Sections. Our cascading algorithm is described in Section II. The algorithm requires solving the power flow equations many times, in a way that mimics the actual dynamics of the grid in a quasi-static causal fashion. We discuss the application of our cascading algorithm to three IEEE systems, consisting of 30, 39 and 118 buses respectively, in Section III. The last Section of the manuscript is reserved for Conclusions and discussing the Path Forward.

II. CASCADING ALGORITHM

Our algorithm models cascades triggered by demand fluctuations around the base OPF solution. We consider a quasi-static ("sequence of steady states") model. This algorithm takes an instance of demand for the input and it outputs a balanced solution, possibly over a sub-grid of the original grid. A flow chart of the general structure of the algorithm is shown in Fig. 1. The algorithm is rather involved. Thus we find it useful to begin by discussing its high-level description first in the main body, followed by more detailed expositions of the algorithm sub-tasks in their respective Subsections. (The order of Subsections in this Section is dictated by convenience of the presentation as well as by the interdependence of the material. It does not necessarily reflect the importance of the sub-tasks or their order in execution of the algorithm.)

The algorithm begins by computing the OPF solution $g^0 = (g_i | i \in \mathcal{G}_g)$, using $d^0 = (d_i | i \in \mathcal{G}_d)$ as the base/reference point. Here, $\mathcal{G} = (\mathcal{G}_0, \mathcal{G}_1)$ is the graph of the power grid consisting of a set of vertexes, \mathcal{G}_0 , and edges, \mathcal{G}_1 . $\mathcal{G}_g \subset \mathcal{G}_0$, $\mathcal{G}_d \subset \mathcal{G}_0$ is the subset of nodes with generators and demands respectively. Then, the OPF solution is perturbed by a random demand fluctuation, $\delta = (\delta_i | i \in \mathcal{G}_d)$, drawn from a distribution parameterized by its root-mean-square deviation from the mean, d^0 . (A detailed discussion of the probability distribution function of the load distribution considered in our simulations

can be found in the beginning of Section III-A.) We apply demand perturbations in a step-wise fashion, with one step corresponding to one cycle of the outer loop in the chart diagram of Fig. 1. The number of steps varies with the perturbation and depends on the severity (number of consecutive violations of line or generator constraints which need to be resolved) in the fluctuation of demand. Each step (loop) consists of the following mini-steps: (a) sequential evaluation of the time rescaling parameter t^* ; (b) redistribution of new loads solely according to droop control (remind that we do not include in this study any control of demand); (c) calculation of Power Flow (PF) solution for the resulting rescaled configuration of demands $d^0 + t^* \delta$, laying (in the multi-dimensional space of demands) strictly in between the reference point d^0 and the investigated configuration $d^0 + \delta$; (d) line check and possible tripping of one violated line based on the DC power flow solution (note that this sub-step involves some randomness in selecting the tripped edge of possibly many violated ones); and (e) checking if the tripping resulted in any islanding. These five mini-steps of the loop are described in detail in Sections II-E, II-D, II-A, II-F and II-C respectively.

It is important to stress at this point that the relatively involved structure of the algorithm mimics the actual microscopic dynamics of the power grid control/adjustment to the demand change. In particular, the gradual modification in demand from the reference point to the final contingency, modeled with t^* increasing monotonically from zero to one in a finite number of steps, reflects the physics of the causal response of the grid. Here we assume that the generator-based control of the power flow takes place on the scale of milliseconds, while any change in demand and resulting tripping events occur at much slower pace (measured in seconds, or even longer intervals). Thus the grid has enough time to respond to each individual contingency establishing a quasi-static equilibrium, modeled in our algorithm by the power flow solver and by the check for islanding steps, after each of the elementary line-tripping or generator-saturation events.

A. DC power flow

The power solver takes injection and consumption of powers at all the nodes of the power grid as well as system parameters as input and it outputs voltages and phases at the nodes, as well as the power transmitted over all the links of the grid. Generally AC power flow is a nonlinear algebraic problem, that becomes a linear problem under a set of additional, so-called Directed Current (DC), assumptions. Our cascading algorithm will work with the most general power solver. However, in this paper we choose to work with DC solver which is simpler in implementation and does not require special algorithms for distributing the losses between the generators. The DC solver evaluates

$$\forall i \in \mathcal{G}_0 : \sum_{j \sim i} p_{ij} = \begin{cases} g_i, & i \in \mathcal{G}_g \\ -d_i, & i \in \mathcal{G}_d \\ 0, & i \in \mathcal{G}_0 \setminus (\mathcal{G}_g \cup \mathcal{G}_d) \end{cases} \quad (1)$$

$$\forall \{i, j\} \in \mathcal{G}_1 : \theta_i - \theta_j = x_{ij} p_{ij} \quad (2)$$

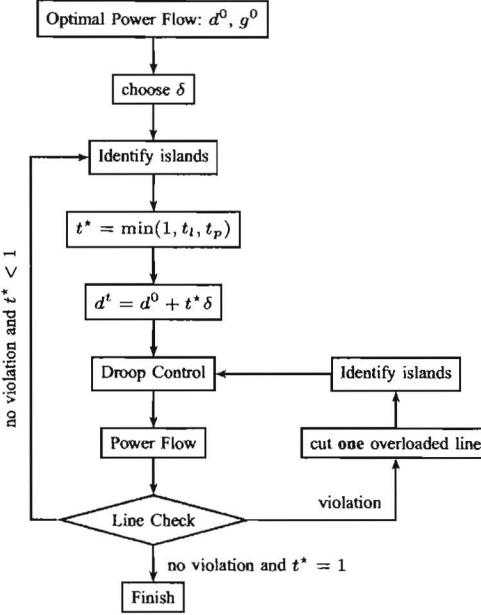


Fig. 1. Flowchart of the proposed cascading algorithm

where $x = (x_{ij} | \{i, j\} \in \mathcal{G}_1)$, $g = (g_i | i \in \mathcal{G}_g)$, $d = (d_i | i \in \mathcal{G}_d)$, $\theta = (\theta_i | i \in \mathcal{G}_0)$, $p = (p_{ij} = -p_{ji} | \{i, j\} \in \mathcal{G}_1)$ are the vector of line inductances, the vector of powers injected at generators, the vector of demands consumed at loads, the vector of phases and vector of line flows, respectively. (Here $\{i, j\}$ is our notation for directed edges and $j \sim i$ indicates that j is the graph neighbor of i .) Note that to streamline notations we used a somehow abbreviated version of the DC power flow equations in (1,2). In particular, we ignore terms associated with tap transformers. In our simulations, we utilize the DC PF solver from the Matlab based MATPOWER package [12] taking into account effects of transformers and other devices included in description of the 30, 39 and 118 nodes IEEE systems.

B. Optimal Power Flow

To set up the system, we solve the standard DC optimal power flow problem finding optimum generator dispatch given the initial load d^0 , cost functions $f = (f_i | i \in \mathcal{G}_g)$ for every generator as well as generation power and line capacity constraints. To execute this task we use MATPOWER [12], and cost functions provided in the description of the IEEE systems studied. The DC optimal power flow, in the simplest nomenclature, corresponds to solving

$$\min_{p, g, \theta} \sum_i f_i(g_i) \quad \left| \begin{array}{l} \text{Eqs. (1,2), where } d \rightarrow d^0 \text{ hold} \\ \forall \{i, j\} : |p_{ij}| \leq p_{ij}^{\max} \\ \forall i : g_i^{\min} \leq g_i \leq g_i^{\max} \end{array} \right. \quad (3)$$

for the branch flows, p , and generation powers, g . The resulting p^0 , g^0 and θ^0 form the base (reference) solution for our cascading algorithm.

C. Identify islands

Our algorithm does not generate a surviving balanced subgrid at once, but instead it resolves it in steps, thus mimicking dynamics of realistic cascades. The temporal evolution of the surviving sub-grid is induced by cutting saturated lines, which might also cause the formation of islands, and removing freshly formed but overloaded islands. We check for islanding (i.e. splitting of the grid into independent components) using the function *grComp* of the MATLAB *grTheory* toolbox [13].

If an island is formed, we do all other computations within the cascading algorithm (including DC Power Flow, droop control and calculation of t^*) independently for every island.

D. Droop control

In the process of evaluating the cascading algorithm, it can happen, due to tripping of overloaded lines, that some loads or generators will become disconnected from the grid or that the grid splits up into islands. Both scenarios require automatic redistribution of generation, done in the so-called droop control fashion.

Droop control is executed at each generator locally in response to an increase or decrease of the system frequency (measured locally as well). If a generator experiences a drop in frequency, which is a symptom of over-consumption and under-generation, it increases its generation power in order to undo the frequency drop. Similarly, if the frequency increases, this is interpreted as a signature of over-generation. Since these deviations are typically small, we stick in this study to a linear response model, assuming a linear relation between the frequency drop/gain and the generation increase/decrease. (Nonlinear modifications can be easily included in a more general scheme.) We do droop control if: (a) The grid changes its structure, i.e. following appearance of new island(s) in the result of line tripping. (b) The demand at any node on the island has changed. Here we assume that the power generation, $g_i(+)$, at node i after at least one of these events is

$$g_i(+) = \frac{g_i(-)}{g_{\Sigma}(-)} d_{\Sigma}(+), \quad (4)$$

where the newly introduced quantities on the right hand side of Eq. (4) are the current power generation, $g_i(-)$, at node i ; the total power generation (before droop control), $g_{\Sigma}(-) = \sum_{j \in \Sigma} g_j(-)$, at the freshly formed island, $\Sigma \subset \mathcal{G}$, the generator belongs to; and the total power demand, $d_{\Sigma}(+) = \sum_{j \in \Sigma} d_j(+)$, of the island observed after event a) or b). Note that if neither case a) nor case b) occurred at the island Σ' , $d_{\Sigma'}(+) = d_{\Sigma'}(-)$ will hold. Since we always make sure that a stable well-balanced solution demand and generation match, $d_{\Sigma'}(-) = g_{\Sigma'}(-)$ should hold and we arrive at $g_i(+) = g_i(-)$ in the result.

Droop control is executed at all the generators of the grid simultaneously. Note that the ratio on the rhs of Eq. (4) changes in the process of our discrete event simulations in accordance with the modification of islands. If at some point in the process a generator becomes saturated, we do not include it anymore in the droop control mechanism described above, but instead keep its generation level constant (at the

maximum generation capacity). Our droop control mechanism also includes accounting for overloads - as long as, using droop control, demand and total power generation can be matched, the island is functioning. However, if the total demand of the island exceeds its total generation capacity, we shut down the entire island. Thus, the transition point from phase two to phase three shown in Fig. 4 is associated with the emergence of a statistically significant number of islands which were shut down.

E. Discrete Time Evolution of Loads

Once the configuration of loads, $d = d^0 + \delta$, is generated we do not increase the demand from the reference point, d^0 , at once but instead break the change into a number of incremental steps. Each of the steps is associated with only a single element modification of the underlying grid (breaking of a single line constraint or an individual generator reaching its capacity). This incremental degradation allows us to model more realistically causal advancement of the stress mounting within the grid. For this discrete event procedure to work, we need to generate a monotonically increasing sequence of (fraction) times $t^* \in (0, 1]$ of load disturbances, each associated with the $d^0 + t^* \delta$ configuration of loads within the island. (We recall that if the grid is islanded, we calculate the discrete time sequence separately for each island and choose the minimum time over all islands as t^* .) One accounts for constraint breaking events of the following two types:

- a) Exceeding the local maximum generation power constraint at time $t^* = t_g$.
- b) Exceeding the line capacity constraint at time $t^* = t_l$.

The time of locally exceeding generation capacity is

$$t_g = \min_i \left(\frac{g_i^{\max} - g_i(-)}{\frac{g_i(-)}{g_{\Sigma}(-)} \delta_{\Sigma}} \right). \quad (5)$$

Indeed, t_g is the time when $g_i(+, t_g) \geq g_i^{\max}$ holds for exactly one special generator site $i \in \Sigma$. Then, the expression for the post-event generation at the special site is

$$g_i(+, t) = \frac{g_i(-)}{g_{\Sigma}(-)} d_{\Sigma}(+, t) \quad (6)$$

$$= \frac{g_i(-)}{g_{\Sigma}(-)} [d_{\Sigma}(-) + t \delta_{\Sigma}] \quad (7)$$

$$= g_i(-) + t \frac{g_i(-)}{g_{\Sigma}(-)} \delta_{\Sigma}, \quad (8)$$

where $d_{\Sigma}(+, t) = d_{\Sigma}(-) + t \delta_{\Sigma}$ is the post-event cumulative demand over the island Σ , and Eq. (5) follows from $g_i(+, t_g) = g_i^{\max}$.

Analogously, the time of exceeding a line constraint is

$$t_l = \min_{(ij)} \left(\frac{p_{ij}^{\max} - p_{ij}(-)}{PF(\mathcal{G}, \delta, \Delta g)} \right), \quad (9)$$

which follows from the following consideration. Let $p(\pm) = PF(\mathcal{G}(\pm), d(\pm), g(\pm))$ denote the vector of power flows, p , over the transmission lines of the grid $\mathcal{G}(\pm)$, obtained by solving Eqs. (1,2) with demand $d(\pm)$ and generation $g(\pm)$, where \pm indicates (as before) relevance to the pre- and post-

droop control state. One derives that, $p(+, t)PF(\mathcal{G}(-), d(-) + t\delta, g(-) + t\Delta g)$, where $(\Delta g)_i$ is defined according to the droop control rule, Eq. (8). Furthermore, since $PF(\dots)$ is linear in d and g , one finds that

$$p(+, t) = p(-) + t \cdot PF(\mathcal{G}(-), \delta, \Delta g), \quad (10)$$

thus arriving, under condition that $p_{ij}(+, t_l) = p_{ij}^{\max}$ holds for exactly one line, at Eq. (9).

F. Line checking and tripping

After solving the DC PF equations we check for lines with violated constraints. The constraints are stated in the two last lines of the conditions for OPF in Eq. (3). It can easily happen, that one encounters degeneracy, in the sense that the constraints are violated at more than one line. If this is the case, we do not trip all the lines with violated constraints at once, but instead exclude only one of them and then do droop control and DC power flow again. (See the small loop in the flowchart Fig. 1.) The order of exclusion is chosen randomly prior to executing the cascading algorithm and it is maintained the same over all iterations and samplings. We stress that this degeneracy can only happen if at a previous step $t_g < t_l$ held, which implies that we change local power generations *instantly*. This course of actions is an approximation and it would of course be more physical to not increase the generation power instantly, but gradually, and to account for singular line tripping events during this gradual increase. However, since dynamic generator data is very often not known, we choose to follow this simple and reasonable scheme.

Here we want to mention, that our decision rule of whether a line is going to be tripped or not is obviously "magnitude only". This differs from other power-modeling approaches, where not only the magnitude of overload of a line (exceeding its rating) is important, but only the duration of overload. Indeed, from a power systems operator point of view the acceptable magnitude of overload is generally dependent on the duration of overload. However, since we study a quasi-static model, we neglect the temporal aspect.

From the quantitative analysis point of view, we only count a line as tripped when at some point the line flow exceeds its capacity. This means that if an island is overloaded and thus removed from the grid, we do not count lines within the removed island (which now do not carry any current) as tripped. After all, these lines show no need for maintenance. This is different from the way we account for tripped generators and demands. When an island shuts down (then cumulative load exceeds generation capacity), we count these unserved demands and the generators set off-line as tripped.

III. RESULTS

In this Section we report the results of testing our cascading algorithm, described in details in Section II, on three IEEE systems.

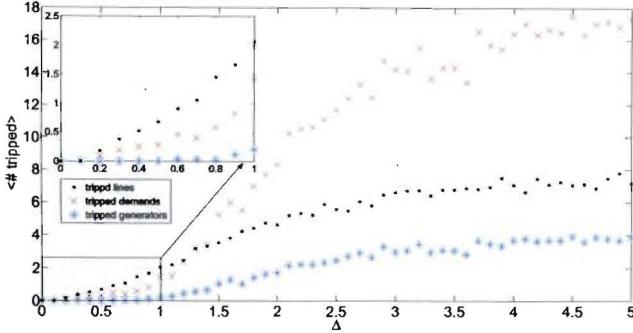


Fig. 2. Average characteristics of outages in the 30 bus network induced by fluctuations in demand explained in the text. Every data point presents the average over 200 i.i.d. samples using the specified distribution. $d_{\Sigma}^{(0)} = 0.565g_{\Sigma}^{(\max)}$. We observe three transition points. Prior to reaching the first transition point the grid is resilient to fluctuations in demand. In between the first and second transition points the probability of having an outage increases slowly. It turns out that the outage in this regime corresponds to tripping of two lines (due to overload) followed by islanding of the adjusted demand node. Here, one observes no cascades yet but only increased probability of line and demand tripping (as also witnessed by the low slope of tripped demands and the stress diagram of the system shown in Fig. 5). Passing the second transition point indicates emergence of a macroscopically significant number of tripped generators, which also results in a faster rise of demand tripping and signifies the start of cascades. In this system we also note a third transition point at which the number of tripped demands exceeds the number of tripped lines, thus indicating that significant number of the unserved demands belong to islands left without power.

A. Effect of random demand distributions

In this Subsection we report tests on standard IEEE systems with 30 and 39 buses respectively. For these two systems all the important system parameters (maximum generation power, g_i^{\max} , maximum line capacities, average demand distribution, $d^{(0)}$) are available in the system specifications documented in [12]. We select the average load according to, $d_G = 0.565g_G^{\max}$, while maintaining the same relative distribution of demands between load buses, as specified in [12]. (The re-scaling reflects a typical day scenario for the reference point.) Then we set the distribution of generation according to the optimal power flow solution of Eq. (3). Fluctuations in the demands are generated using the half normal distribution allowing only positive fluctuations, in demand $\forall i \in \mathcal{G}_0$:

$$\mathcal{P}(\delta_i) = \begin{cases} \frac{\exp(-(\delta_i)^2/(2d_i^0\Delta))}{\sqrt{\pi d_i^0\Delta/2}}, & d_i^0 + \delta_i > d_i^0 \\ 1/2, & d_i^0 + \delta_i = d_i^0 \\ 0, & d_i^0 + \delta_i < d_i^0 \end{cases} \quad (11)$$

Figs. 2 and 3 show average characteristics of outages, caused by the distribution of demand Eq. (11) under different values of dispersion, for the systems of 30 and 39 buses respectively. Although the qualitative forms of the curves observed are quite similar, one also finds interesting differences in the way how cascades evolve in the two systems. Whereas in the 30-bus system the first phase is associated with line tripping also resulting in isolation of a few demands, this state is virtually absent in the 39-bus system. The cascading behavior in this system begins with the second phase. The increase in the number of unserved demands is here in fact induced by generator tripping. Furthermore phase two and

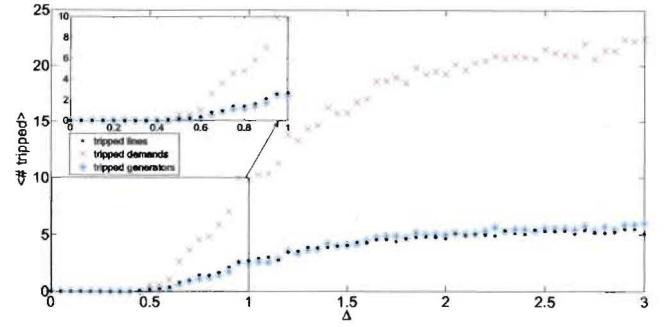


Fig. 3. Outages in the 39bus network induced by disorder in the demands. Every data point presents the average over 200 samples of Eq. (11). $d_{\Sigma}^{(0)} = 0.565g_{\Sigma}^{(\max)}$.

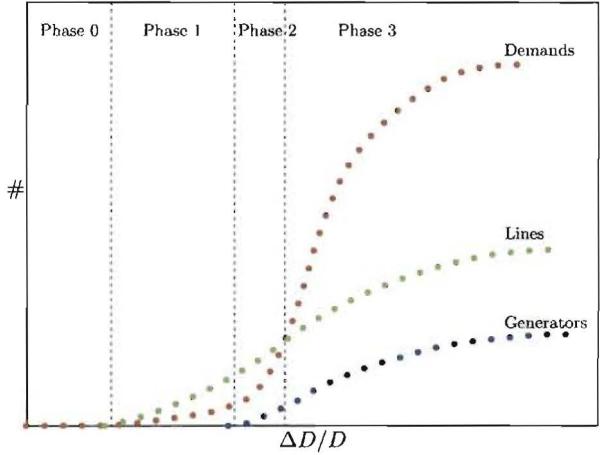


Fig. 4. Illustration of possible cascade phases. In phase 0 the grid is resilient against fluctuations in demand. Phase 1 shows tripping of demands due to tripping of overloaded lines. This has an overall "de-stressing" effect on the grid. In phase 2 generator nodes start to become tripped, mainly due to islanding of individual generators. With the early tripping of generators the system becomes stressed and cascade evolves much faster (with increase in the level of demand fluctuations) when compared with a relatively modest increase observed in phase 1. Outages in phase 3 are associated with removal from the grid of complex islands, containing both generators and demands.

phase three almost coincide in the 39-bus system. However, as the stress-diagrams (see Fig. 6 and Fig. 5) of both systems show, islanding is an important effect. The improved behavior of phase one in the 39-bus system appears to be due to the fact that created islands are (a) rather big, (b) include generators, and (c) are stressed and becomes powerless almost immediately after emergence. In phase one of the first system, islands are small and immediately removed from the grid (as not containing generators). Somehow surprisingly, this early removal of many small islands has a positive, de-stressing effect on the remaining part of the grid. Figure 4 provides a qualitative illustration of the three phases. The stress diagrams, shown in Fig. 6 and Fig. 5, provide additional evidence supporting the aforementioned explanations.

A qualitative similar behavior to the one shown in Figs. 2 and 3 was also observed for some other choices of demand distributions Eq. (11), e.g. for fluctuations allowing a decrease

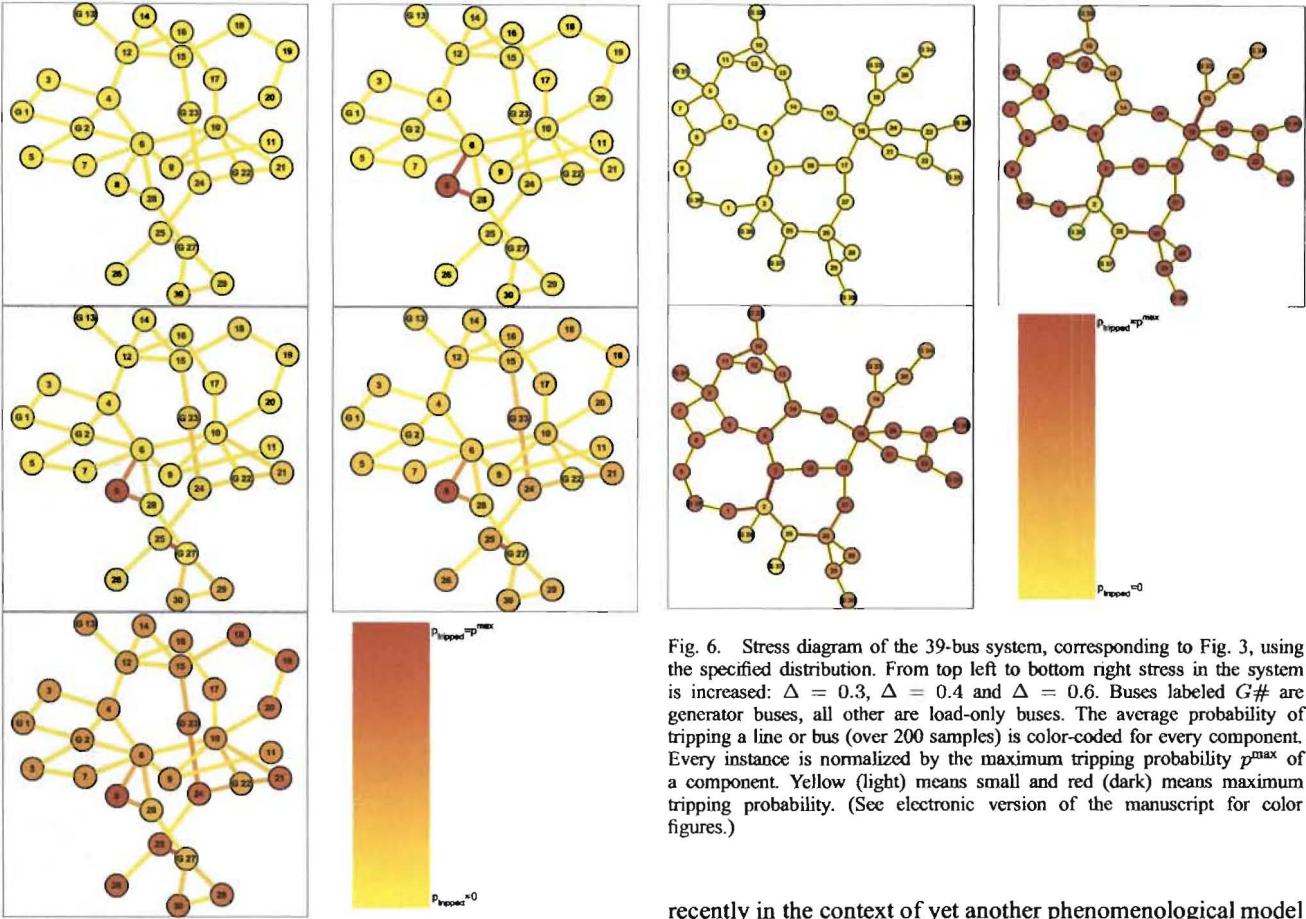


Fig. 5. Stress diagram of the 30-bus system, corresponding to Fig. 2, using the specified distribution. From top left to bottom right stress in the system is increased: $\Delta = 0.1$, $\Delta = 0.2$, $\Delta = 0.9$, $\Delta = 1.2$ and $\Delta = 2.0$. Buses labeled $G\#$ are generator buses, all other are load-only buses. The average probability of tripping a line or a bus (over 200 samples) is color-coded for every component. Every instance is normalized by the maximum tripping probability p^{\max} of a component. Yellow (light) means small and red (dark) means maximum tripping probability. (See electronic version of the manuscript for color figures.)

in demands (yet keeping the total demand positive). Hence one can speculate that the qualitative picture of outage growth induced by increase in load fluctuations is universal.

At this point we would like to comment on one interesting similarity between our microscopic results and the results of Dobson et al. reported for the phenomenological CASCADE model [7]. The CASCADE model is an abstract representation of the power grid, considering equivalent components failing according to some pre-defined distribution. If a component fails, a certain amount of load is distributed equally to all other components. The model is structureless and as such it carries no explicit relation to the power flow equations. It is argued in [7] that an increase in the parameter mimicking increase in the total load, results in an abrupt increase in the size of the damage starting at some finite threshold value of the parameter. This observation is akin to the transition from phase one to phase two (see Fig. 4) observed in our microscopic model. Moreover, a similar observation was made

Fig. 6. Stress diagram of the 39-bus system, corresponding to Fig. 3, using the specified distribution. From top left to bottom right stress in the system is increased: $\Delta = 0.3$, $\Delta = 0.4$ and $\Delta = 0.6$. Buses labeled $G\#$ are generator buses, all other are load-only buses. The average probability of tripping a line or a bus (over 200 samples) is color-coded for every component. Every instance is normalized by the maximum tripping probability p^{\max} of a component. Yellow (light) means small and red (dark) means maximum tripping probability. (See electronic version of the manuscript for color figures.)

recently in the context of yet another phenomenological model discussed in [14]. In contrast to the model of [7], the one of [14] is based on a network with some spatial structure. Triggered by one random overload (failure of a generator), cascading behavior is observed as the overload is equally distributed among generators, which are graph neighbors of the failed one. Here as well, a phase transition was observed. Overall, we conclude that transition from phase one (in which almost all generators are functional) to phase two (where a significant, $O(1)$, fraction of generators is in outage) is observed universally across the models.

B. Effects of line capacities

In this Section we describe results of our cascading algorithm test on the larger IEEE 118-bus system. In the contrast with the 30-bus and 39-bus system, specification of the 118-bus system available in MATPOWER does not have line capacities. To resolve this problem, we generated line capacities ourselves. Somehow surprisingly, these experiments lead us to the conclusion that variability in the line capacities affects the dynamics of cascades in a strong way. More precisely, we observed that the distribution of line capacities influences the structure of emerging islands.

This observation, made first on the 118-bus system, led us to experiment with the variability of capacities in the smaller systems. We show in Figs. 8 how the cascading behavior is influenced by different distributions of line capacities in the 30-bus system with 6 generation and 24 demand nodes.

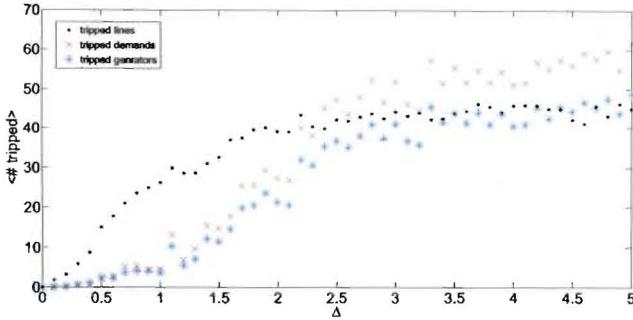


Fig. 7. Outages in the IEEE 118-bus network induced by disorder in the demands. Every data point presents the average over 25 i.i.d. samples using the specified distribution. $d_{\Sigma}^{(0)} = 0.565g_{\Sigma}^{(max)}$. Since for this systems the IEEE standard does not specify line capacities, we assign line capacities randomly according to the relative-capacity-distribution from the 30-bus case.

We consider three univalued cases, when all the capacities are set to the same value, equal to the smallest, largest and mean characteristics of the original distribution (available in the MATPOWER specification for the system). When the univalued line capacity is maximal (top left figure), no line outage or islanding is observed. Outages in such systems are solely due to generators exceeding their capacities, which obviously becomes more likely with increase in Δ . Referring to our classification scheme, phase 1 and 2 are absent in this case. This observation is also consistent with results obtained earlier in [7] or [14]. In contrast, when the univalued line capacity is minimal (top right figure), the cascading behavior is not seen. In this case, any (sensible) initial distribution leads to islanding and blackout of some of these islands. (Note that in this case the complete grid is never outaged completely, instead we reach a stable point with 4 loads and 2 generators remaining for a rather wide variability range in Δ . We attribute this peculiar result to the specific topology of the grid and the initial demand distribution, d^0 . The islands are formed in a way that two remaining generators provide powers to the remaining loads without exceeding line capacities.) In the univalued case corresponding to the mean value of the original distribution, illustrated in the bottom left of Fig. 8, phases 1 and 2 are missing again, suggesting that islanding did not lead to any relief (de-stressing). Looking at all these three univalued examples from a quantitative perspective, we observe that the average size of the outage at the maximum fluctuations in the demand considered, $\Delta = 5$, is significantly larger than in the original case of a realistic distribution of capacities. We associate this negative effect of the univalued capacity with the lack of heterogeneity in the islands formed under stress. To conclude, we find that setting the line capacities to the same value (large, small or averaged) leads to overestimation of the strength of the outage in comparison with the case of a more realistic distribution of capacities.

Therefore, to generate a realistic study of the 118-bus system, we distributed line capacities according to the relative-capacity-distribution of the 30-bus system, i.e. according to

$$\frac{p_{ij}^{max}}{p_{ij}^0} (30 \text{ bus}) \sim \text{Dist}_{ij}(30 \text{ bus}). \quad (12)$$

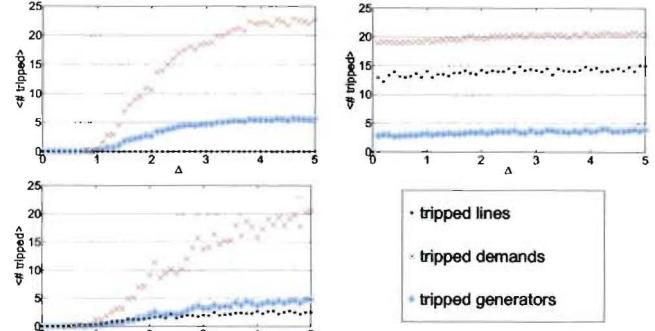


Fig. 8. Outages in the IEEE-30bus network with univalued line capacities induced by fluctuations in demands. Every data point is a result of averaging over 200 i.i.d. samples from Eq. (11). $d_{\Sigma}^{(0)} = 0.565g_{\Sigma}^{(max)}$. Top left: Set all line capacities equal to the maximum line capacity from the original distribution. Top right: Set all line capacities equal to the minimum line capacity of the original distribution. Bottom left: Set all line capacities equal to the mean of the original distribution.

We obtain the maximum possible power flow (line capacity) over line α in the system as

$$p_{ij}^{max}(118 \text{ bus}) \sim \text{Dist}_{ij}(30 \text{ bus}) \cdot p_{ij}^0(118 \text{ bus}). \quad (13)$$

Fig. 7 shows the resulting outage diagram observed in this synthetic system. These results are consistent with the simulations of the smaller systems and also with the qualitative scheme described in Fig. 4. Note, that the same simulations conducted with univalued capacities (not shown) give a significantly different picture, qualitatively consistent with the results reported in Fig. 8 for univalued capacity tests in smaller systems.

Based on the results of this Subsection, we conclude that in order to capture realistically cascading effects, like islanding, it is crucial to take the non-uniformity of line capacities into account.

IV. DISCUSSION AND CONCLUSION

In this manuscript we, proposed a new microscopic model of cascades in power grid. The model was tested on three IEEE systems. We solved the power flow dynamics (in DC the approximation), analyzed structural evolution of the operational part of the grid associated with islanding, and observed the emergence of cascades caused solely by fluctuations in loads. Analyzing the statistics of the damage, we identified four distinct phases, observed in response to variations in the level of demand fluctuations. Phase #0, described by small demand fluctuations, does not lead to any significant damage. Phase #1, described by modest fluctuations, results in the removal of some number of lines as well as the removal of a few loads (the formation of islands which do not have any generation and are thus blacked out immediately), while generators remain largely unaffected. This phase can have a de-stressing effect on the grid seen in a significant reduction of damage increase with increasing demand fluctuations. Phase #2, described by sizable fluctuations, is characterized by the appearance of some tripped generators (surrounded by tripped lines), while the relations, $\langle \# \text{ of tripped lines} \rangle > \langle \# \text{ of tripped loads} \rangle > \langle \# \text{ of tripped generators} \rangle$, remain

valid. Finally, phase #3, described by large demand fluctuations, is characterized by multiple islands, of which a sizable $O(1)$ portion is outaged: $\langle \# \text{ of tripped loads} \rangle > \langle \# \text{ of tripped lines} \rangle > \langle \# \text{ of tripped generators} \rangle$. Typical instances, contributing to phases #1-#3, develop in multiple sequential steps, and as such can all be interpreted as cascades of severity increasing with the numerical index of the phase.

Analyzing the dependence of the phase structure on line capacities, we observed that selecting line capacities with sufficient variability over the grid is important for capturing realistic dynamics of outages. In particular, our simulations suggest that introducing sufficient variability in line capacities (expressing realm of existing power grids) reinforces the grid, creating multiple islands, and thus making the resulting grid more resistant to a correlated large scale blackout. One observes that a cascade model, which does not account for variations in line capacities, would overestimate the damage.

All in all our study suggests that to describe dynamics and statistics of outages in the power grid faithfully, one must account for (a) fluctuations (and eventual increases) in demand (leading locally and globally to exceeding the generation capacities) and (b) islanding influenced by the distribution of line capacities.

Obviously, this study constitutes only the beginning of a strategy for analyzing power grid cascades. One natural extension would be to replace the DC power flow solver by a more realistic AC solver. We also intend to study mixed models combining the effects of demand fluctuations with effects of incidental line tripping. Then, with an eye toward aiding efforts in grid reinforcement, we plan to continue our analysis of the effect of capacity inhomogeneities on islanding. Finally, our long-term goal is to build a novel phenomenological model and theory of cascades based on a detailed microscopic analysis of the type discussed in this manuscript.

ACKNOWLEDGMENT

We are thankful to all the participants of the "Optimization and Control for Smart Grids" LDRD DR project at Los Alamos and the Smart Grid Seminar Series at CNLS/LANL for multiple fruitful discussions and to Prof. Daniel Bienstock for sharing results of his recent research [11] prior to publication. Research at LANL was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. DE C52-06NA25396. RP and MC acknowledge partial support of NMC via NSF collaborative grant CCF-0829945. The work of MC on this project was also partially supported by a DTRA basic research grant under Topic 08-D BRCALL08-Per3-D-1-0026.

REFERENCES

- [1] N. A. E. R. Council. A review of system operations leading up to the blackout of august 14, 2003. http://www.nerc.com/docs/docs/blackout/Operations_Report_FINAL.pdf.
- [2] I. Dobson, B. A. Carreras, V. E. Lynch, and D. E. Newman, "Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 17, no. 2, pp. 026103-13, 06 2007. [Online]. Available: <http://link.aip.org/link/?CHA/17/026103/1>
- [3] D. J. Watts, "A simple model of global cascades on random networks," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 99, no. 9, pp. 5766-5771, 04 2002. [Online]. Available: <http://www.pnas.org/content/99/9/5766.abstract>
- [4] I. Dobson, B. Carreras, V. Lynch, and D. Newman, "An initial model for complex dynamics in electric power system blackouts," in *Proceedings of the 34th Annual Hawaii International Conference on System Sciences (HICSS-34)-Volume 2*. Washington, DC, USA: IEEE Computer Society, 2001, pp. 2017-. [Online]. Available: <http://portal.acm.org/citation.cfm?id=820521.820527>
- [5] B. A. Carreras, V. E. Lynch, I. Dobson, and D. E. Newman, "Complex dynamics of blackouts in power transmission systems," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 14, no. 3, pp. 643-652, 09 2004. [Online]. Available: <http://link.aip.org/link/?CHA/14/643/1>
- [6] D. Nedic, I. Dobson, D. Kirschen, B. Carreras, and V. Lynch, "Criticality in a cascading failure blackout model," *International Journal of Electrical Power & Energy Systems*, vol. 28, no. 9, pp. 627-633, 2006.
- [7] I. Dobson, B. A. Carreras, and D. E. Newman, "A loading-dependent model of probabilistic cascading failure," *Probability in the Engineering and Informational Sciences*, vol. 19, no. 01, pp. 15-32, 2005.
- [8] —, "A branching process approximation to cascading load-dependent system failure," *Hawaii International Conference on System Sciences*, vol. 2, p. 20055c, 2004.
- [9] J. Thorp, A. Phadke, S. Horowitz, and S. Tamronglak, "Anatomy of power system disturbances: importance sampling," *International Journal of Electrical Power & Energy Systems*, vol. 20, no. 2, pp. 147 - 152, 1998. [Online]. Available: <http://www.sciencedirect.com/science/article/B6V2T-3SYS134-8/2/239d800c5cf1b8e4f95958b9d76ae9ec>
- [10] J. Chen, J. S. Thorp, and I. Dobson, "Cascading dynamics and mitigation assessment in power system disturbances via a hidden failure model," *International Journal of Electrical Power & Energy Systems*, vol. 27, no. 4, pp. 318-326, 5 2005. [Online]. Available: <http://www.sciencedirect.com/science/article/B6V2T-4FMBK9S-1/2/184816c8e98176b895b74453d7245724>
- [11] D. Bienstock, "Optimal adaptive control of cascading power grid failures," 2010.
- [12] R. D. Zimmerman, C. E. Murillo-Sanchez, and R. J. Thomas, "Matpower: Steady-state operations, planning, and analysis tools for power systems research and education," *Power Systems, IEEE Transactions on*, vol. PP, no. 99, pp. 1 -8, 2010.
- [13] S. Iglin, "grtheory - graph theory toolbox for matlab," available at Matlab Central.
- [14] S. Kadloor and N. Santhi, "Understanding cascading failures in power grids," <http://arxiv.org/abs/1011.4098v1>, 2010.

René Pfitzner is pursuing a Diploma (Masters) degree in Theoretical Physics (Dipl. Phys.) at Friedrich-Schiller-University Jena, Germany. He is a Diploma-student with M. Chertkov, working on several issues of non-equilibrium statistical physics, at New Mexico Consortium and the Center for Nonlinear Studies at Los Alamos National Laboratory, Los Alamos, NM, USA.

Konstantin Turitsyn is an Assistant Professor at MIT, Department of Mechanical Engineering. He received his PhD degree in Theoretical Physics from Landau Institute, Moscow in 2007. Before joining MIT he worked as a postdoctoral scholar at the James Franck Institute, University of Chicago (2007-2009) and as a J.R. Oppenheimer Fellow at Los Alamos National Laboratory (2009-2010). His scientific interests encompass a broad range of non-equilibrium statistical phenomena that occur in natural and artificial complex systems.

Michael Chertkov areas of interest include statistical and mathematical physics applied to information theory, computer science, hydrodynamics, optics, communication and infrastructure networks. Dr. Chertkov received his Ph.D. in physics from the Weizmann Institute of Science in 1996 and then spent three years at Princeton University as a R.H. Dicke Fellow at the Department of Physics. He joined Los Alamos National Lab in 1999, initially as a J.R. Oppenheimer Fellow in the Theoretical Division. He is now a technical staff member in the same division. Dr. Chertkov has published more than 100 papers in these research areas and is currently leading "Physics of Algorithms" and "Optimization and Control Theory for Smart Grids" projects at LANL.