IEEE Power Engineering Society General Meeting, Tampa FL USA, June 2007

Where is the edge for cascading failure?: challenges and opportunities for quantifying blackout risk

Ian Dobson, *Fellow IEEE*

Abstract— **Quantifying the overall risk of blackout due to cascading failure and determining the corresponding safe limits for power system design and operation are challenging problems. Large blackouts involve long, complicated and diverse cascades of events that often are unlikely or unexpected. An exhaustive and detailed analysis of these cascading events before the blackout occurs is impossible because of the huge number of possible combinations of unlikely events. Despite these challenges, approaches to quantify the overall risk of blackouts are emerging and we give a tutorial account reviewing these emerging approaches and their prospects.**

We summarize the implications for blackout risk of the power-law region in the observed distribution of sizes of North American blackouts. High-level probabilistic models of cascading failure and power system simulations suggest that there is a critical loading at which expected blackout size sharply increases and there is a power law in the distribution of blackout size. This critical loading could serve as a reference point for determining the "edge" for cascading failure risk. We model cascading failure as an initial disturbance that sometimes propagates to become much more widespread. The size of the initial disturbance and the average amount of propagation of the failures can be estimated from data from simulated cascades. We suggest that these estimates could be used to efficiently quantify the blackout risk. We summarize initial testing on power system simulations of cascading overloads and speculate that extending this approach to process data from series of cascades occurring in the power system could lead to direct monitoring of power system reliability.

I. Introduction

Cascading failure occurs when a sequence of failures successively weaken the electric power transmission system and make further failures more likely. Large blackouts typically become widespread by a complicated sequence of cascading failures [35]. Here we use the term failure of a power system component in a general sense of an event that includes misoperation or being unavailable to transmit power due to tripping by automatic or manual protection or damage. The failures are typically rare and unanticipated because the likely and anticipated failures have already been accounted for in power system design and operation. The determined efforts to analyze and mitigate the anticipated and likely failures have usually been successful and this accounts for the historically high reliability of bulk power generation and transmission systems in developed economies.

I. Dobson is with the ECE department, University of Wisconsin, Madison WI 53706 USA; email dobson@engr.wisc.edu. This paper is an account of work sponsored in part by the Power Systems Engineering Research Center (PSERC). We gratefully acknowledge support in part from NSF grants ECCS-0606003 and SES-0623985.

Cascading failure is very challenging to analyze because of the huge number of possible rare and unanticipated failures and the dependence of the failures on the previous failures in the sequence of failures. There are many ways in which failures can influence further failures, including overloads, hidden failures of protection systems, software, control or operator error, flawed or inapplicable operating procedures, and a variety of dynamic and stability phenomena including transients, oscillations, transient stability and voltage collapse. Each large blackout typically combines several of these diverse types of interactions between failures. One common aspect of these diverse interactions is that the interactions tend to be stronger when the power system is more stressed or loaded. For example, tripping a highly loaded transmission line causes a larger transient and a larger steady-state redistribution of power flow. Moreover, the impact on the remaining operating margins of other components is larger if the power system is more stressed.

Cascading failure can be to some extent mitigated by design and operation to limit the start of cascades. In the n-1 criterion, a contingency list is established and the power system is operated so that no single contingency will lead to further failures. Sometimes a chosen subset of multiple contingencies is also considered. In power systems of practical size, the extension of the n-1 criterion to systematically account for 3 or 4 contingencies is precluded by a combinatorial explosion in the number of possibilities. This paper addresses the complementary problem of monitoring the propagation of cascades after they are started.

One useful way to analyze cascading failure is to examine the detailed sequence of failures of a particular blackout after it has happened. One takes a deterministic point of view and works out the way in which the particular sequence of failures occurred. The analysis is time-consuming, but it provides useful engineering data for strengthening weaker parts of the system. Careful attention to the details of previous blackouts motivates good practice and this is vigorously pursued by engineers for reasons of professionalism and in response to significant pressures from society and industry. In summary, "blackouts cause reliability".

It is also practical to simulate some of the detailed interactions of cascading failure and we summarize some of the simulation capabilities in the Appendix. This simulation capability is significant in providing an approximate direct assessment of cascading failure that can identify both particular cascading sequences and some sampling of possible cascading sequences. Moreover, these simulations can be used to initially test methods for assessing cascading risk from real data.

In this paper, instead of examining the details of individual blackouts, we discuss the prospects for an overall risk analysis of blackouts of all sizes. This is necessarily a global and top-down analysis of bulk system properties that neglect the details of the cascading failure sequence and adopts a probabilistic point of view. This probabilistic and high-level approach should be seen as complementary to the more traditional detailed and deterministic analyses.

We begin in section II by discussing the observed statistics of blackouts. Section III discusses blackout cost and blackout risk. Section IV describes the critical loading at which a change in the behavior of the blackout risk can be identified. Section V describes branching and other high-level probabilistic models for cascading. Section VI describes Monte Carlo and branching process parameter estimation methods to estimate the distribution of blackout size from cascading failure simulations. The paper concludes and discusses challenges such as direct monitoring of blackout risk in section VII.

A review of the literature in cascading failure blackouts may be found in [23].

II. Observed data for the frequency of blackouts

Large blackouts are rarer than small blackouts, but how much rarer are they? There is empirical data that addresses this question by plotting blackout frequency or probability as a function of blackout size. One might expect a probability distribution of blackout sizes to fall off exponentially. That is, for the larger blackouts, doubling the blackout size squares its probability and so, after many squarings, the very largest blackouts have vanishingly small probability. However, analyses of the published international data show that the probability distribution of the blackout sizes does not decrease exponentially with the size of the blackout, but rather has an approximate power law region (North America [13], [7], [14], [38], Sweden [28], Norway [5], New Zealand [2], China [37]). For example, Figure 1 plots on a log-log scale the empirical probability distribution of energy unserved in North American blackouts. The fall-off with blackout size is close to a power law dependence with an exponent between –1 and –2. (Note that a power law dependence with exponent –1 implies that doubling the blackout size only halves the probability. A power law dependence with exponent –1 appears as a straight line of slope –1 on a log-log plot.) Thus the North American data suggests that large blackouts are much more likely than might be expected from the common probability distributions that have exponential tails. The power law region is of course always limited in extent in a practical power system by a finite cut off corresponding to the largest possible blackout.

The approximate power law region can be qualitatively attributed to the dependency of failures in the blackout. As the blackout progresses, the power system usually becomes more stressed, and it becomes more likely that further failures will happen. This weakening of the power system as failures occur makes it more likely that a smaller blackout will evolve into a larger blackout.

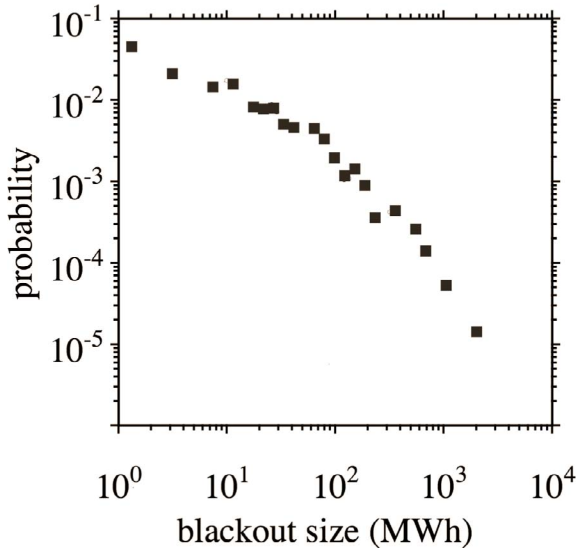

Fig. 1. North American blackout size probability distribution from NERC Disturbance Analysis Working Group data.

There are several useful measures of blackout size. Energy unserved and power or customers disconnected are measures that impact society. An example of a measure of disturbance size internal to the power system is number of transmission lines tripped.

III. Blackout risk

Risk is often defined as the product of frequency and cost. Therefore the risk of a blackout near a given size is the product of the frequency of blackouts near that size times the cost of those blackouts.

The direct cost of blackouts is most simply modeled as proportional to the energy unserved but there are numerous uncertainties in estimating this cost [6]. Blackouts halt the economy and impose direct costs on the public and business. Especially for large blackouts, there are sometimes large indirect costs such as those resulting from social breakdown and impacts on other infrastructures. Utilities can incur reputational, legal and regulatory costs and the costs of upgrades or personnel to improve reliability. A backstop feedback for reliability is provided by elected officials, who are certain to act if there are repeated large blackouts.

In the case of an exponential dependence of blackout probability on blackout size, large blackouts become rarer much faster than blackout costs increase so that the risk of large blackouts is negligible. However, in the case of a power law dependence of blackout probability on blackout size, the larger blackouts can become rarer at a rate similar to the rate at which costs increase, and then the risk of large blackouts is comparable to the risk of small blackouts [10]. For example, if the power law dependence has exponent –1, then doubling the blackout size halves the probability of the blackout and doubles its cost and therefore the risk of the blackout remains the same. Thus power laws in blackout size distributions significantly affect the risk of large blackouts and this risk justifies the study and analysis of large blackouts. That is, although large blackouts are rarer than small blackouts, their costs are higher and they become rarer slowly enough that they can have comparable risk.

There can be tradeoffs between the frequencies of small and large blackouts [10]. One very extreme and uneconomic example is that operating a power system without using any interarea tie lines would eliminate large blackouts and increase the frequency of small blackouts. Another example is that a policy of never shedding load in emergencies to solve local problems would tend to reduce small blackouts but would tend increase larger blackouts. These considerations, together with the roughly comparable risk of small and large blackouts, suggest that engineering efforts to mitigate blackouts should aim to manage the blackout risk by considering the joint reduction of the frequency of small, medium and large blackouts as well as trading off the blackout risk with the economic benefits of maximizing the use of the transmission system. To achieve this aim, it is necessary to be able to quantify the overall risk of blackouts of various sizes, particularly the blackout frequency and cost for the various sizes. The normalized blackout frequency as a function of blackout size is also known as the probability distribution of blackout size.

Another promising direction for cascading failure risk analysis computes the initial portions of the highest risk cascading sequences since these are candidate multiple contingencies for mitigation efforts [16], [32], [17].

IV. Critical loading

Consider cascading failure in a power transmission system in the impractically extreme cases of very low and very high loading. At very low loading, any failures that occur have minimal impact on other components and these other components have large operating margins. Multiple failures are possible, but they are approximately independent so that the probability of multiple failures is approximately the product of the probabilities of each of the failures. Since the blackout size is roughly proportional to the number of failures, the probability distribution of blackout size will have an exponential tail. The probability distribution of blackout size is different if the power system is operated recklessly at a very high loading in which every component is at its loading limit. Then any initial disturbance will cause a cascade of failures leading to total or near total blackout. The probability distribution of blackout size must somehow change continuously from the exponential tail form to the certain total blackout form as loading increases from a very low to a very high loading. We are interested in the nature of the transition between these two extremes. Our results presented below suggest that the transition occurs via a critical loading at which there is a power law region in the probability distribution of blackout size. This concept is shown in Figure 2.

There are two attributes of the critical loading:

1) A sharp change in gradient of some quantity such as expected blackout size as one passes through the critical loading.
2) A power law region in probability distribution of blackout size at the critical loading.

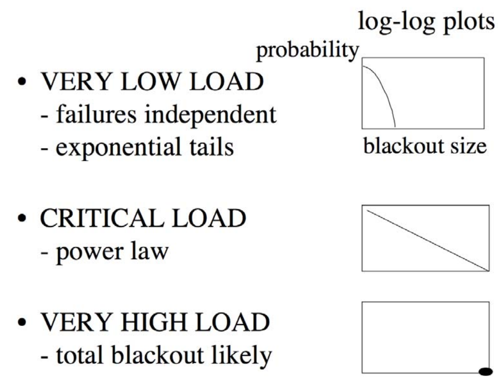

Fig. 2. Log-log plots sketching idealized blackout size probability distributions for very low, critical, and very high power system loadings.

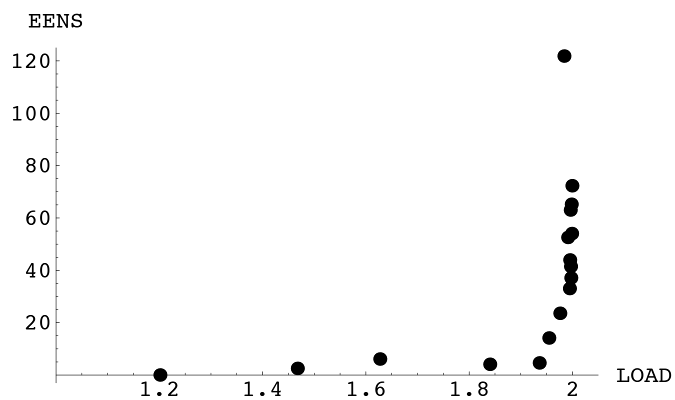

Fig. 3. Average blackout size (expected energy not served) in Manchester blackout simulation as loading increases on a realistic 1000 bus model. Critical loading occurs at kink where average blackout size sharply increases. Data from [33].

We use the terminology "critical" because this behavior is analogous to a critical phase transition in statistical physics. Critical behavior in power system blackout models as load is increased was first described in 1992 [8], [19], [9], [15] using models of cascading overloads and hidden failures. Observation of the critical loading in the most realistic power system model to date used the Manchester model on a 1000 bus European system example as shown in Figures 3 and 4 [33]. The average blackout size in Figures 3 remains low and changes little as load is increased until the critical loading is reached and then it sharply increases. This shows that, although a low average blackout size could be used to confirm that the power system is below criticality, average blackout size is not a good index of blackout frequency or risk.

The critical loading defines a reference point for increasing risk of cascading failure. Monitoring the proximity to the critical loading would give an indication of the risk of large, cascading failure blackouts. Although there is a substantial risk of large blackouts near criticality, it cannot be assumed based

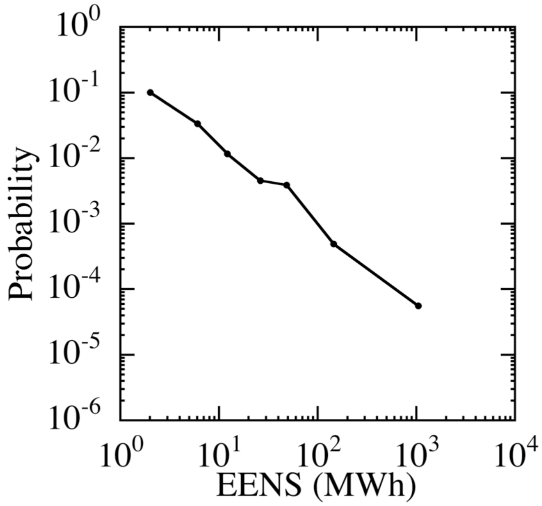

Fig. 4. Log-log plot of blackout size probability distribution obtained from Manchester blackout simulation at critical loading on a realistic 1000 bus model. Data from [33].

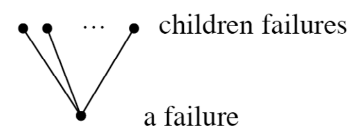

Fig. 5. Branching from a failure showing a random number 0,1,2,3,... of children failures. The mean number of children failures is λ.

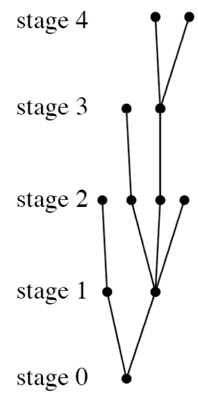

Fig. 6. Example of failures produced in stages by a branching process. Each failure independently has a random number of children failures in the next stage.

on current knowledge that power system operation near criticality is undesirable; there are substantial economic benefits in maximizing the use of the power transmission system that also have to be considered in an overall risk analysis. Indeed, one of the main aims of the research discussed in this paper is to make such a risk analysis practical so that the tradeoffs between economics and security involved in positioning the power system with respect to cascading failure risk can be quantified.

Although most authors have increased power system total load to stress the power system so that it moved through criticality, there may be quantities other than total load that more directly reflect the system stress with respect to cascading failure. This is an open question and these quantities are of great interest because of their potential for monitoring and mitigating cascading failure risk. Several papers have started to investigate the impact of generation margin [9], [10], [15] and generalized line outage distribution factors [11].

Over time, system stress or loading tends to increase due to load growth and tends to decrease due to the system upgrades and improvements that are the engineering responses to blackouts. It has been suggested that these opposing forces tend to slowly shape the power system towards a near power law distribution of blackout size [12]. Moreover, it is argued in [10] that blackout mitigation should take account of these opposing forces in order to achieve positive reliability results in the long-term.

V. Probabilistic models of cascading

Two high-level probabilistic models for cascading failure are the CASCADE model and branching processes. Both models have a large but finite number of identical components that can fail. The total number of components failed in each cascade varies randomly and the models are simple enough that there are analytic formulas for the probability distribution of the total number of components failed. Neither model directly represents the power system and their purpose is to provide simple models of a generic cascading process to better understand power system cascading.

We briefly summarize the CASCADE model [23]. The CASCADE model has components that fail when their load exceeds a threshold, an initial disturbance loading the system, and the additional loading of components by the failure of other components. The initial overall system stress is represented by upper and lower bounds on a range of random initial component loadings. The model neglects the length of times between failures, the structure of the power grid, and the diversity of power system components and interactions. In a parameter range of the most interest, the CASCADE model is well approximated by a branching process [20], so we proceed to explain the branching process in more detail.

Branching processes have long been used in a variety of applications to model cascading processes [27], [4], but their application to the risk of cascading failure is novel and was first suggested in [20], [21]. Here we informally describe the simplest branching process model that is called a single-type Galton-Watson process. The branching process gives a probabilistic model of the number of failures. There are a random number of initial failures that then propagate randomly to produce subsequent failures in stages. The initial failures are modeled as a random number of failures with mean θ. The subsequent failures are produced in stages or generations starting from the initial failures. Each failure in each stage (a "parent" failure) independently produces a random number 0,1,2,3,... of failures ("children" failures) in the next stage as shown in figure 5. The distribution of the number of children is called the offspring distribution. The children failures then become parents to produce the next generation and so on. If the number of failures in a stage becomes zero, then all subsequent stages have zero failures and the cascade stops. The mean number of children failures for each parent is the parameter λ. λ quantifies the tendency for the cascade to propagate. The intent of the modeling is not that each parent failure in some sense "causes" its children failures; the branching process simply produces random numbers of failures in each stage that can statistically match the outcome of cascading processes.

The branching process theory gives analytic formulas for the probability distribution of the total number of failures. For example, when the number of initial failures is a Poisson distribution of mean θ and the number of children failures for each parent failure is a Poisson distribution of mean λ, the total number of failure follows a generalized Poisson distribution that is parameterized by θ and λ [20]. Larger numbers of initial failures (larger θ) or larger cascade propagation (larger λ) will increase the frequency of larger numbers of total failures.

The eventual behavior of the branching process is governed by the propagation parameter λ. In the subcritical case of $\lambda < 1$ (each parent failure has on average less than one child), the failures will die out and this usually corresponds to either no blackout or a small blackout. In the supercritical case of $\lambda > 1$ (each parent failure has on average more than one child), although it possible for the process to die out, often the failures increase exponentially until the system size or saturation effects are encountered. At the critical case of $\lambda = 1$, the branching process has a power law distribution of the total number of failures with exponent -1.5.

An accelerated propagation model for the number of transmission line failures is described by Chen and McCalley in [18]. They also examine the fit of the accelerated propagation model and other models such as the generalized Poisson distribution to combined data for North American transmission line failures from [1].

VI. Estimating blackout size distribution from simulations

A. Monte Carlo approaches

There are a variety of blackout simulations that can produce samples of cascading failures (see Appendix). Without quantitative statistical analysis of these sample cascades, it is not clear how robust the power system is to cascading failure. For example, if one of the sample cascades is a large blackout, does this indicate a vulnerable power system, an unrepresentative rare event, or simply bad luck?

A straightforward way to obtain the distribution of blackout size empirically is to run the simulation exhaustively with a sample of initial conditions and/or Monte Carlo evaluation of any probabilistic decisions. There are technical challenges in choosing a good sample of initial conditions with sufficient range and uniformity over the possibilities of interest. Moreover, simulation is inherently time-consuming, especially in obtaining accurate estimates of the rare events in the (potentially) heavy tails of the distribution of blackout sizes. Importance sampling [15] and correlated sampling [30] are useful methods to reduce the simulation time.

B. Estimating blackout size distribution via estimating λ

If the cascading can be approximated by a branching process, we can estimate the parameters such as λ and θ of the branching process from the simulated cascade data and then use the branching process theory to compute the distribution of blackout size from the parameters [22]. This procedure is much more efficient than exhaustively running the simulation because it is model based and because the parameters being estimated relate to distributions such as the offspring distribution which do not have power laws. Of course, this approach relies on the branching process representing the gross behavior of the cascade with sufficient accuracy.

The power systems simulation produces some quantity such as number of transmission lines tripped in stages and we use this data to estimate the branching process parameters. The mean θ of the number of initial failures is estimated simply as the total number of initial failures divided by the number of sample cascades. The propagation λ is the mean number of children failures per parent failure. λ may be estimated from a sample of cascades by dividing the total number of children failures in the sample cascades by the total number of parent failures. However, this standard estimate [25] may require adjustment for saturation effects [22].

The propagation λ and mean initial failures θ are useful metrics describing the cascading in the simulation data. $\lambda = 1$ indicates criticality. Moreover, these parameters provide an estimate of the distribution of the total number of failures using formulas from the branching process theory [22]. This gives a way to verify the assumption that a branching process approximates the simulation results. One can simply run the simulation exhaustively to obtain an empirical distribution of total number of failures. This empirical distribution can then be compared to the estimated distribution of total number of failures. If the match is acceptable then the estimation via the branching process can be used to approximate the estimated distribution of total number of failures. An example of the match obtained in a subcritical test case from [22] is shown in Figure 7.

Tracking the numbers of transmission lines that fail in each stage of the cascade gives an integer number of failures in each stage. This is modeled by the Galton-Watson branching process explained above. However, it is also useful to track in each stage of the cascade continuously variable measures of blackout size such as power shed. Cascades of continuously varying quantities are modeled by continuous state branching processes and these can be applied in a similar way to estimate branching process parameters and hence compute the probability distribution of the power shed [39].

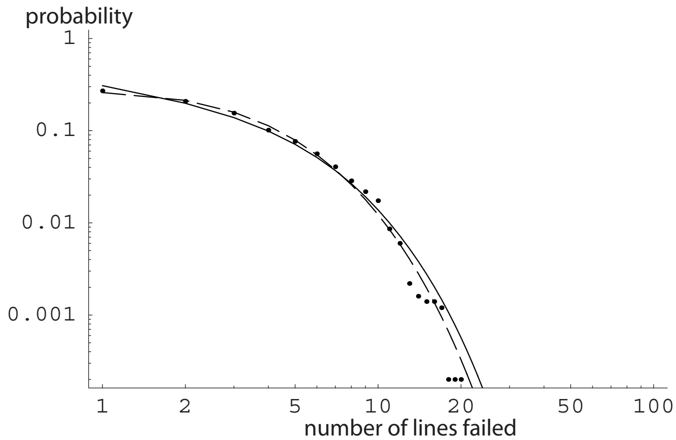

Fig. 7. Probability distributions of total number of transmission lines failed obtained by different methods. The dashed line is obtained by estimating parameters ($\lambda = 0.4$ and $\theta = 1.5$) from simulation data and assuming a branching process model of the cascading. The dots are obtained empirically from the same simulation data. The simulation is the OPA model of cascading line outages in blackouts [9] and the test case is the IEEE 118 bus system. Figure reprinted from [22].

VII. Conclusion and challenges

In this paper, we review some emerging approaches to the challenging problem of cascading blackout risk. International data suggests an increased (but still rare) frequency of large blackouts that makes the risk of large blackouts substantial. We argue that the problem of avoiding blackouts should be posed as jointly mitigating the risk of blackouts of all sizes. There is evidence of a change in cascading blackout risk at a critical loading or stress that provides a reference point for power system design and operation with respect to cascading failure. There are a range of power system simulations and high-level models of cascading failure that could be used to develop risk analysis methods appropriate to cascading failure. One of the high-level models is a branching process that models cascading failure probabilistically as initial failures followed by propagation of these failures. The amount of propagation λ can be quantified from data and is under investigation as a metric for cascading failure that can help predict overall blackout risk.

We now comment on some key challenges in quantifying and monitoring cascading blackout risk.

A. Access to data

Blackout and failure data is highly sensitive information. However, if researchers cannot access real blackout and failure data, the emerging methods to estimate blackout risk will founder because there will be no way to eliminate wrong hypotheses, develop sound theory, and validate methods and simulations. There can be optimism that the conflict between secrecy and research access can be resolved because the high-level methods of this paper inherently tend to ignore most of the detail and specifics of the data and therefore it should be technically feasible to filter the data to drastically reduce or eliminate its sensitivity. Continuing discussions between industry, national labs, government, engineering societies and researchers would be most helpful. Almost all the references of this paper directly or indirectly relied on access to data. I note with gratitude the essential role of the NERC Disturbance Analysis Working Group database summarizing North American blackout data in catalyzing and enabling all the work on the complex systems study of blackouts and cascading failure.

B. Develop insight, models and methods

Although a promising start has been made in several aspects of understanding and modeling cascading failure and large blackouts, much work remains to develop further insights and systematically test and validate ideas and methods. The critical loading behavior of blackout simulations is complicated [9] and the limitations and capabilities of high-level models such as branching processes are not yet established. It would be very useful to identify the parameters that most strongly control the proximity of the system to critical loading. These are the parameters that should be used to implement the mitigation of blackout risk and it might be feasible to monitor and control some of these parameters in operations. Practical metrics for risk analysis such as the propagation λ are starting to emerge but are not fully tested. The goal is a comprehensive risk analysis that is practical for cascading blackouts and further comparisons between real data, detailed simulations and high-level models need to be vigorously pursued.

After a transmission system upgrade, the patterns of use of the network can evolve to change the balance between maximizing use of the transmission system and reliability [29], [34]. For initial work in understanding how the slow complex dynamics of network upgrade and responses to recent blackouts shapes the blackout risk see [10], [12].

C. Improve cascading blackout simulation

In addition to the usual issue of the appropriate tradeoffs between model detail and simulation speed, there is uncertainty about which aspects of cascading failure should be modeled and what detail is required. Cascading failure is influenced by operational policies, software and human errors, upgrade and the responses to previous blackouts in addition to the varied and complicated ways in which failures interact via physical laws on the power grid. It remains to be seen to what extent these aspects can be modeled and how closely the statistics of real blackouts can be reproduced. There is also a need for simple, high-level models to explain the phenomena observed in the detailed simulations and in real data.

There are also challenges in running the simulations so that appropriate samples of initial conditions and probabilistic decisions are made and in developing ways to extract useful risk metrics from the results and to communicate the risks.

D. Can we directly monitor system reliability?

It seems that simulations will not be able to fully represent all the intricacies of real blackouts in quantifying reliability. Therefore, although simulations clearly have a vital role in planning and operations and in testing new methods, it is desirable to also develop methods of monitoring reliability directly from real power system data. One tactic is to first develop and test efficient methods of quantifying blackout risk from simulations and then adapt these methods to work on real data. The efficiency of the methods is important because practicality dictates that the monitoring quantify reliability quickly enough. For example, monitoring failure data for about one year to approximately determine the distribution of blackout sizes might be acceptable. Waiting many decades in order to empirically gather enough data for good statistics on the rare but most costly blackouts is not acceptable.

We speculate that the branching process parameter estimation methods, after they are more thoroughly validated on simulations, could be developed to directly monitor power system reliability with respect to cascading failure. One possibility is statistically estimating branching process parameters from transmission line outage records to compute the distribution of blackout size and then combining this with blackout cost to estimate the distribution of blackout risk. There are a number of challenges to be overcome to achieve this, but the current research does indicate a path towards this goal.

Power system reliability with respect to the more complicated and dependent failures has in the past been approached by studying selected real or simulated failures in detail and devising design and operating guidelines that position the power system a judicious margin away from likely failures. These guidelines are based on expert judgement and have proven effective in the past, but the methods do not measure reliability in a way that allows the tradeoffs involved in reliability investments to be quantified. That is, there is a difference between being able to recommend prudent measures that help reliability and being able to quantify the effect of reliability measures in terms of the costs of various sizes of blackouts. If methods to directly monitor and quantify power system blackout reliability are developed, this would enable new possibilities in quantifying these tradeoffs and determining the value of reliability investments.

References

[1] R. Adler, S. Daniel, C. Heising, M. Lauby, R. Ludorf, T. White, An IEEE survey of US and Canadian overhead transmission outages at 230 kV and above, IEEE Transactions on Power Delivery, vol. 9, no. 1, Jan. 1994, pp. 21 -39.

[2] G. Ancell, C. Edwards, V. Krichtal, Is a large scale blackout of the New Zealand power system inevitable?, Electricity Engineers Association 2005 Conference "Implementing New Zealands Energy Options", Aukland, New Zealand, June 2005.

[3] M. Anghel, K.A. Werley, A.E. Motter, Stochastic model for power grid dynamics, 40th Hawaii International Conference on System Sciences, Hawaii, January 2007.

[4] K.B. Athreya, P.E. Ney, Branching Processes, Dover NY 2004 (reprint of Springer-verlag Berlin 1972).

[5] J.Ø.H. Bakke, A. Hansen, and J. Kertész, Failures and avalanches in complex networks, Europhysics Letters, vol. 76, no. 4, pp. 717-723, 2006.

[6] R. Billington, R.N. Allan, Reliability evaluation of power systems (2nd ed.), New York: Plenum Press, 1996

[7] B.A. Carreras, V.E. Lynch, M. L. Sachtjen, I. Dobson, D. E. Newman, Modeling blackout dynamics in power transmission networks with simple structure, 34th Hawaii International Conference on System Sciences, Maui, Hawaii, January 2001.

[8] B.A. Carreras, V.E. Lynch, I. Dobson, D.E. Newman, Dynamics, criticality and self-organization in a model for blackouts in power transmission systems, Thirty-fifth Hawaii International Conference on System Sciences, Hawaii, January 2002.

[9] B.A. Carreras, V.E. Lynch, I. Dobson, D.E. Newman, Critical points and transitions in an electric power transmission model for cascading failure blackouts, *Chaos*, vol. 12, no. 4, December 2002, pp. 985-994.

[10] B.A. Carreras, V.E. Lynch, D.E. Newman, I. Dobson, Blackout mitigation assessment in power transmission systems, 36th Hawaii International Conference on System Sciences, Hawaii, 2003.

[11] B.A. Carreras, V.E. Lynch, I. Dobson, D.E. Newman, Dynamical and probabilistic approaches to the study of blackout vulnerability of the power transmission grid, 37th Hawaii International Conference on System Sciences, Hawaii, 2004.

[12] 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.

[13] 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.

[14] J. Chen, J.S. Thorp, M. Parashar, Analysis of electric power system disturbance data, Thirty-fourth Hawaii International Conference on System Sciences, Maui, Hawaii, January 2001.

[15] 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.

[16] Q. Chen, K. Zhu, J.D. McCalley, Dynamic decision-event trees for rapid response to unfolding events in bulk transmission systems, 2001 IEEE Porto Power Tech Proceedings, vol. 2, Portugal 2001.

[17] Q. Chen, J.D. McCalley, Identifying high risk n-k contingencies for online security assessment, IEEE Transactions on Power Systems, vol. 20, no. 2, May 2005, pp. 823-834.

[18] Q. Chen, C. Jiang, W. Qiu, J.D. McCalley, Probability models for estimating the probabilities of cascading outages in high-voltage transmission network, IEEE Transactions on Power Systems, vol. 21, no. 3, August 2006, pp. 1423-1431.

[19] I. Dobson, J. Chen, J.S. Thorp, B.A. Carreras, D.E. Newman, Examining criticality of blackouts in power system models with cascading events, 35th Hawaii International Conference on System Sciences, Hawaii, January 2002.

[20] I. Dobson, B.A. Carreras, D.E. Newman, A branching process approximation to cascading load-dependent system failure. 37th Hawaii International Conference on System Sciences, Hawaii, January 2004.

[21] I. Dobson, B.A. Carreras, D.E. Newman, Branching process models for the exponentially increasing portions of cascading failure blackouts, 38th Hawaii International Conference on System Sciences, January 2005, Hawaii.

[22] I. Dobson, K.R. Wierzbicki, B.A. Carreras, V.E. Lynch, D.E. Newman, An estimator of propagation of cascading failure, 39th Hawaii International Conference on System Sciences, January 2006, Kauai, Hawaii.

[23] I. Dobson, B.A. Carreras, D.E. Newman, A loading-dependent model of probabilistic cascading failure, Probability in the Engineering and Informational Sciences, vol. 19, no. 1, January 2005.

[24] 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.

[25] P. Guttorp, Statistical inference for branching processes, Wiley, NY, 1991

[26] R.C. Hardiman, M.T. Kumbale, Y.V. Makarov, An advanced tool for analyzing multiple cascading failures, Eighth International Conference on Probability Methods Applied to Power Systems, Ames Iowa, September 2004.

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

[28] Å.J. Holmgren, S. Molin, Using disturbance data to assess vulnerability of electric power delivery systems, Journal of Infrastructure Systems, December 2006, pp. 243-251.

[29] D.S. Kirschen, G. Strbac, Why investments do not prevent blackouts, UMIST (now University of Manchester) August 2003. available at www.ksg.harvard.edu/hepg/Standard_Mkt_dsgn/ Blackout_Kirschen_Strbac_082703.pdf

[30] D.S. Kirschen, D. Jawayeera, D.P. Nedic, R.N. Allan, A probabilistic indicator of system stress, IEEE Transactions on Power Systems, vol. 19, no. 3, 2004, pp. 1650-1657.

[31] H. Liao, J. Apt, S. Talukdar, Phase transitions in the probability of cascading failures, Electricity Transmission in Deregulated Markets, conference at Carnegie Mellon University, Pittsburgh PA USA Dec. 2004.

[32] D.P. Nedic, D.S. Kirschen, Discovering mechanisms of disturbance development, IREP Conference on Bulk Power System Dynamics and Control VI, Cortina D'Ampezzo, Italy, August 2004.

[33] 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.

[34] N.D. Reppen, Increasing utilization of the transmission grid requires new reliability criteria and comprehensive reliability assessment, Eighth International Conference on Probabilistic Methods Applied to Power Systems, Ames, Iowa USA September 2004.

[35] U.S.-Canada Power System Outage Task Force, Final Report on the August 14th blackout in the United States and Canada. United States Department of Energy and National Resources Canada, April 2004.

[36] Transmission reliability evaluation for large-scale systems (TRELSS): version 6.0 User's manual, EPRI, Palo Alto, CA: 2000. 1001035

[37] X. Weng, Y. Hong, A. Xue, S. Mei, Failure analysis on China power grid based on power law, Journal of Control Theory and Applications, vol. 4, no. 3, August 2006, pp. 235-238.

[38] R. Weron, I. Simonsen, Blackouts, risk, and fat-tailed distributions, in Practical Fruits of Econophysics, ed. H. Takayasu, Springer-Tokyo, 2005.

[39] K.R. Wierzbicki, I. Dobson, An approach to statistical estimation of cascading failure propagation in blackouts, CRIS, Third International Conference on Critical Infrastructures, Alexandria, Virginia, September 2006.

Appendix

We briefly summarize some of the available cascading failure simulations and indicate their modeling detail reported in publications in Table I. There are several models of cascading overloads of transmission lines representing the power grid at the level of DC load flow and LP dispatch of generation [9], [12], [15], [31]. Some special capabilities of these models are that the OPA model represents the slow complex dynamics of grid upgrade [12], the model developed at Cornell University represents hidden failures of the protection system [15], and the model developed using the PSA suite at Los Alamos National Laboratory represents the timing of blackout events and operator actions [3]. The state of the art in terms of model detail is to use an AC load flow and have some approximate representations of protection, operator actions and voltage collapse as in the Manchester model [30] and TRELSS [36], [26]. The Manchester model has been run on industrial data and the cascading mode of TRELSS is used by industry to identify cascading failure problems.

	OPA	hidden failure	Manchester	CMU	PSA	TRELSS
overloads	X	X	X	X	X	X
generator redispatch	X	X	X		X	X
hidden failure		X		X		
protection group						X
AC network		X	X			X
generator trip						X
voltage collapse			X			X
transient stability			X			
under frequency load shed			X			X
islanding	X		X			X
operator response			X		X	X
blackout time intervals & repair					X	
load increase & grid upgrade	X					
approx. max number of buses	400	300	1000	2500	?	13000
reference	[9]	[15]	[30]	[31]	[3]	[36], [26]

TABLE I

Modeling detail in cascading failure simulations