

Towards Quantifying Cascading Blackout Risk

Ian Dobson Kevin R. Wierzbicki Janghoon Kim Hui Ren
ECE Department, University of Wisconsin, Madison, WI 53706 USA

Abstract

Blackouts become widespread by initial failures propagating in a diverse and intricate cascade of rare events. We describe this complicated cascade using a bulk probabilistic model in which the initial failures propagate randomly according to a branching process. The branching process parameters can be statistically estimated from observed data or simulations and then used to efficiently predict the probability distribution of blackout size. We review the current testing of these methods on simulations and observed data and discuss the next steps towards achieving verified and practical methods for quantifying cascading failure of electric power systems. The ability to efficiently quantify cascading blackout risk from observed data and simulations could offer new ways to monitor power transmission system reliability and quantify the reliability benefit of proposed improvements.

1 Introduction

It would be very useful to be able to efficiently quantify overall blackout risk from simulated or real power system data. Established analytic methods of power system risk analysis can model the detail of some likely and foreseen combinations of failures and estimate their risk. This is very useful in finding and mitigating likely failures, but it does not address quantifying the overall risk of large cascading blackouts, in which there is combinatorial explosion of potential rare, unforeseen, and interacting events ranging from diverse power system physical effects through software failures to deficiencies in planning, operation, organization, and maintenance. Here the term “failure” includes outages in which the component or process is intact but temporarily unavailable to transmit power or function properly as well as outages in which the component or process is damaged or impaired.

Studying the intricate details of particular blackouts [30], in addition to quickly reinforcing the case for the complexity of these events, also provides a useful method for finding and mitigating weaknesses in the power system. However, larger blackouts are infrequent events in developed economies, and, after a large blackout, it is difficult to as-

certain whether the blackout should be attributed to bad practice or bad luck, especially with everyone’s attention riveted on only one sample from the huge number of potential blackouts. A detailed analysis of the chain of events after the blackout is useful in suggesting specific weaknesses that can be rectified, but gives little guidance on the overall problem of whether society is rationally balancing the blackout risks with the costs of investing in increased reliability. Quantifying the overall blackout risk would allow this balancing by putting an approximate value on reliability.

The most straightforward way to estimate the probabilities of various sizes of blackouts is simply to wait a long time to observe enough blackouts to get good estimates of blackout probabilities. But in practice the wait is too long for most purposes because the large blackouts occur too rarely. (Observed blackouts in several countries have power law regions in the distribution of blackout size [17] and empirical estimation of these distributions requires many samples.) Moreover, the power system upgrades over time so that the observed statistics represent the average over a considerable time of an evolving system. Therefore, we aim for efficient methods of estimating blackout risk that could work in a time scale of about one year. While there are certainly challenges in estimating the direct and indirect costs of blackouts and both cost and probability of blackouts are needed to estimate blackout risk, this paper focusses on estimating the probability of blackouts, or, more precisely, the probability distribution of blackout size. The probability distribution of blackout size can be combined with the cost as a function of blackout size to yield the distribution of blackout risk as a function of blackout size. We maintain that problem of blackout risk is better framed in terms of the distribution of risk of various sizes of blackouts rather than the risk of blackouts in general [17].

We suggest that any practical approach to efficient estimation of blackout probabilities must be a bulk “top-down” statistical approach that incorporates a verified understanding of cascading failures in power systems. This is a departure from methods of risk analysis that rely on detailed analysis of enumerated interactions, but it is complementary to these detailed analyses. The purpose of this paper is to review progress in developing and testing a bulk statistical method to estimate the probability of cascading blackouts.

Specialized simulations can sample a subset of the intricacies of large cascading blackouts [2, 4, 7, 19, 23, 25, 26, 27, 28, 31] and, while they may never be able to capture all the interactions in blackouts, simulations are vital in testing statistical methods for monitoring the power system. Moreover, efficient estimation of blackout probabilities from data produced by cascading failure simulations makes efficient use and leverages understanding of the simulation results. In particular, reducing the number of simulation runs would enable the effects of proposed reliability upgrades to be quickly assessed.

Our top-down modeling views cascading as a random initial disturbance followed by a random propagation of failures. The outcome of each cascade is probabilistic, but its statistics are governed by the size of the initial disturbance and the average tendency for the failures to propagate. A simple way to capture this mathematically is to use a branching process model.

Why consider a branching process model?

1. Branching processes are a standard model for cascades in many other subjects, including genealogy, cosmic rays, and epidemics [20, 3]. This makes branching processes an obvious first choice for modeling cascading failures.
2. Observed [6, 17] and simulated [4, 7, 27, 24] power system data shows qualitative features such as distributions of blackout sizes with near power law regions and criticality that can be produced by branching processes as illustrated in section 2.1 and [12].
3. The CASCADE model [15] is another high-level probabilistic model of cascading failure. CASCADE is well approximated by Galton-Watson branching process models in suitable parameter ranges [12].
4. Branching processes are simple and tractable models and it is good to test simple models first. In particular, branching process models can be tested against real and simulated power system data as described in this paper.

If branching processes are useful models of cascading processes, this would open up several opportunities. In addition to providing the essential understanding of the overall features of cascading, the probability distribution of blackout size could be quantified more efficiently from much smaller samples of cascades. One reason is that estimating the parameters of a branching process model and then computing the distribution of blackout size using the model requires much smaller samples than directly estimating the distribution of blackout size exhaustively, especially since large cascading blackouts are rare events. The smaller number of samples is important when observing cascades in the power system because smaller samples enable a shorter observation time that can make the approach practical. The

smaller number of samples is similarly important when simulating cascading failure because smaller samples enable shorter run times.

However, we still need to exhaustively determine the distribution of blackout size by simulation and observation in order to test whether branching process models are valid and applicable. This paper shows how to test branching process models on observed or simulated power system data. A large number of cascades are observed or simulated in order to exhaustively determine the distribution of blackout size. Then the same data is used to estimate the branching process model parameters and compute the distribution of blackout size with the branching process model and see how well this matches the empirical distribution. If the match is acceptable, then the branching process model is useful for determining the distribution of blackout size.

Throughout the paper there are two measures of blackout size of interest to monitor the progress and outcome of the cascade; namely, number of transmission lines failed and the load shed. The lines failed are more easily tracked and analyzed and are of internal interest to the utilities and system operators. Moreover, lines may fail in “precursor” events with no load power shed. On the other hand, load power shed matters to the utilities, society, industry and government.

Section 2 reviews and explains several branching process models and the statistical estimation of their parameters. Section 3 tests branching process models on data from a power systems simulation and section 4 tests a branching process model on observed power system cascades. Section 5 examines issues and challenges in further work towards quantifying cascading failure risk in power systems and section 6 concludes the paper.

2 Branching processes

Branching processes have long been used in a variety of applications to model cascading processes [20, 3], but their application to the risk of cascading failure is novel [12, 13]. The branching processes of primary interest produce failures in stages starting from some initial distribution of failures. There are several types of branching processes that we describe below: A Galton-Watson process has discrete numbers of failures such as numbers of transmission lines failed. A continuous-state branching process has a continuously varying amount of failure in each stage such as load power shed in each stage. Other branching processes evolve in continuous time.

2.1 Galton-Watson branching process

The Galton-Watson branching process gives a probabilistic model of the number of failures. There are a random




Fig. 1: Example of failures produced in stages by a Galton-Watson branching process. Each “parent” failure independently has a random number of “child” failures in the next stage.

number of initial failures that then propagate randomly to produce subsequent failures in stages. The mean number of initial failures is 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 $0,1,2,3,\dots$ of failures (“child” failures) in the next stage as illustrated in Figure 1. The children failures then become parents to produce the next generation and so on. If the number of failures in a stage becomes zero, the cascade stops. The mean number of child failures for each parent is the parameter λ . λ quantifies the average 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 [9, 10, 12]. There are general arguments supporting the choice of Poisson distributions [9, 10]. The Poisson distribution is a good approximation when each failure propagates to a large number of components so that each parent failure has a small, fairly uniform probability of independently causing child failures in a large number of other components. This assumption seems reasonable for cascades in power systems, especially in the initial portions of the cascade when there are many unfailed components that are stressed by the failed components.

There are assumed to be N transmission lines and there are Z_0 initial failures. Line failures occur in stages with Z_n the number of failures in stage n and Y_n the total number of failures up to and including stage n .

$$Y_n = Z_0 + Z_1 + Z_2 + \dots + Z_n \quad (1)$$

The process saturates when $S \leq N$ lines fail. Each of the Z_n failures in stage n independently causes a further number of failures in stage $n+1$ according to a Poisson distribution with mean λ , except that if the total number of failures exceeds S , then the total number of failures is limited to S . That is, the j th failure in stage n causes $Z_{n+1}^{[j]}$ failures in stage $n+1$ according to the Poisson distribution and the total number of failures in stage $n+1$ is

$$Z_{n+1} = \min \left\{ Z_{n+1}^{[1]} + Z_{n+1}^{[2]} + \dots + Z_{n+1}^{[Z_n]}, S - Y_n \right\}, \quad (2)$$

where $Z_{n+1}^{[1]}, Z_{n+1}^{[2]}, \dots, Z_{n+1}^{[Z_n]}$ are independent. (A different form of saturation is described in [12, 16].)

The modeling of saturation is essential for λ near or exceeding 1 because then there is a significant probability of cascading to all N lines of the system failing (total blackout). That is, practical application of the branching process requires an upper limit $S \leq N$ of components failing for λ near or exceeding 1. Moreover there may be saturation effects encountered before all N components fail and then S can be chosen less than N , either because the number failures tend to bunch at S or because the modeling is uncertain or variable after S components fail. These saturation effects are not well understood and it remains to be seen how or whether saturation effects are significant for cascading failure in practical systems.

There are several variants of analytic formulas for the distribution of the total number of failures Y [13], depending on the assumed distribution of initial failures Z_0 , whether saturation is assumed, and whether the distribution of Y is conditioned on Y being nonzero. If the initial failures Z_0 are distributed according to a Poisson distribution with mean θ and saturation S is assumed, then the total number of failures conditioned on Y being nonzero is distributed according to a saturating generalized Poisson distribution:

$$P[Y = r \mid Y > 0] = \begin{cases} \theta(r\lambda + \theta)^{r-1} \frac{e^{-r\lambda - \theta}}{r!(1 - e^{-\theta})}; & 1 \leq r < S \\ 1 - \sum_{s=1}^{S-1} P[Y = s]; & r = S \end{cases} \quad (3)$$

Some qualitative features of the saturating generalized Poisson distribution (3) are shown in Figure 2. For small propagation $\lambda = 0.1$, the distribution drops off sharply and there is negligible probability of a large number of line failures. As λ increases, the distribution extends towards a large number of line failures until at $\lambda = 1$, there is a power law region of slope approximately -1.5 . Moreover, a graph of

average blackout size as a function of λ (not shown here) shows a change in gradient at $\lambda = 1$, indicating criticality at $\lambda = 1$. The example in Figure 2 assumes saturation at $S = 1000$ lines. If an unsaturating branching process would have proceeded beyond 1000 lines failures, then the saturating branching process records that event as 1000 lines failures. The probability of 1000 lines failing increases as the system becomes more stressed. At $\lambda = 1$ the probability of 1000 lines failed is 0.025 and this increases to 0.797 at $\lambda = 2$. The very low probability of an intermediate number of failures at $\lambda = 2$ is due to almost all the cascades that have an intermediate number of failures continuing to grow until there are 1000 lines failed.

If there is an arbitrary distribution of nonzero initial failures $P[Z_0 = z_0]$ for $z_0 = 1, 2, 3, \dots$, and no saturation, then the total number of failures is distributed according to a mixture of Borel-Tanner distributions:

$$P[Y=r] = \sum_{z_0=1}^r P[Z_0=z_0] z_0 \lambda (r\lambda)^{r-z_0-1} \frac{e^{-r\lambda}}{(r-z_0)!} \quad (4)$$

The neglect of saturation implies that (4) is only valid for subcritical $\lambda < 1$.




Fig. 2: Probability distributions of total number of line failures Y according to the generalized Poisson distribution with saturation at $S = 1000$ lines for three values of propagation λ . The distribution has an approximate power law region at criticality when $\lambda = 1$. The probability of 1000 lines failing is 0.025 for $\lambda = 1$, and 0.797 for the supercritical case $\lambda = 2$.

The simulation is run or data are observed K times to produce K independent realizations of the cascade. Since a cascades with $Z_0 = 0$ are discarded, all cascades have $Z_0 > 0$ and all statistics are conditioned on the nontrivial start of a cascade. The failures in the k th run are written as $Z_0^{(k)}$,

$Z_1^{(k)}, Z_2^{(k)}, Z_3^{(k)}, \dots$. The data can be tabulated as follows:

	stage 0	stage 1	stage 2	stage 3	\dots
run 1	$Z_0^{(1)}$	$Z_1^{(1)}$	$Z_2^{(1)}$	$Z_3^{(1)}$	\dots
run 2	$Z_0^{(2)}$	$Z_1^{(2)}$	$Z_2^{(2)}$	$Z_3^{(2)}$	\dots
run 3	$Z_0^{(3)}$	$Z_1^{(3)}$	$Z_2^{(3)}$	$Z_3^{(3)}$	\dots
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
run K	$Z_0^{(K)}$	$Z_1^{(K)}$	$Z_2^{(K)}$	$Z_3^{(K)}$	\dots

Define the cumulative number of failures in run k up to and including stage n as

$$Y_n^{(k)} = Z_0^{(k)} + Z_1^{(k)} + Z_2^{(k)} + \dots + Z_n^{(k)} \quad (6)$$

Each run has a stage at which the number of failures is zero and remains zero for all subsequent stages, either because the cascade dies out, or the cascade has saturated. The number of failures at which saturation occurs is $S \leq N$. Define

$$s(k, S) = \max\{n \mid Y_n^{(k)} < S \text{ and } Z_{n-1}^{(k)} > 0\} \quad (7)$$

That is, $s(k, S)$ is either the first stage at which there are zero failures or the last stage before a total of S failures.

We define an estimator of λ as

$$\begin{aligned} \hat{\lambda} &= \frac{\sum_{k=1}^K (Z_1^{(k)} + Z_2^{(k)} + \dots + Z_{s(k,S)}^{(k)})}{\sum_{k=1}^K (Z_0^{(k)} + Z_1^{(k)} + \dots + Z_{s(k,S)-1}^{(k)})} \\ &= \frac{\sum_{k=1}^K (Y_{s(k,S)}^{(k)} - Z_0^{(k)})}{\sum_{k=1}^K Y_{s(k,S)-1}^{(k)}} \end{aligned} \quad (8)$$

The estimator $\hat{\lambda}$ is asymptotically unbiased as $K \rightarrow \infty$ if the offspring distribution is Poisson. This asymptotic unbiasedness property shows that the estimator $\hat{\lambda}$ is an improved variant of the estimator of λ proposed in [14].

We determined the bias and variance of $\hat{\lambda}_s$ numerically for a saturating branching process with one initial failure by computing $\hat{\lambda}$ 1000 times and computing the sample mean $\mu(\hat{\lambda})$ and standard deviation $\sigma(\hat{\lambda})$ of $\hat{\lambda}$.

For saturation $S = 20$, $0 < \lambda < 2$, and number of runs $10 \leq K \leq 1000$, $|\mu(\hat{\lambda}) - \lambda| \leq 0.05$ and

$$\sigma(\hat{\lambda}) \leq \frac{0.81}{\sqrt{K}}. \quad (9)$$

For saturation $S = 100$, $0 < \lambda < 2$, and number of runs $10 \leq K \leq 150$, $|\mu(\hat{\lambda}) - \lambda| \leq 0.07$ and

$$\sigma(\hat{\lambda}) \leq \frac{0.51}{\sqrt{K}}. \quad (10)$$

For example, $K = 25$ runs gives $\sigma(\hat{\lambda}) \leq 0.1$.

When significant saturation effects are encountered, it is essential to use the estimator (8) to avoid underestimating λ [14]. However, when the branching process is subcritical ($\lambda < 1$) and there is no saturation or negligible saturation ($S = \infty$), $\hat{\lambda}$ reduces to the standard Harris estimator for λ [20, 36, 11, 18]:

$$\hat{\lambda} = \frac{\sum_{k=1}^K (Z_1^{(k)} + Z_2^{(k)} + \dots)}{\sum_{k=1}^K (Z_0^{(k)} + Z_1^{(k)} + \dots)} = \frac{\sum_{k=1}^K (Y_\infty^{(k)} - Z_0^{(k)})}{\sum_{k=1}^K Y_\infty^{(k)}} \quad (11)$$

The Harris estimator (11) is an asymptotically unbiased maximum likelihood estimator [20, 36]. The asymptotic standard deviation of the Harris estimator can be worked out using the methods of [36], in the case of Poisson offspring distribution with one initial failure, to be

$$\sigma(\hat{\lambda}) \sim \frac{\sqrt{\lambda(1-\lambda)}}{\sqrt{K}} \leq \frac{0.5}{\sqrt{K}} \quad (12)$$

The Harris estimator (11) is intuitive: Think of “parent” failures in each generation giving rise to “child” failures in the next generation. The child failures then become parents to produce the next generation and so on. Then λ is the average family size; that is, the average number of child failures for each parent. Since Z_0, Z_1, \dots are all parent failures and Z_1, Z_2, \dots are all child failures, the Harris estimator (11) is simply the total number of children in all the cascades divided by the total number of parents in all the cascades.

If the initial failures Z_0 are approximated by a Poisson distribution conditioned on a nonzero number of failures, then an estimate of the mean initial failures $\hat{\theta}$ can be obtained by solving

$$\frac{\hat{\theta}}{1 - e^{-\hat{\theta}}} = \frac{1}{K} \sum_{k=1}^K Z_0^{(k)}. \quad (13)$$

It is straightforward to derive (13) by maximum likelihood methods or the method of moments.

2.2 Continuous state branching process

We summarize the application of continuous state branching processes to model the cascading of load shed in [35]. Continuous state branching processes have a similar theory to Galton-Watson branching processes, but the computations of the distribution of the total load shed are more difficult. See [21, 29] for systematic accounts of continuous state branching processes. Throughout this subsection, we assume the subcritical case $\lambda < 1$ and no saturation.

There is an initial amount of load shed Z_0 in stage 0 that is given by a probability density function (pdf) $G(z)$ and the

load shed amounts Z_1, Z_2, \dots in stages 1, 2, \dots are produced using an offspring distribution $H(z)$. $H(z)$ is defined to be the pdf of load shed in any stage if the load shed in the preceding stage is 1. We write Z for a random variable with pdf $H(z)$. The expected value of Z is λ .

For the subsequent stages $n = 1, 2, 3, \dots$, the load Z_n shed in stage n has pdf that is the convolution of $H(z)$ with itself Z_{n-1} times, or, equivalently, the pdf of the sum of Z_{n-1} independent copies of Z . These pdfs are computed in the frequency domain using their cumulant generating functions (cgf’s). The cgf $h(s)$ of the offspring distribution is the negative logarithm of the Laplace transform of $H(z)$:

$$h(s) = -\ln \int_0^\infty e^{-sz} H(z) dz = -\ln Ee^{-sZ}.$$

and, since we choose $H(z)$ to be a gamma distribution, $h(s)$ has the form

$$h(s) = \frac{\lambda^2}{\sigma_{\text{off}}^2} \ln \left(1 + s \frac{\sigma_{\text{off}}^2}{\lambda} \right). \quad (14)$$

The distribution of the total load shed Y can be computed from the offspring distribution as follows [35]. Computer algebra is used to perform the symbolic operations. Let Z_0 have a cgf of the form

$$m(s) = \frac{\theta^2}{\sigma_{\text{init}}^2} \ln \left(1 + s \frac{\sigma_{\text{init}}^2}{\theta} \right) \quad (15)$$

so that Z_0 also has a gamma distribution. Then the cgf of Y is

$$k(s) = m(k_\bullet(s)) \quad (16)$$

where $k_\bullet(s)$ satisfies

$$k_\bullet(s) = s + h(k_\bullet(s)). \quad (17)$$

The implicit equation (17) can be solved for $k_\bullet(s)$ by the Lagrange inversion method [33] and then the pdf $K(x)$ of the total load shed Y can be obtained as the inverse Laplace transform of $e^{-k(s)}$ using the Post-Widder method [34].

Some cascades may have negligible or zero load shed in initial stages. These initial stages are discarded from the data so that all nontrivial cascades start with $Z_0 > 0$. (Trivial cascades with no load shed are also discarded.) Then the estimator $\hat{\lambda}$ is the same as (8) except that saturation is neglected ($S = \infty$), and we impose a maximum of 10 stages so that (7) becomes

$$s(k, S) = \max\{n \mid n \leq 10 \text{ and } Z_{n-1}^{(k)} > 0\} \quad (18)$$

The condition on $Z_{n-1}^{(k)}$ may be omitted without changing the value of $\hat{\lambda}$ in (8).

2.3 Continuous time branching process

There are many interesting generalizations of basic Galton-Watson and continuous state branching process, notably those that evolve in continuous time instead of stages and those with variable propagation. In this paper we focus on basic Galton-Watson and continuous state branching processes because of their simplicity and their basic approach to the timing of failures.

Many of the current cascading failure simulations do not model the detailed timing of failures (exceptions include [2] and the sequences of discrete events simulated in voltage collapse simulations [32]) and naturally compute failures in stages. Observed failure data can be grouped into stages by simple methods such as grouping together failures occurring closely in time as explained in section 4. The stages occur in a discrete time that advances by one unit with each successive stage but has no correspondence with the real time of the failures. These simple approaches avoid explicitly modeling the time evolution of the failures. However, it remains a possibility to model the time evolution of the failures in a branching process. For example, a continuous time Markov branching process is used in [13] to try to fit the time evolution of failures in large blackouts.

3 Estimating λ and blackout size distribution from simulations

We test branching process models using simulated power system cascades as follows:

1. Run the power system simulation a large number of times to produce cascading failure data in stages and the blackout size.
2. Estimate the empirical distribution of the blackout size. This is a brute force evaluation that relies on the large number of runs in step 1.
3. Use the cascading failure data in stages to estimate the parameters of the branching process, such as the average propagation λ . (The large number of runs will produce accurate parameter estimates, so that we do not have to address in this testing of the branching process model the tradeoff that will arise in practice when the branching process model is applied to a much smaller number of runs and some parameter inaccuracy is tolerated.)
4. Use the estimated branching process parameters and the branching process theory to predict the distribution of blackout size.
5. Compare the empirical blackout size distribution obtained by brute force with the blackout size distribution obtained via the branching process. The extent that these distributions match indicates the success of

the branching process model in capturing the overall features of the cascading process in the power system simulation.

3.1 OPA cascading failure simulation

The OPA model produces cascading failures in stages resulting from a random initial set of line failures [4, 17]. The power transmission system is modeled using DC load flow and LP generator dispatch, and cascading line overloads and failures are represented. Each cascade produces a number of lines failed and the load shed in each stage of the cascade. The power system is assumed to be fixed with no transmission line upgrade process.

For each case considered, OPA was run so as to produce 5000 cascades with a nonzero number of line failures. These 5000 runs yield both line failure data and load shed data in the form (5). All our statistics are conditioned on a nonzero number of line failures.

3.2 Estimating line outage probability distribution

This subsection tests whether the line outages produced by the OPA simulation are governed by a Galton-Watson branching process with Poisson offspring distribution. The results are very similar to those in [14], but here we use the improved estimator $\hat{\lambda}$ in (8). The parameters are the mean of the initial distribution of failures θ and the mean of the offspring distribution λ . The parameter estimators $\hat{\lambda}$ and $\hat{\theta}$ were obtained using equations (8) and (13) with saturation $S = 100$. According to (10), the number of runs $K = 5000$ gives a negligible standard deviation of $\sigma(\hat{\lambda}) < 0.01$. For each case, the probability distribution of line failures was predicted using the estimates $\hat{\lambda}$ and $\hat{\theta}$ in the saturating generalized Poisson distribution formula (3). This tests the validity of the branching process model underlying (3) and the estimates $\hat{\lambda}$ and $\hat{\theta}$.

The first three cases used the IEEE 118 bus system at average load levels of 0.9, 1.0, and 1.3 times the base case loading. (The OPA parameters (explained in [4]) are $\gamma = 1.67$, $p_0 = 0.0001$ and $p_1 = 1$.) The estimated parameters are shown in Table 1. Good matches between the empirical and estimated distributions are shown in Figures 3-5.

The last two cases used the 190 bus tree-like test system [4] at average load levels of 1.0 and 1.2 times the base case loading. (The OPA parameters are $\gamma = 1.94$, $p_0 = 0.005$ and $p_1 = 0.15$.) The estimated parameters are shown in Table 1. The matches between the empirical and estimated distributions are shown in Figures 6 and 7. The match is not as good for the 190 bus system for 1.2 times the base case loading. The branching process captures the general trend of the empirical distribution, but does not capture the peaks at 15 and 30 lines. It is conceivable that the




Fig. 3: IEEE 118 bus system with loading factor 0.9. Distribution of line outages estimated with branching process (solid line) compared with OPA empirical distribution (dots). Note the log-log scales.

peaks are caused by saturation effects or by nonuniformity in the (artificially regular) design of the 190 bus tree-like test system.

The results suggest that good predictions of the probability distributions of the number of line failures can be obtained as long as saturation effects are not significant. We do not understand how to accurately model the saturation effects at present. The ability to predict the probability distribution of the number of line failures in non saturating cases supports the applicability of branching models to cascading failure before saturation is reached and the usefulness of the estimate $\hat{\lambda}$ of failure propagation. $\hat{\lambda}$ can be obtained much more efficiently than empirical probability distributions of line failures obtained by brute force.

Table 1: Estimators for line failure data produced by OPA

power system	loading factor	$\hat{\theta}$	$\hat{\lambda}_s$
IEEE 118 bus	0.9	1.10	0.19
IEEE 118 bus	1	1.66	0.41
IEEE 118 bus	1.3	12.20	0.44
tree-like 190 bus	1.0	1.49	0.53
tree-like 190 bus	1.2	6.21	0.62

3.3 Estimating load shed pdf

This subsection summarizes testing in [35] that tests whether the load shed produced by the OPA simulation is governed by a continuous state branching process with a gamma initial distribution and a gamma offspring distribution. The load shed is measured as a fraction of the total load so that the maximum load shed possible is 1, or total blackout.

The estimated propagation $\hat{\lambda}$ and other parameters of the




Fig. 4: IEEE 118 bus system with loading factor 1.0. Distribution of line outages estimated with branching process (solid line) compared with OPA empirical distribution (dots).




Fig. 5: IEEE 118 bus system with loading factor 1.3. Distribution of line outages estimated with branching process (solid line) compared with OPA empirical distribution (dots).




Fig. 6: Tree-like 190 bus system with loading factor 1.0. Distribution of line outages estimated with branching process (solid line) compared with OPA empirical distribution (dots).




Fig. 7: Tree-like 190 bus system with loading factor 1.2. Distribution of line outages estimated with branching process (solid line) compared with OPA empirical distribution (dots).




Fig. 8: IEEE 118 bus system with loading factor 0.85. Probability density function of fraction of load shed Y . PDF estimated from branching process (dashed line) compared with empirical PDF from OPA simulation (dots).

gamma distributions are estimated for the OPA data. The estimated propagation $\hat{\lambda}$ at each load level is shown in the second column of Table 2. As expected, $\hat{\lambda}$ increases with loading. All cases considered are subcritical ($\lambda < 1$) as required in subsection 2.2.

Figure 8 compares the empirical and estimated PDFs for loading factor 0.85, and Figure 9 compares the empirical and estimated PDFs for loading factor 1.0. The blackout size is plotted on a log scale over two decades, from a small blackout $Y = .01$ (shedding of 1% of total load) to $Y = 1$ (shedding of 100% of total load and total blackout).

3.4 Comparing λ estimated from line and load shed data

There are several ways to think of the cascading processes that occur during a blackout, but one possible way is to think of a single cascading process that is monitored by two different measurements, lines failed and load shed. From this point of view, it would be natural for the propagation




Fig. 9: IEEE 118 bus system with loading factor 1.0. Probability density function of fraction of load shed Y . PDF estimated from branching process (dashed line) compared with empirical PDF from OPA simulation (dots).

λ of the cascade measured by line failures to be the same as the propagation λ of the cascade measured by the load shed. (Note that since λ is a ratio of similar quantities, it is dimensionless and λ estimates using different quantities can be directly compared.) Using the same OPA cascades, Table 2 compares $\hat{\lambda}$ computed from the load shed with $\hat{\lambda}$ computed from the lines failed. The $\hat{\lambda}$ for load shed matches well with the $\hat{\lambda}$ for line failures. This result is consistent with the lines failed and the load shed reflecting a common cascading process.

This result that the propagation of line failures and load shed agree is obtained here for only a few cases, but if the result is repeated for other simulations and cases, it would be useful because then one could use the monitoring of the propagation of line failures to predict the propagation of load shed.

Table 2: Estimated propagation $\hat{\lambda}$ from load shed and line failure data produced by OPA on IEEE 118 bus system

loading factor	load shed $\hat{\lambda}$	line failures $\hat{\lambda}$
0.85	0.128	0.115
0.9	0.159	0.188
0.95	0.264	0.288
1.0	0.429	0.430

4 Estimating line outage distribution from observed data

This section shows how to test branching process models on observed line outage data.

4.1 Observed line outage data

We processed one decade of transmission line fault data from a regional electric power system with approximately 100 buses and 200 high voltage lines. The voltage levels considered are 220 kV and 500 kV; outages at lower voltage levels are not considered because of the potential number of unrecorded cases. There are several types of line outages in the data, including three phase and single phase and outages with successful or unsuccessful auto-reclosing. In this initial work processing the data, both voltage levels and all types of line outages are regarded as the same and the detailed causes of the line outages (line fault, busbar fault, or other fault or operations) are neglected. This is consistent with an initial, “top-down” analysis of the overall cascading process. Large flashover events in the data with approximately 260 outages over two days are neglected because they lack time tags. The data for each transmission line outage include the time (to the nearest minute), voltage level, and the auto-recloser’s action.

4.2 Grouping outages into cascades and stages

We group the line outages first into different cascades, and then into different stages within each cascade using a simple method based on the time of the outages. Since operator actions are usually completed within one hour, we assume that successive outages separated in time by more than one hour belong to different cascades; since transients or auto-recloser actions are completed within one minute, we assume that successive outages in a given cascade separated in time by more than one minute are in different stages within that cascade.

Table 3: Total number of outages in each stage

stage	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
outages	296	45	18	14	10	3	1	1	1	1	1	1	2	1	1	0

Table 3 is obtained by summing over all the 226 cascades the number of outages in each stage. The initial outages are the 296 outages in stage 0. The probability distribution of the number of initial outages is shown in Fig.10a. The distribution in Fig.10a has a peak at 6 outages that prevents it being well fit by a Poisson distribution. One reason for the peak is that some cascades are initiated by a bus outage, and the relay trips off all transmission lines connected to that bus simultaneously at the start of the cascade.




Fig. 10: Probability distributions of number of outages. (a) Initial outages from data; (b) Total number of outages estimated using branching process (line) and from empirical observed data (dots).

4.3 Estimating λ and the distribution of outages

Assuming a Galton-Watson branching process model with no saturation, we use (11) to estimate λ based on the data in table 3:

$$\hat{\lambda} = \frac{45+18+14+10+3+1+1+1+1+1+2+1+1+0}{296+45+18+14+10+3+1+1+1+1+1+1+2+1+1} = 0.25$$

That is, each outage produces an average of $\lambda = 0.25$ outages in the next stage. This result is insensitive to the grouping of outages into stages (redefining the minimum time between successive outages in different stages to be 2 minutes and recomputing λ yields $\lambda = 0.24$).

To test how well the branching process model describes the data, we use the branching process with $\lambda = 0.25$ and the empirical distribution of initial outages to predict the distribution of the total number of outages Y using (4), and compare this with the distribution of the total number of outages directly obtained from the data. The comparison is shown in Figs. 10b and 11. A chi-squared goodness-of-fit test shows that the distributions are consistent at the 5% confidence level (the test groups together 5 or more outages). A heavy tail in the distribution of the total number of line outages is also observed in North American data in [8, 1], but our data has a heavier tail than [8, 1].

Observing outages for one year in this power system would yield an average of 25 cascades. To show how accurately λ could be estimated from one year of data, we took 9 non-overlapping random samples of 25 cascades and estimated λ for each sample of 25 cascades. A typical result is that the estimated λ has a standard deviation of 0.14. That is, assuming normality, an estimate of λ from one year of data lies within 0.14 of the true value about 68% of the time. This accuracy can be improved by collecting data over a longer time or over a larger region.




Fig. 11: Distribution of total number of line outages estimated using branching process (line) and from empirical observed data (dots); this log-log plot shows a heavy tail.

To summarize, the line failures are grouped into cascades and stages according to the failure times and we estimate the distribution of the initial number of failures and the propagation λ of the failures and hence estimate the distribution of the total number of line failures using a Galton-Watson branching process model. For this observed data, the empirical distribution of the total number of line outages is well approximated by the initial line outages propagating according to a branching process with parameter λ . Estimating λ requires significantly less data than estimating the heavy tail of the empirical distribution so that the distribution of blackout size may be efficiently estimated from data observed over a much shorter time. The efficiency of this estimation opens up possibilities of practical monitoring of power system reliability based on direct observations.

5 Discussion and future directions

This section considers some challenges to be addressed in modeling cascading failure with branching processes and future work.

5.1 Efficiency

Methods for quantifying cascading failure become practical when they require a modest number of cascades in a sample. There is a tradeoff between an efficiently small number of cascades in a sample and the accuracies of the estimated branching process parameters and the blackout size distribution based on these estimates.

Formula (12) for the asymptotic standard deviation of the estimator $\hat{\lambda}$ suggests that a standard deviation in $\hat{\lambda}$ of 0.1 requires about 25 cascades in the sample if there is one initial failure. The line outages observed in the 200 line power

system of section 4 occur at about 25 cascades per year and each year of data yields a standard deviation in $\hat{\lambda}$ of 0.14. Halving the standard deviation would require quadrupling the number of cascades in the sample to 100, either by observing data over four years or by monitoring the line outages over a larger portion of the power system with 800 lines. That is, improving the accuracy of $\hat{\lambda}$ requires either waiting longer and having less time resolution or gathering statistics over a wider area with a loss of spatial resolution. It is inherent in the method that the estimated value $\hat{\lambda}$ is averaged over the time period and spatial extent over which the cascades were observed.

Describing these temporal and spatial limitations in quantifying cascading failure requires further experience, especially with observed data, before firm conclusions can be drawn. However, it seems that the branching process estimation methods will offer substantial improvements in efficiency over estimating blackout probability distributions exhaustively. In particular, in North America, the largest blackouts occur on a time scale of decades and waiting to accumulate enough statistical data on these blackouts is much slower than using branching process models that could give approximate results on a time scale of about one year.

5.2 Elaborations in modeling and data processing

In this paper, we start with the simplest branching process models and estimate a few model parameters without distinction between types of outages and averaged over the duration of the cascade, the portion of the network studied, and the sample of cascades. Parsimony in parameters and simplicity of modeling are desirable, but it is not clear whether the branching process models require elaboration, or what is the best processing of data before the branching process model is applied. For example, it is not clear in processing line outage data whether a higher voltage line outage should be given more weight than a lower voltage line outage. Another example is that if the power system stress varies over the day, it might be desirable to accumulate data for high stress and low stress periods separately. Experimenting with the processing is needed.

We are applying a model of random propagation that is uniform over the duration of the cascade with parameters estimated by data averaged over the cascade. There could be changes in the cascading phenomenon as the cascade proceeds. For example, it is conceivable that λ could decrease over the duration of the cascade as blackout inhibition or saturation effects start to apply. The current calculations for the Galton-Watson branching process do allow for estimation of λ before a saturation limit of S is reached, but we do not yet have a good understanding of possible saturation effects or whether or how they impact the computations of cascading failure risk. Saturation effects have been observed as peaks in some simulations of

smaller power systems of the order of 100 buses [16], but it is unknown whether saturation effects are important in large power systems. It seems likely that any model of cascading could be expected to break down or become inaccurate at some extreme level of network disconnection, but simulations on smaller test systems do not provide data to address this question and the rarity of the largest possible blackouts of entire interconnections limits the applicability of observations.

Considerations of method efficiency and applicability typically depend on the range of λ considered. For example, lower values of λ make saturation unimportant, the standard deviation of the estimator $\hat{\lambda}$ depends on λ , and the continuous state branching process methods presented in this paper assume that $\lambda < 1$. Further work with realistic simulations and especially with observed data can determine the range of λ that we are trying to measure.

While there are reasonable arguments for choosing Poisson offspring distributions for the Galton-Watson branching processes, a good choice of the offspring distribution for continuous state branching processes is an open question. Directly estimating the offspring distribution from data could be difficult.

The role of time in the branching processes is one area of possible elaboration of the processing and the modeling, both in the possibility of modeling the time at which the failures occur and improving the methods of grouping failures into stages.

5.3 Further testing

This paper initially tests branching process models on cascading line failures in systems of the order of 100 buses with line outage data observed over one decade and data simulated by the OPA model of cascading line failures. The results are promising and further testing on other cases and larger systems is indicated. It would be particularly useful to test the results on cascading failure simulations with more modeling detail such as the Manchester model [28, 27] and TRELSS [31, 19] and any real data summarizing cascading that can be made available for research. If there is some universality to the gross features of cascading failure in power system blackouts, then many cases need to be tested to gain confidence in this universality.

The speculation that the propagation λ of line failures is the same as the propagation λ of load shed that is supported by the cases considered in section 3.4 should be tested further.

6 Conclusions

We have shown how to test branching process models on observed and simulated power system cascading failure data. The observed data is grouped into stages according to the failure times and the simulated data is naturally produced in stages by the OPA model of cascading line overloads. Then the propagation λ and other parameters of the branching process model are estimated and used to compute the probability distribution of blackout size. The estimator for λ has less bias than the estimator proposed in [14] when there are saturation effects. Both the number of line failures and the load shed can be used as measures of blackout size. The models corresponding to these measures are a Galton-Watson branching process with saturation and a continuous state branching process respectively. In the case of the observed outages, the method amounts to predicting from industry data the statistics of the size of cascading line outages when the initial line outage distribution is specified.

For the cases examined, which include cascading line overloads on the IEEE 118 bus test system and observations of line failures on a power system with approximately 100 buses, the probability of blackout size obtained via the branching process compares well with the probability of blackout size obtained by observing the power system for a long time or by exhaustive simulation. Quantifying cascading failure by first estimating branching process parameters requires significantly fewer cascades to be observed or simulated. This efficiency is key to practical application to monitoring the risk of cascading failure or using simulations to efficiently quantify the reliability benefits of proposed upgrades.

The initial results in this paper are promising and further testing of branching process models is warranted using other cascading failure simulations and other observations of cascading failures. The bulk statistical analysis of cascading failure using branching process models is developed here for blackouts, but could also be tested on data for cascading failure in or between other infrastructures.

7 Acknowledgements

We gratefully acknowledge that this paper is an account of work sponsored in part by the Power Systems Engineering Research Center (PSERC) and support in part from NSF grants ECCS-0606003 and SES-0623985. Janghoon Kim gratefully acknowledges support in part from scholarships from the Korean electric power corporation. Hui Ren gratefully acknowledges support in part from scholarships from the People's Republic of China and North China Electric Power University.

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 Power Delivery*, vol. 9, no. 1, Jan. 1994, pp. 21-39.
- [2] M. Anghel, K. A. Werley, A. E. Motter, Stochastic Model for Power Grid Dynamics, 40th Hawaii International Conference on System Sciences, Hawaii, January 2007.
- [3] K.B. Athreya, P.E. Ney, *Branching Processes*, Dover NY 2004 (reprint of Springer-verlag Berlin 1972).
- [4] 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.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] Q. Chen, C. Jiang, W. Qiu, J.D. McCalley, Probability models for estimating the probabilities of cascading outages in high-voltage transmission network, *IEEE Trans. Power Systems*, vol. 21, no. 3, August 2006, pp. 1423-1431.
- [9] P.C. Consul, M.M. Shoukri, Some chance mechanisms related to a generating Poisson probability model, *American Journal of Mathematical and Management Sciences*, vol. 8, nos. 1 and 2, pp. 181-202, 1988.
- [10] P.C. Consul, *Generalized Poisson Distributions*, Dekker, NY 1989.
- [11] J.-P. Dion, N. Keiding, Statistical inference in branching processes, in *Branching Processes*, editors A. Joffe, P. Ney, Marcel Dekker, New York 1978, pp. 105-140.
- [12] 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, 2004.
- [13] 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.
- [14] 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.
- [15] 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, 2005.
- [16] 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.
- [17] I. Dobson, B.A. Carreras, V.E. Lynch, D.E. Newman, Complex systems analysis of series of blackouts: cascading failure, critical points, and self-organization, *Chaos*, vol. 17, no. 2, June 2007.
- [18] P. Guttorp, *Statistical inference for branching processes*, Wiley, NY, 1991
- [19] 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.
- [20] T.E. Harris, *Theory of branching processes*, Dover NY 1989.
- [21] P.J.M. Kallenberg, *Branching processes with continuous state space*, Mathematical Centre Tracts 117, ISBN 90 6196 188 2, Amsterdam 1979.
- [22] D.S. Kirschen, D. Jayaweera, 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.
- [23] 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.
- [24] S. Mei, Yadana, X. Weng, A. Xue, Blackout model based on OPF and its self-organized criticality, Proceedings of the 25th Chinese Control Conference, Harbin, Heilongjiang, China, August 2006.
- [25] L. Mili, Q. Qui, A.G. Phadke, Risk assessment of catastrophic failures in electric power systems, *International Journal of Critical Infrastructures*, vol. 1, no. 1, pp.3863, 2004.
- [26] 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.
- [27] 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.
- [28] M.A. Rios, D.S. Kirschen, D. Jayaweera, D.P. Nedic, R.N. Allan, Value of security: modeling time-dependent phenomena and weather conditions. *IEEE Transactions on Power Systems*, vol. 17, no. 3, pp. 543-8, 2002.
- [29] E. Seneta, D. Vere-Jones, On the asymptotic behaviour of subcritical branching processes with continuous state space, *Zeitschrift fur Wahrscheinlichkeitstheorie und verwandte Gebiete*, vol. 10, pp. 212-225, 1968.
- [30] 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.
- [31] Transmission reliability evaluation for large-scale systems (TRELSS): version 6.0 User's manual, EPRI, Palo Alto, CA: 2000. 1001035
- [32] T. Van Cutsem, Voltage instability: phenomena, countermeasures, and analysis methods, *Proceedings of the IEEE*, vol. 88, no. 2, Feb 2000, pp. 208-227.
- [33] E.T. Whittaker, G.N. Watson, *A Course of Modern Analysis*, Cambridge University Press, Cambridge, UK, 1948.
- [34] D.V. Widder, *The Laplace Transform*, Princeton University Press, Princeton, NJ, 1946.
- [35] 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.
- [36] N.M. Yanev, On the statistics of branching processes, *Theory of Probability and its Applications*, vol. 20, 1975, pp. 612-622.