2015 48th Hawaii International Conference on System Sciences

Lambda-gaga: Toward a New Metric for the Complex System State of the Electrical Grid

B. A. Carreras
Depart. Fisica
Universidad Carlos III
Madrid, Spain
bacarreras@gmail.com

D. E. Newman
Physics Dept.
University of Alaska
Fairbanks AK 99775
denewman@alaska.edu

I. Dobson
ECpE Department
Iowa State University
Ames IA 50011
dobson@iastate.edu

Abstract

The power transmission grid, as well as many other complex critical infrastructure systems, display characteristics of a critical or near critical complex system with the risk of large cascading failures. Understanding this risk and its relation to the system state as it evolves, could allow for a more realistic risk assessment over time and even for mitigation or at least preparation if in a high risk state. In order to facilitate this type of analysis in the context of the power grid as a complex system, we develop a new measure of the complex system state, the generalized autonomous generational average Lambda-gaga, which correlates with the risk. The Lambda-gaga measure is an extension of the standard cascading propagation measure lambda but avoids the contamination of that measure by small events.

1. Introduction

Complex critical infrastructure systems such as the power transmission grid are prone to cascading failures of all sizes [1,2]. This is characteristic of complex systems operating near their critical point [3-7]. However, in estimating the risk of large blackouts, a question that remains open is how close a system is operating to the critical point or operational limit. Ideally, we would like to know this both on average and on a short time scale to be able predict the risk of large failures. It is not clear that in complex systems of high dimensionality such as the power grid there is well defined short time "critical point" or if there is rather more of a "critical region".

We can look at fluctuations of a typical power system property such as the loading. The loading oscillates in time on a daily basis. See for instance the oscillating red line in Fig. 1. Let us assume that there is a critical loading (black line in Fig. 1), and we want to evaluate how close the system gets to the critical point (the length of the arrows). Much of the time the system is well below the critical point and any initiating failures that occur during this time will tend to have very short cascades that have little effect. However, if the initiating failures happen to coincide with the more highly loaded conditions, then there is a chance of a large cascade. We want to be able to estimate the proximity of the system to criticality independently of the regular fluctuations well below criticality.

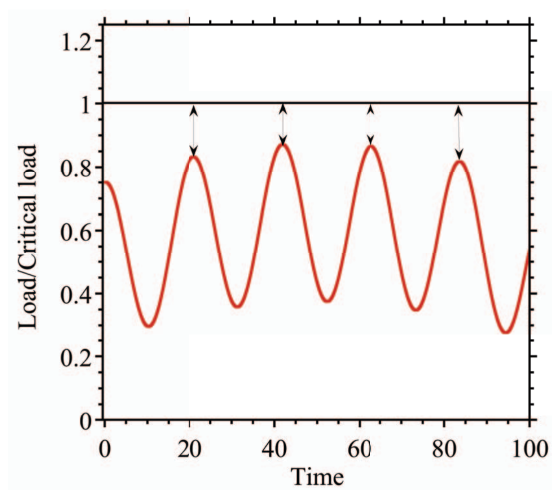

Fig. 1. Time evolution of the system loading (in arbitrary units) compared to the critical loading.

One of the standard measures of the proximity to the critical point is the average propagation parameter λ. λ is easily calculated from the time series of failures during a cascade. The cascading failures are grouped

1530-1605/15 $31.00 © 2015 IEEE
DOI 10.1109/HICSS.2015.329

into generations so that the failures in one generation (parents) give rise to (children) failures in the next generation. When simulating cascading failure, the generations correspond to iterations of the main loop calculating the failures.

The standard Harris estimator of the propagation λ is calculated by taking the average of the ratio of child failures (generation i) to parent failures (generation i-1) over all the cascading events [8-11]. If the ratio is bigger then 1, the propagation can grow very large; if the ratio is smaller then 1, it eventually dies off. (As an interesting side note, the average propagation was a measure originally developed to investigate the survival and potency of the lines of descent of the British upper classes [12]). One problem with this measure (in addition to its original application) is that it averages over all fluctuations of the system well below criticality.

Here, we will develop the **g**eneralized **a**utonomous **g**enerational **a**verage (gaga) λ_{gaga} measure based on the propagation of the cascading failures in the power system modified from the average propagation λ just discussed.

Although the λ_{gaga} measure is general, the simulated results illustrating the new measure presented here are obtained with the OPA power transmission model [3, 6, 15]. The OPA model for a fixed network configuration represents transmission lines, loads and generators with the usual DC load flow approximation using linearized real power flows with no losses and uniform voltage magnitudes.

There are two basic timescales modeled in OPA. For the slow, long time scale part, the OPA blackout model represents the essentials of slow load growth, cascading line outages, and the increases in system capacity coming from the engineering responses to blackouts. The short timescale part captures the cascading line outages leading to a blackout, which are regarded as fast dynamics and are modeled as follows. Starting from a solved base case, blackouts are initiated by random line outages with a probability p_0. Whenever a line is outaged, the generation and load is re-dispatched using standard linear programming methods. This is because there is more generation power than the load requires and one must choose how to select and optimize the generation that is used to exactly balance the load. The cost function is weighted to ensure that load shedding is avoided where possible. If any lines were overloaded during the optimization, then these lines are outaged with probability p_1. The process of re-dispatch and testing for outages is iterated until there are no more outages. The total load shed is, then, the power lost in the blackout.

The slow dynamics model the growth of the load demand and the engineering response to the blackout by upgrades to the grid transmission capability. The slow dynamics is carried out by the following small changes applied each time a potential cascading failure is simulated: All loads are multiplied by a fixed parameter that represents the rate of increase in electricity demand. If a blackout occurs, then the lines involved in the blackout have their line flow limits increased slightly. The grid topology remains fixed in the upgrade of the lines for model simplicity. In upgrading a grid it is important to maintain coordination between the upgrade of generation and transmission. The generation is increased at randomly selected generators subject to coordination with the limits of nearby lines when the generator capacity margin falls below a threshold.

The test networks we use here include 200, 400, 800 and 1600 node networks that are randomly constructed to have properties similar to real power networks, but allow the effect of changing network size to be studied. These artificial networks are constructed based on the methods in [13]. We also include some results on observed and simulated data from the WECC grid. In order to collect reasonable statistics, we typically simulate for 40000 simulation steps or more depending on the system size.

2. Defining λ_{gaga}

To start, we consider a time interval long enough to have many cascades. The cascade length is simply the number of iterations in a cascading failure and j_M is the maximum number of iterations in the cascades during the time evolution. Then, for every $k < j_M$, we can calculate

$$\lambda_k(i) = \frac{O_k(i)}{O_k(i-1)} \tag{1}$$

where $O_k(i)$ is the sum of the number of overloaded lines in iteration i for all cascades with length k or greater than k. Note that $\lambda_0(i)$ is our usual definition of the λ parameter. Then, for k > 2 (to have values to average over), we define an average value of λ over the iterations; that is,

$$\langle \lambda \rangle_k = \frac{1}{k-2} \sum_{i=2}^{k-1} \lambda_k(i) \tag{2}$$

The value of $< \lambda >_k$ generally saturates quickly for k > 3 as shown in Fig. 2.

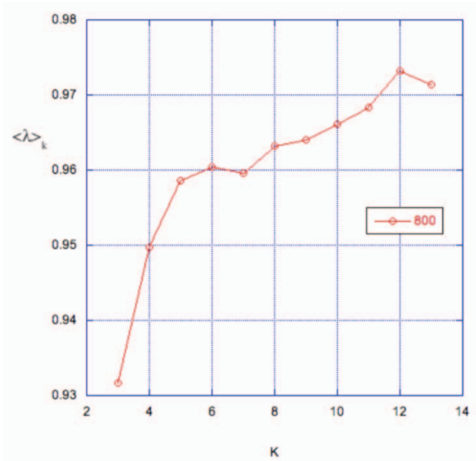

Fig. 2. Lambda k values for the 800-node network as a function of k.

The new measure lambda gaga is the averaged value over all values of the cutoff k > 2. That is, λ_{gaga} is

$$\langle\langle\lambda\rangle\rangle_1 = \frac{1}{j_m - 3}\sum_{k=3}^{j_m}\langle\lambda\rangle_k \tag{3}$$

This is a measure of the averaged propagation value for the longer cascades. It is then an averaged value of the propagation λ when the system is closer to critical that is independent or autonomous from the subcritical fluctuations of the complex system with time. That is, the measure is a generalized autonomous generational average (gaga). Its value can give a measure of the proximity to criticality.

To compare different evolutions of the system from different system states, measuring the propagation of the cascade is not sufficient. The maximum length of the cascades is another factor that indicates the proximity to the critical point. Therefore, to compare different states, we take the maximum value of j_M for the different states, let us call this value J_M, then we define an alternative λ_{gaga} by

$$\langle\langle\lambda\rangle\rangle_2 = \frac{1}{J_m - 3}\sum_{k=3}^{j_m}\langle\lambda\rangle_k \tag{4}$$

This is equivalent to including zeros for the values of lambda at the missing iterations, weighting the result toward states with longer cascades even if the propagation λ is the same.

For instance, we can consider the OPA results for the 400-node network. We divided the total length of time, 10^5 days, in periods of 1000 days and in each of these periods we calculate the two definitions of λ_{gaga}. The result for the full time evolution is shown in Fig. 3.

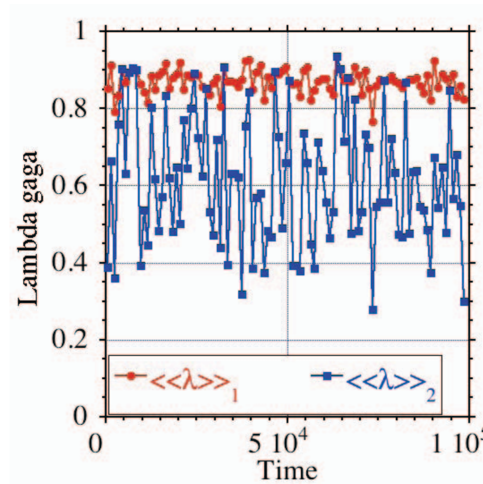

Fig. 3. Both λ_{gaga} values for the 400-node network

We can see that the value of $<<\lambda>>_1$ has small oscillations in time with a mean value of 0.87 and standard deviation 0.03. On the other hand, the value of $<<\lambda>>_2$ has much larger fluctuations in time with a mean value of 0.61 and standard deviation 0.17. The latter shows a greater variation because it combines propagation with the length of the cascade. When combined these two measures give complementary information about the system state with $<<\lambda>>_1$ giving information about the average state and $<<\lambda>>_2$ giving more time-localized information. For example, the peaks in $<<\lambda>>_2$ where the value is the same as or close to $<<\lambda>>_1$ the cascades are dominantly of the maximum length, while in the valleys, where the $<<\lambda>>_2$ value is much less then $<<\lambda>>_1$ the lengths are on average much less then the maximum. As we approach the maximum length of the cascades, there are few cascades of that length so the statistics become sparse.

It is still an open question what the shortest time interval one can use to make an optimal determination of these measures. The criteria to use for these determinations is still being investigated, so for now we use 1000 days which seems sufficiently converged for measure 1 and yet short enough to show significant differences in measure 2.

We can now ask if there is any correlation between λ_{gaga} and the size of the blackouts. In every time interval that we have calculated lambda gaga, we have also evaluated the averaged load shed normalized to the power demand and its maximum value. In Fig. 4, we can see failure size as a function of the two λ_{gaga} measures for the 400 and 800 node networks.

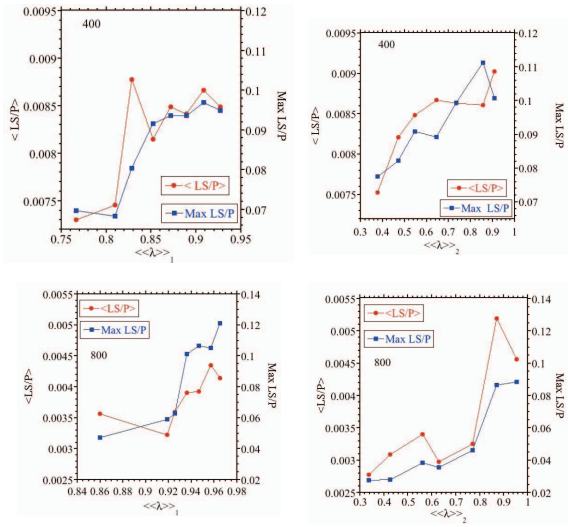

Fig. 4. Averaged load shed normalized to the power demand and its maximum value as a function of λ_{gaga} for the 400 and 800 node networks

Fluctuations notwithstanding, we can see a clear correlation in the results plotted in Fig. 4. The resolution is not very good because we do not have many samples of values of λ_{gaga} as can be seen in Fig. 3. However, higher values of λ_{gaga} lead to larger blackouts, both in average normalized load shed and in the maximum load shed as shown in figure 5. Interestingly, there is no observable correlation with the overall probability of a blackout, for the 400-node network.

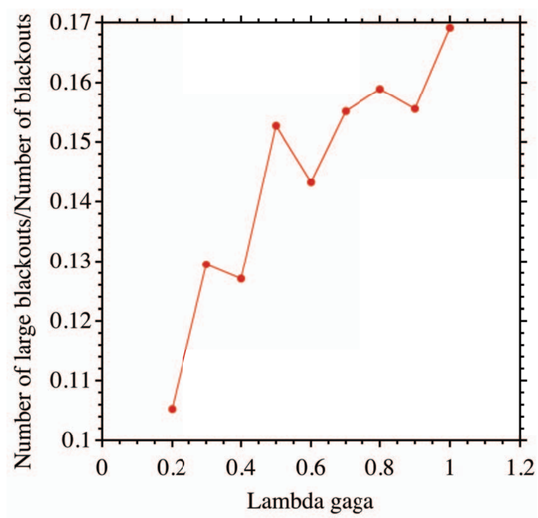

Fig. 5. Likelihood of a blackout being large as a function of λ_{gaga} for the 400-node network

This result is to be expected since we have averaged over 500 consecutive days in order to calculate λ_{gaga}. This implies averaging over many different states of the system and we obtain an averaged value of the probability. Therefore, this way of calculating λ_{gaga} is good as an overall measure of proximity to the critical point, but it does not function well as a metric to be correlated with short time risk of a blackout.

3. A second approach to the calculation of λ_{gaga}: state of the system vs global properties of the system

While the calculation of the λ_{gaga} measures averaged over time intervals is both reasonable and useful (as obtaining the time evolution on as short a time scale as possible is a goal of this work), one potential problem encountered in the previous evaluation is the time averaging that combines multiple system states. A possible alternative is to order the different cascades by value of <M> (the spatially averaged normalized line loading) and do the averaging of cascades within a bin of <M>. That could insure a more homogeneous (by the <M> measure) set of system states. We have done this for the three artificial networks and the results are plotted in Fig. 6. The figure shows that the value of $\langle\langle\lambda\rangle\rangle_1$ is nearly constant and practically independent of the value of <M>. However, $\langle\langle\lambda\rangle\rangle_2$ shows some degree of correlation with <M>. Therefore, the second definition may be a better representation of the "instantaneous" proximity to the critical point while the first definition is giving an overall state measure. There are oscillations seen in these correlations because the statistics of large events are relatively sparse.

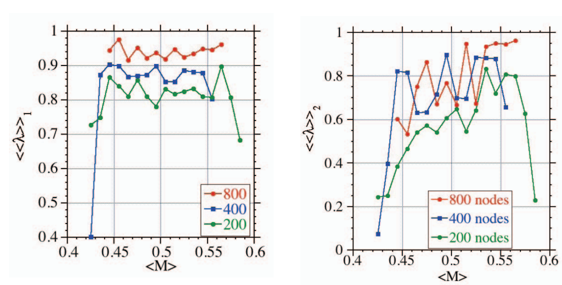

Fig. 6. The two measures of λ_{gaga} as a function of <M> for the three artificial networks.

To be able to take better account of sample sparsity, in Fig. 7, we show the number of events per bin for the three networks. We can see in the second Fig. 6 that the 200-node case with more samples shows a smoother behavior and better correlation. To avoid distortions at the two ends of the distribution because of smaller number of samples in the bins we will exclude bins with less than 200 events.

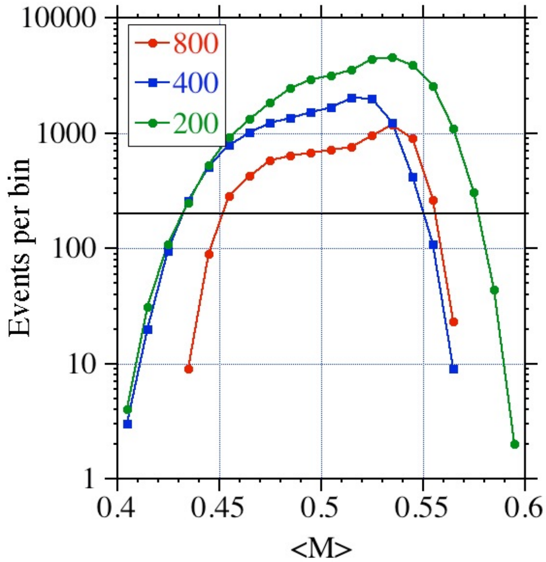

Fig. 7. Number of events per <M> bin for the three artificial networks

After accounting for sample size effects, it remains clear that $<<\lambda>>_1$ measures only propagation as it does not change much when <M> is varied. This measure is an overall property of the system as large failure events may happen at any value of <M> although their probability can depend on <M>. However, the averaged length of a cascade shows similar dependence on <M> to the one we have seen for the blackout size, as shown in Fig. 8, and $<<\lambda>>_2$ carries some of this information which is why this measure has a correlation with <M>.

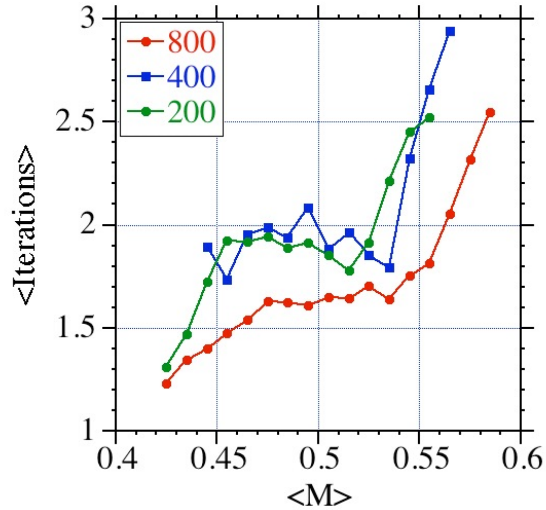

Fig. 8. Averaged number of iteration of the cascading failures as a function of <M> for the three artificial networks

4. λ_{gaga} and criticality

The next question to ask is can we relate the λ_{gaga} measures to other standard measures of criticality? One way is to look at the relation between those parameters and the existence and value of a power law in the PDF and Rank functions of the size of the blackouts. To explore this, we can vary lambda gaga using p_1 as the control parameter. A lot of things change when we change p_1, but we will focus just on <<λ>>.

To determine the exponent of the possible power tail, we use the Clauset method [14] because it gives us both the exponent and a criterion for when such a power tail may exist. Results for the 400-node network are summarized in Fig. 9.

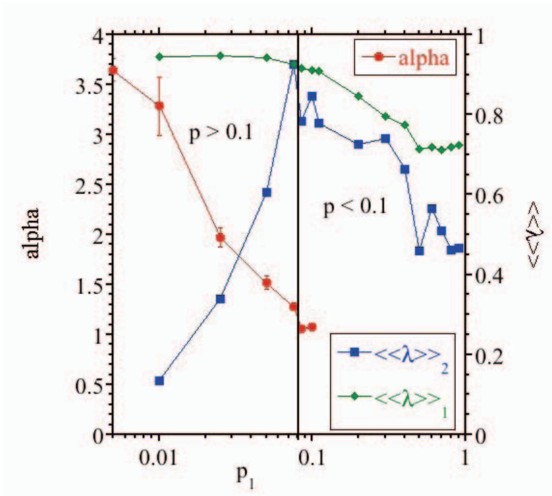

Fig. 9. Exponent of the rank function of the load shed normalized to power demand with the two expressions for λ_{gaga} as a function of p_1 for the 400-node network.

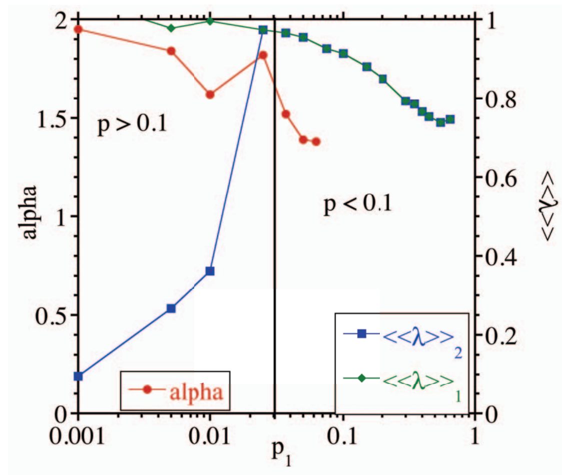

Fig. 10. Exponent of the rank function of the load shed normalized to power demand with the two expressions for λ_{gaga} as a function of p_1 for the 800-node network

Figure 9 contains a great deal of information, including the power law exponent alpha determined by Clauset's method The figure is divided in two regions following the criterion of Clauset on the credibility of the power law. On the left, the power law is credible and the value of the exponent alpha decreases with increasing p_1. At the boundary the exponent is close to 1, a characteristic of a critical point. On the right the tail of the Rank function is dominated by an exponential decay. In this region the dynamics of the system have some of the characteristics of a supercritical system.

In the figure, we can see that $\langle\langle\lambda\rangle\rangle_2$ peaks at the critical point, while $\langle\langle\lambda\rangle\rangle_1$ does not show a dependence on p_1 in the region of criticality. It is practically constant in the region where the Rank function has a power tail.

A similar plot is shown in Fig. 10 for the 800-node network. However, we can see very similar features to the 400-node case. It is notable, that the point at which $\langle\langle\lambda\rangle\rangle_1$ and $\langle\langle\lambda\rangle\rangle_2$ come together is the point at which the system becomes supercritical. This could make the combination a useful diagnostic for measuring proximity to the supercritical region. It would also be interesting to investigate if different systems have different proximities to this point such as summer vs winter, US vs Europe, developed vs less developed systems, etc.

5. Is Lambda gaga related to blackout risk?

We have already seen that the blackout risk is related to metrics that describe the state of the system but not always to λ_{gaga}, which gives a global property of the system. However, we could ask further if one system is associated with a larger λ_{gaga} than a second one, is the first system at higher risk of blackouts?

To change λ_{gaga} in OPA, we have two possible parameters p_1 and the system size. We consider here a scan on each of these parameters. For each scan, we calculate the probability of a blackout and the averaged size of the blackout. Knowing both, we can construct a risk function [16]. The results for the p_1 scan are given in Fig. 11 and compared to λ_{gaga}. We have only plotted $\langle\langle\lambda\rangle\rangle_2$ because it was the measure more strongly correlated with the criticality results.

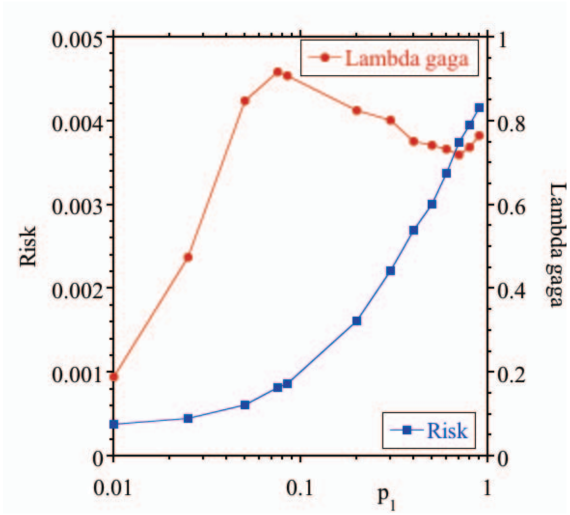

Fig. 11. Risk of a blackout and λ_{gaga} for the p_1 scan. The p_1 scan was done for the 400-node network.

From the figure we can see that there is a strong correlation between Lambda-gaga and risk for small values of p_1 (consistent with the left sides of Figs. 9 and 10) but no correlation between lambda gaga and the blackout risk for values of p_1, greater then 0.1, which is consistent with a supercritical state due to an unrealistically large p_1.

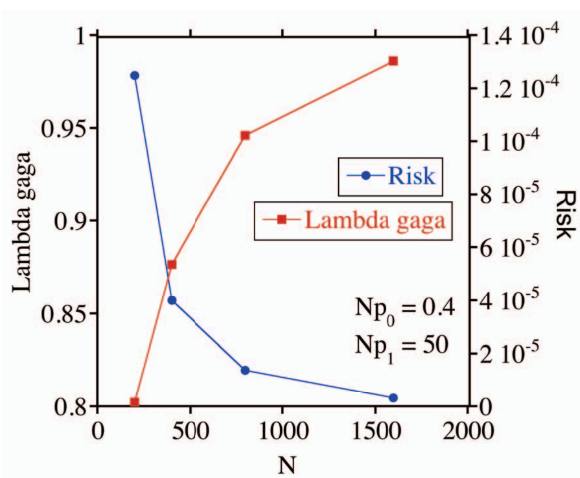

Fig. 12. Risk of a blackout and λ_{gaga} for the size scan, with constant Np_0 and Np_1.

While, Fig. 12 shows no correlation (or an anti-correlation) between λ_{gaga} and Risk This is likely because Np_1 and Np_0 are being held constant as the size N is increased. In this case while the system might be getting more critical, the frequency and propagation of the failures could be decreasing due to the shrinking p_1 and p_0. This is consistent with Fig. 13 in which p_0 and p_1 are held constant as the system size increases showing a marked increase in risk with size and a strong correlation with Lambda gaga. It is worth pointing out that what is held constant when looking at size scalings is more subtle then might appear. The two ways of approaching size scaling shown here are likely each reasonable depending on whether a system is growing in size or more detail is being put into the system. For the constant Np_0 and Np_1 scaling, Lambda gaga grows as the size increases because larger systems show more critical behavior, however the risk decreases because p_0 and p_1 (the failure probability parameters) both get smaller, decreasing the failure frequency and propagation. In contrast, when p_0 and p_1 are kept constant, as the system gets larger the risk increases as Lambda gaga does with some evidence that Lambda gaga saturates at the largest sizes. This too can be understood as coming from a higher overall probability of random failure triggers, since Np_0 gives that probability and a higher likelihood of propagation as Np_1 is related to that overall probability.

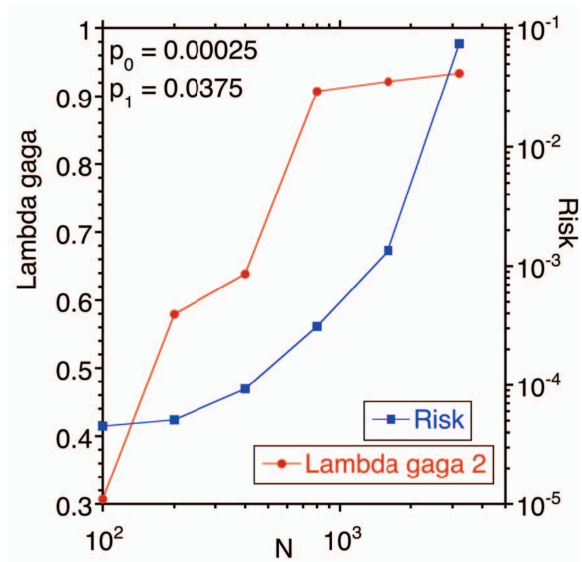

Fig. 13. Risk of a blackout and λ_{gaga} for the size scan, with constant p_0 and p_1.

The risk calculations shown were done following ref [16] and with scanning sizes from 100 to 1600 nodes.

6. Real and model data

The λ_{gaga} computations can be applied to real data as well as model data. Figure 14 shows λ_{gaga} as a function of the cutoff k for observed automatic line outage data from a utility in the Western USA interconnection. This line outage data is explained in detail in [10]. It can be clearly seen that the system reaches it's saturated value at a minimum cascade length k of ~ 3 and the saturated value is the highly critical value of 1. The line outage data sets we have are too short to perform the time series analysis so for comparison in Fig. 14, we show the same analysis for both a series of artificial grids of sizes 200-800 (on the left) and 2 WECC model grids (on the right). The WECC model grids have 2508 and 1553 nodes and are described in [15].

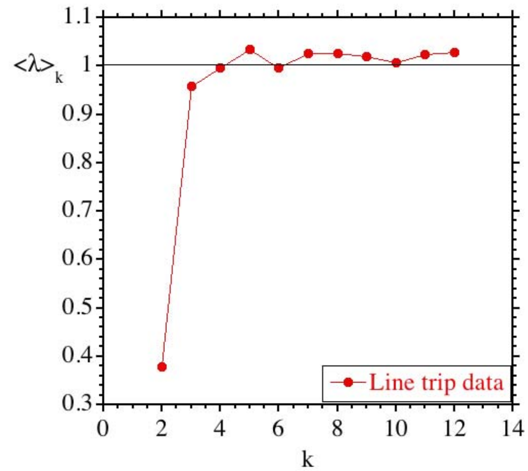

Fig. 14. λ_{gaga} versus k (the cutoff) for utility line trip data showing a critical state.

The left side of Fig. 15 shows that for the largest network, λ_{gaga} also saturates by k of three at a value close to 1. This suggests that the minimum size needed for the system to show the critical characteristics is ~800 nodes. The WECC grid models (Fig 14 right hand side) approach the saturated value more slowly then the utility data or the artificial network data. This could be due to the fact that the utility data was for one utility rather then the entire WECC and was therefore confined to a fairly homogeneous sub-region of the grid making it more like the homogeneous artificial grids. This is in contrast to the entire WECC which is rather heterogeneous. Additionally, despite the large number of nodes in the model WECC grids, the fairly homogeneous sub-regions of the network are smaller then 800 nodes, which could pose an additional constraint on the grid dynamics.

The application to the utility data is promising but needs much more data to be able to say anything definitive about the overall state.

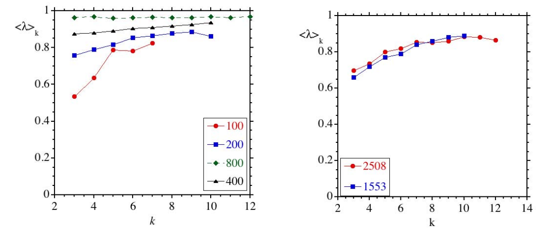

Fig. 15. λ_{gaga} versus k (the cutoff) for a set of homogeneous artificial grids (left) and a pair of WECC grids (right).

7. Conclusions

Two new measures of system state have been explored, showing the possibility of providing both an overall measure of the proximity to a critical point and, at a shorter time scale, the "instantaneous" state. These have been applied to both OPA model data from a series of artificial networks with parameters p_0 and p_1 [15] similar to real power grids. A preliminarily application to real data has yielded similar results. These measures appears to do a good job of quantifying the systems proximity to the "critical region" and therefore can give some indication of the "risk" associated with the state or operational regime. It would be interesting to calculate these measures in different operational periods, winter vs summer, regulated vs deregulated, etc in order to see if there are any differences in the propagation characteristics. The combination of the two measures and how close they are to each other appears to be a particularly promising metric for the system state and its overall proximity to the supercritical regime.

An additional application of these measures is for validation of the models used to study the grid. Since we can apply these to the real world cascading data as well as model results, comparison of the two gives a good measure of how well the models capture the true cascade dynamics.

Though we have used the power transmission grid as our example system, the measure is general and can be applied to any cascading system, perhaps even for investigating survival and potency of royal bloodlines.

ACKNOWLEDGMENTS

We gratefully acknowledge support in part from NSF grant CPS-1135825 (ID) and in part by the California Energy Commission, Public Interest Energy Research Program. This paper does not necessarily represent the views of the Energy Commission, its employees or the State of California. It has not been approved or disapproved by the Energy Commission nor has the Energy Commission passed upon the accuracy or adequacy of the information. Two of us (BAC and DEN) thank the financial support of Universidad Carlos III and Banco Santander through a Càtedra de Excelencia.

7. References

[1] B.A. Carreras, D.E. Newman, I. Dobson, A.B. Poole, Evidence for self organized criticality in a time series of electric power system blackouts, IEEE Transactions on Circuits and Systems I, vol. 51, no. 9, September 2004, pp. 1733-1740.

[2] P. Hines, J. Apt, S. Talukdar, Large blackouts in North America: Historical trends and policy implications, Energy Policy, vol.37, 2009, pp. 5249-5259.

[3] B.A. Carreras, V.E. Lynch, I. Dobson, D.E. Newman, Complex dynamics of blackouts in power transmission systems, Chaos, vol. 14, no. 3, September 2004, pp. 643-652.

[4] J. Chen, J.S. Thorp, I. Dobson, Cascading dynamics and mitigation assessment in power system disturbances via a hidden failure model, International Journal of Electrical Power and Energy Systems, vol 27, no 4, May 2005, pp. 318-326.

[5] D.P. Nedic, I. Dobson, D.S. Kirschen, B.A. Carreras, V.E. Lynch, Criticality in a cascading failure blackout model, International Journal of Electrical Power and Energy Systems, vol 28, 2006, pp 627-633.

[6] D.E. Newman, B.A. Carreras, V.E. Lynch, I. Dobson, Exploring complex systems aspects of blackout risk and mitigation, IEEE Trans. Reliability, vol. 60, no. 1, March 2011, pp. 134-143.

[7] D.E. Newman, B.A. Carreras, N.S. Degala, I. Dobson, Risk metrics for dynamic complex infrastructure systems such as the power transmission grid, Forty-fifth Hawaii International Conference on System Sciences, Maui, Hawaii, January 2012.

[8] T.E. Harris, Theory of branching processes, Dover NY 1989.

[9] I. Dobson, B.A. Carreras, V.E. Lynch, B. Nkei, D.E. Newman, Estimating failure propagation in models of cascading blackouts, Probability in the Engineering and Informational Sciences, vol 19, no 4, October 2005, pp 475-488.

[10] I. Dobson, Estimating the propagation and extent of cascading line outages from utility data with a branching process, IEEE Transactions on Power Systems, vol. 27, no. 4, November 2012, pp. 2146-2155.

[11] J. Kim, K.R. Wierzbicki, I. Dobson, R.C. Hardiman, Estimating propagation and distribution of load shed in simulations of cascading blackouts, IEEE Systems Journal, vol. 6, no. 3, September 2012, pp. 548-557.

[12] H. W. Watson, F. Galton, On the probability of the extinction of families, Journal of the Anthropological Institute of Great Britain and Ireland, vol. 4, 1875, pp. 138-144.

[13] Z. Wang, R.J. Thomas, A. Scaglione, Generating random topology power grids, Forty-first Hawaii International Conference on System Sciences, Hawaii, January 2008.

[14] A. Clauset, C.R. Shalizi, M.E.J. Newman, Power-law distributions in empirical data, SIAM Review, vol. 51, no. 4, Nov. 2009, pp.661-703.

[15] B.A. Carreras, D.E. Newman, I. Dobson, N.S. Degala, Validating OPA with WECC data, Forty-sixth Hawaii International Conference on System Sciences, Maui, Hawaii, January 2013.

[16] B.A. Carreras, D.E. Newman, I. Dobson, Does size matter?, Chaos: An Interdisciplinary Journal of Nonlinear Science 24 (2), 023104 (2014)