

2016

North American Blackout Time Series Statistics and Implications for Blackout Risk

Benjamin A. Carreras
BACV Solutions Inc.

David E. Newman
University of Alaska, Fairbanks

Ian Dobson
Iowa State University, dobson@iastate.edu

Follow this and additional works at: http://lib.dr.iastate.edu/ece_pubs



Part of the [Power and Energy Commons](#)

The complete bibliographic information for this item can be found at http://lib.dr.iastate.edu/ece_pubs/141. For information on how to cite this item, please visit <http://lib.dr.iastate.edu/howtocite.html>.

This Article is brought to you for free and open access by the Electrical and Computer Engineering at Iowa State University Digital Repository. It has been accepted for inclusion in Electrical and Computer Engineering Publications by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

North American Blackout Time Series Statistics and Implications for Blackout Risk

Abstract

We use North American Electric Reliability Corporation historical data to give improved estimates of distributions of blackout size, time correlations, and waiting times for the Eastern and Western interconnections of the North American grid. We then explain and estimate the implications of the power law region (heavy tails) in the empirical distribution of blackout size in the historical data for the Western interconnection. Annual mean blackout size has high variability and the risk of large blackouts exceeds the risk of medium size blackouts. Ways to communicate blackout risk are discussed.

Keywords

Power transmission reliability, risk analysis

Disciplines

Electrical and Computer Engineering | Power and Energy

Comments

This is a manuscript of an article published as Carreras, Benjamin A., David E. Newman, and Ian Dobson. "North American blackout time series statistics and implications for blackout risk." IEEE Transactions on Power Systems 31, no. 6 (2016): 4406-4414. doi: [10.1109/TPWRS.2015.2510627](https://doi.org/10.1109/TPWRS.2015.2510627). Posted with permission.

Rights

© 2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

North American blackout time series statistics and implications for blackout risk

Benjamin A. Carreras, David E. Newman, Ian Dobson, *Fellow IEEE*

Abstract—We use North American Electric Reliability Corporation historical data to give improved estimates of distributions of blackout size, time correlations, and waiting times for the Eastern and Western interconnections of the North American grid. We then explain and estimate the implications of the power law region (heavy tails) in the empirical distribution of blackout size in the historical data for the Western interconnection. Annual mean blackout size has high variability and the risk of large blackouts exceeds the risk of medium size blackouts. Ways to communicate blackout risk are discussed.

Index Terms—Power transmission reliability, risk analysis.

I. INTRODUCTION

The Disturbance Analysis Working Group of the North American Electric Reliability Corporation (NERC) published online records of blackout size and duration that give a historical time series of reportable blackouts in North America. This NERC data is foundational for quantifying and understanding blackout risk. We use 22 years of this data from 1984 to 2006. The data includes the power shed, number of customers disconnected, and duration of reportable blackouts. In this paper we analyze only the load shed and customers disconnected because of uncertain interpretation and missing data for the blackout durations. The data arise from government blackout reporting requirements. The thresholds for a reportable blackout include uncontrolled loss of 300 MW or more of firm system load for more than 15 minutes from a single incident, load shedding of 100 MW or more implemented under emergency operational policy, loss of electric service to more than 50 000 customers for one hour or more, and other criteria detailed in the United States Department of Energy form EIA-417. The data has some unevenness, so that very precise conclusions are elusive, but general conclusions can be drawn that describe the nature of electric power transmission system reliability. Note that the NERC data and this paper do not address the more common and smaller distribution system blackouts.

B. A. Carreras is with BACV Solutions Inc., Oak Ridge, TN USA; email bacarreras@gmail.com, D. E. Newman is with Physics Dept., University of Alaska, Fairbanks AK 99775 USA; email ffdn@uaf.edu, and I. Dobson is with the ECpE Dept., Iowa State University, Ames IA 50010 USA; email dobson@iastate.edu. We gratefully acknowledge funding in part by the California Energy Commission, Public Interest Energy Research Program. This paper does not necessarily represent the views of the Energy Commission, its employees or the State of California. It has not been approved or disapproved by the Energy Commission nor has the Energy Commission passed upon the accuracy or adequacy of the information. BAC also gratefully acknowledges the support from a “Cátedra de Excelencia” from Universidad Carlos III-Banco de Santander Project. ID also gratefully acknowledges the support of NSF grant CPS-1135825.

Starting in January 2000 [1], analyses or presentations of portions of the NERC data set have been published by several research groups [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15]. All these research groups report a heavy tail in the distribution of blackout size. Similar results appear in historical blackout records from other countries [16]. The heavy tail of the distribution of blackout size shows the significant contribution of large blackouts to blackout risk as explained, for example, in [5], [13]. Possible trends or changes in blackouts over time are variously analyzed in [6], [8], [9], [10]. Seasonal variations are analyzed in [7], [8], [9], [10]. A larger data set is analyzed for North America in [10]. Distributions of the times between blackouts are suggested in [5], [7] and evaluated in [10]. Imperfections in the NERC data are discussed more extensively in [12]. Data mining is applied in [11].

This paper addresses the following general questions using the NERC data:

- What is the statistical variation in blackout size?
- How can we usefully quantify blackout size with statistics and communicate these statistics?
- Are mean blackout sizes useful statistics?
- Do large blackouts have greater risk than smaller blackouts?
- Do the initiating events that lead to blackouts occur at random times?
- To what extent are blackouts correlated with past blackouts?

In particular, we make the following contributions to further analyze and understand the NERC data and its implications:

- Sections II and III take advantage of the larger data set and correct the size of several large blackouts in order to obtain better estimates of the power law exponents and long-time correlations for blackout size and to separately analyze the Eastern and Western interconnections. In particular this improves on initial work in [5] for the power law exponents and the long-time correlations and with respect to [4] for the power law exponents and the separate analyses of the Eastern and Western interconnections.
- Section IV obtains a better description of the statistics of the times between blackouts as a novel nonhomogeneous Poisson process, corrects some misconceptions in the literature, and concludes that the new model is consistent with randomly occurring blackout initiating events.
- Section V presents a new way to quantify the variability of mean blackout size directly from the data, finds the mean blackout size to be highly variable, and discusses

the consequences for blackout statistics, including advice to avoid using mean blackout size statistics.

- Section VI estimates in a novel, data-driven way the risk of various sizes of blackouts for the Western interconnection. While overall transmission grid reliability is good, it is important to maintain this reliability, and to quantify which sizes of blackouts pose the greater risk. In this case, the risk of large blackouts exceeds the risk of medium-size blackouts.
- Section VII briefly discusses how to communicate blackout risk to technical and non-technical audiences.

Since the results are sensitive to the largest blackouts, we corrected the NERC load shed data for the Western interconnection for the largest blackouts. We examined NERC annual reports [17], [18] and an expert account of the 1996 blackout [19] and used the larger load shed in these sources for four of the largest blackouts as detailed in Table V at the end of the paper. (These sources are more credible than the NERC data, partly because the NERC data reports have to be filed quickly after the blackout by multiple parties.) No changes to the Eastern interconnection NERC data were made.

One general point about observed blackout data is that because it includes large blackouts that are rare but consequential, there is inherently a shortage of data that limits the precision of the conclusions. This highlights the importance of this paper analyzing a longer data set and correcting the larger blackout data. Also, some of the methods we apply (such as Clauset's method and bootstrap resampling) can test for or circumvent the paucity of data.

II. DISTRIBUTION OF BLACKOUT SIZE

The observed distribution of blackout size, such as in terms of load shed or customers disconnected, is one of the most foundational descriptions of transmission grid reliability. There are several reasons for the centrality of the distribution of blackout size for power transmission reliability. First, large blackouts cannot be neglected; although rare, large blackouts have a large impact on society as well as the most consequential regulatory and reputational impacts on the power industry. We show in section VI that the risk of large blackouts (here defined as more than 1000 MW shed) can exceed the risk of medium-size blackouts. This arises from the heavy-tailed behavior in the distribution of blackout size that we quantify with power-law exponents in this section; blackouts do become rarer as their size increases, but at a rate slower than the rate at which the cost of blackouts increases with size. Second, although it is easier to mitigate small blackouts since they are much simpler, and some mitigation measures for small blackouts also mitigate large blackouts, other mitigation measures that mitigate small blackouts, can, sometimes after a delay, increase the probability of large blackouts [13]. Therefore the blackout mitigation problem is best framed as jointly mitigating blackouts of all sizes, and this objective is naturally expressed as maintaining or shaping the distribution of blackout size. Third, as shown in section V, there can be pitfalls in attempting to summarize the distribution of blackout size with statistics such as the mean blackout size.

This section advances the state of the art by quantifying the power-law exponents of the distribution of blackout size with more data and better methods. In particular, the availability of more data and the correction of the data for the largest blackouts in the Western interconnection enables better calculation of the power law exponents of the distribution of blackout size separately for the Eastern and Western interconnections.

The distribution of blackout size can be shown as a complementary cumulative distribution function, abbreviated as CCDF. The CCDF is the probability that a blackout is larger than a given size as a function of blackout size. In particular, the CCDF evaluated at blackout size x is the integral of the probability density function from x to infinity. The CCDF is estimated from data by ordering the blackouts in order of decreasing size, plotting the rank of the blackout in this ordering against the blackout size, and renormalizing the vertical scale so that it indicates probability. While it is easier to compare the relative probabilities of different blackout sizes with the probability density function, estimating the probability density function requires decisions about binning of data that introduces some arbitrariness. Therefore in this paper we use the CCDF and the rank function, the unnormalized version of the CCDF. Fig. 1 shows rank functions for blackout size measured by load shed.

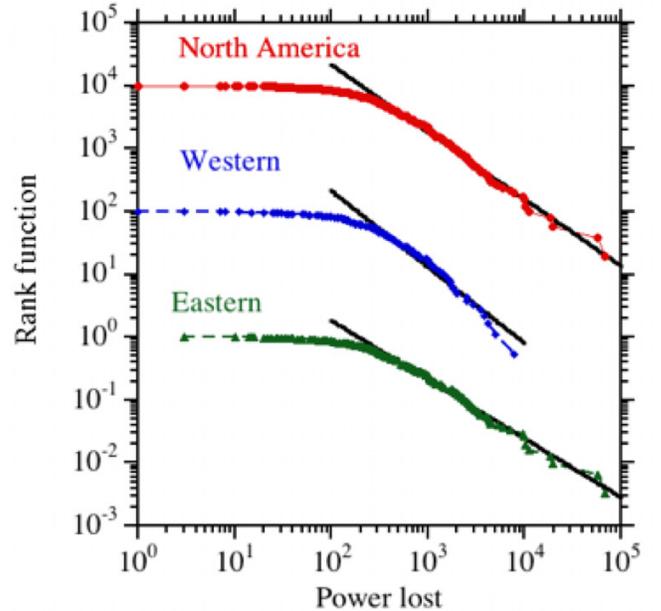


Fig. 1. Rank of blackout load shed. North American and Western data are multiplied by 10^4 and 10^2 to separate the plots.

Probability distributions of blackout size show power law regions. A power law region is a substantial interval of blackout size over which the blackout probability changes as a power function of the blackout size. A log-log plot of the probability distribution over a power law region shows as a straight line. Since CCDFs are integrals of probability density functions, the exponent of the power law and the slope of the straight line depends on whether the probability distribution

is specified as a probability density function or a CCDF. That is, over some range of x , if CCDF $F(x) \sim x^{-\alpha}$, then the corresponding probability density function $f(x) \sim x^{-\alpha-1}$. This paper quotes power law exponents as the exponent α for the CCDF.

Figs. 1 and 2 show the CCDF with power law fits for the distribution tails. The combined results for North America also include data for the Texas grid ERCOT.

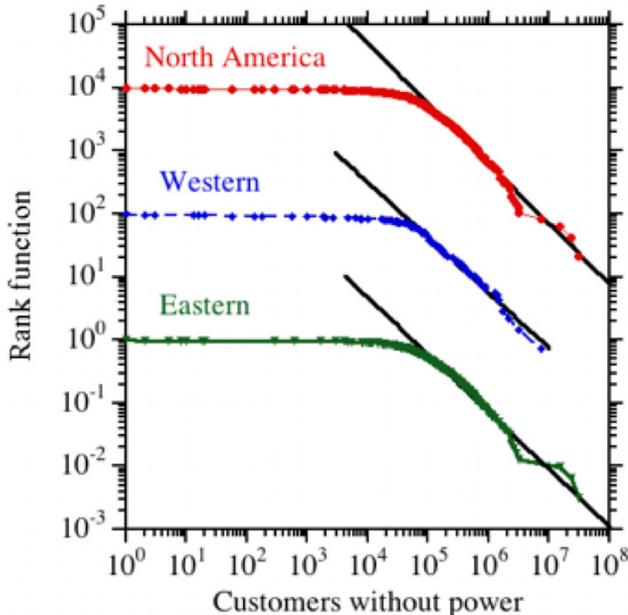


Fig. 2. Rank of blackout number of customers without power. North American and Western data are multiplied by 10^4 and 10^2 to separate the plots.

We use two methods to compute the power law exponents of the rank functions of blackout size, and the results are shown in Tables I and II. The power law tail region to be determined from the data is all the data with blackout size x greater than x_{\min} . The first method judges x_{\min} by inspection and fits a line to the data in this region. This method is simple but has subjective aspects. The second, more accurate and objective method is due to Clauset et al. [15]. The Kolmogorov-Smirnov distance D between the CCDF of the data for x greater than x_{\min} and the pdf we are trying to fit $(x/x_{\min})^{-\alpha}$ varies with the choice of x_{\min} . In Clauset's method, the distance D is calculated for a range of choices of x_{\min} and then x_{\min} is chosen to minimize D . Then a maximum likelihood estimator is used to estimate α .

Table I shows the exponents of the power tails. For the Eastern interconnect the rank function of the load shed shows a power tail over two decades with exponent 0.94, whereas the Western interconnect shows a power tail with exponent 1.02. The North American exponent for load shed of 1.07 agrees with [10]. For all cases the rank function of the customers disconnected has a power law tail with exponent approximately 0.9.

Clauset's method has a goodness of fit parameter p , where approximately $p > 0.1$ indicates a good fit. All the exponents

TABLE I
POWER LAW EXPONENTS FOR LOAD SHED

	North America	Western	Eastern
Fit tail by inspection: power law exponent α	1.07	1.02	0.94
Method of Clauset: power law exponent α	1.16 ± 0.10	0.98 ± 0.10	0.82 ± 0.06
x_{\min}	850.	260.	245.
D	0.076	0.07	0.14
p	0.45	0.68	0.04
tail data points n	123	102	204
total data points	512	185	307

TABLE II
POWER LAW EXPONENTS FOR CUSTOMERS DISCONNECTED

	North America	Western	Eastern
Fit tail by inspection: power law exponent α	0.95	0.88	0.91
Method of Clauset: power law exponent α	0.82 ± 0.05	0.92 ± 0.11	0.84 ± 0.07
x_{\min}	90 000	82 000	110 000
D	0.057	0.44	0.078
p	0.59	0.99	0.37
tail data points n	252	70	153
total data points	467	138	314

have a good fit, with the exception of the exponent for the power shed for the Eastern interconnection. This poorer fit can be attributed to insufficient large blackout data to guarantee the statistical validity of the fit. Uncertainties on all these exponents are still significant. We note that Hines et al. [9] apply Clauset's method to 22 years of North American blackout size data normalized to year 2000 according to population size ratio and obtain for power shed $x_{\min} = 1012$ MW and $\alpha = 1.2$, and for number of customers disconnected $x_{\min} = 291\,000$ and $\alpha = 1.14$. And Clauset analyzes 18 years of North American blackout data on number of customers disconnected and obtains $x_{\min} = 230\,000$ and $\alpha = 1.3$ [15].

The importance of the power law behavior in the distribution of blackout size can be seen by contrasting the power law behavior with the behavior of distributions of blackout size that decrease exponentially or faster than exponential as blackout size increases [13]. An exponential decrease in probability implies that the probabilities of the largest blackouts are so vanishingly small that they would not occur in practice¹. Blackout risk is the product of blackout probability and blackout cost. Generally, as blackout size increases, blackout probability decreases and blackout cost increases, and the behavior of blackout risk depends on which factor dominates. Blackout costs are uncertain, but the direct blackout costs increase at least linearly with blackout size [13]. In the case of blackout size probability decreasing exponentially with increasing blackout size, the blackouts get rarer much faster than blackout costs increase, so that the risk of large blackouts is very small. In contrast, in the power law case,

¹The popular press describes this condition as a "perfect storm," even when it is known that the phenomenon is heavy tailed and that extreme events are to be expected occasionally.

with increasing blackout size, the blackouts get rarer slowly enough that the increased blackout costs can make the risk of large blackouts exceed the risk of small blackouts [13]. To quantify these effects, section VI estimates this risk based on data.

There is some variation in usage in terms for power law regions. In both probability and physics, a power law tail in a probability density function $f(x)$ with an unbounded x variable indicates a power law or an asymptotic power law as x tends to infinity. However, if the distribution is bounded in x , physicists may also use the term power law tail, assuming that it is obvious that the power law region is limited, whereas a probabilist would be likely to avoid this term. All probability distributions for blackout size are bounded by the largest possible blackout in which entire interconnection blackouts out.²

The power law region of the distribution of blackout size has been explained using complex systems concepts of criticality and self-organization [5], [20], [16], [21], [13], [22], [14].

III. LONG-RANGE TIME CORRELATIONS

As complex systems evolve they can show correlations over various time scales. For example, a large blackout always leads to grid upgrades, and these upgrades can affect the blackouts for some considerable time in the future. These correlations can be detected as autocorrelation functions decaying slowly in a power law fashion, and the power law can be measured with Hurst exponents³ [23], [24]. A Hurst exponent greater than 0.5 signifies a positive correlation between blackouts at one time and blackouts at all other times, a Hurst exponent less than 0.5 indicates negative correlations, and a Hurst exponent equal to 0.5 indicates no correlations. A positive or negative correlation is sometimes referred to as system memory. It means that the memory of a blackout is retained in the operational procedures or engineering upgrades which will then impact failures at later times. Our paper [5] previously found moderate long range time correlations (Hurst exponent 0.6) in the NERC blackout data.

To improve on [5], we recompute for the larger data set the Hurst exponents that describe the long-time correlations

²Historical data for North American blackouts does not show the nature of the blackout distribution near the largest possible blackout because blackouts of the entire or almost the entire interconnections have not occurred. It is not clear whether these blackouts are so rare that they have not happened over the time period that the interconnections have existed, or whether they are extremely unlikely to happen. Theory models for cascading with subcritical propagation show a power law region for intermediate blackouts and an exponentially decaying region for the largest blackouts [16] However, larger probabilities of the largest blackouts are also possible in these theory models if the propagation is larger.

³The long-range dependence in a time series can be quantified with the rescaled range R/S statistic. One first integrates the time series, and then computes the R/S statistic which is the range of blocks of m successive points divided by the block standard deviation. The idea is to quantify how this rescaled range R/S statistic grows as the time lag m increases. For a time series with an autocorrelation function that has a power law tail, the R/S statistic scales proportionally to m^H , where H is the Hurst exponent. That is, H is the asymptotic slope on a log-log plot of the R/S statistic versus the time lag m . Formulas for the R/S statistic are in the appendix of [5].

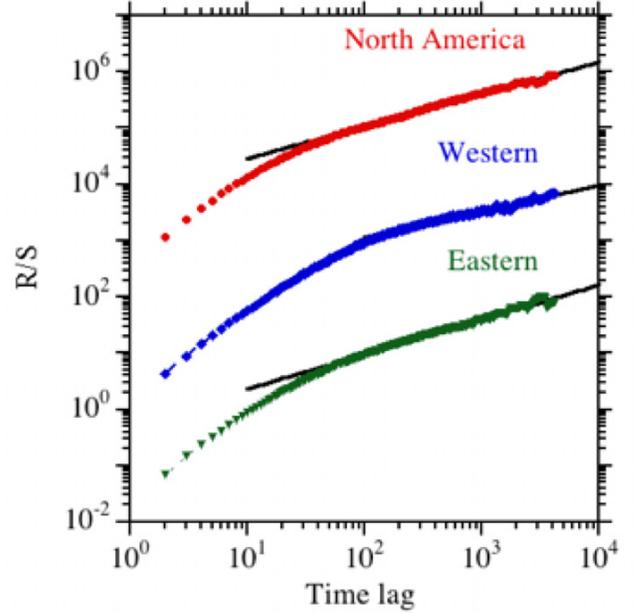


Fig. 3. R/S plots for power shed. North American and Western data are multiplied by 10^4 and 10^2 to separate the plots.

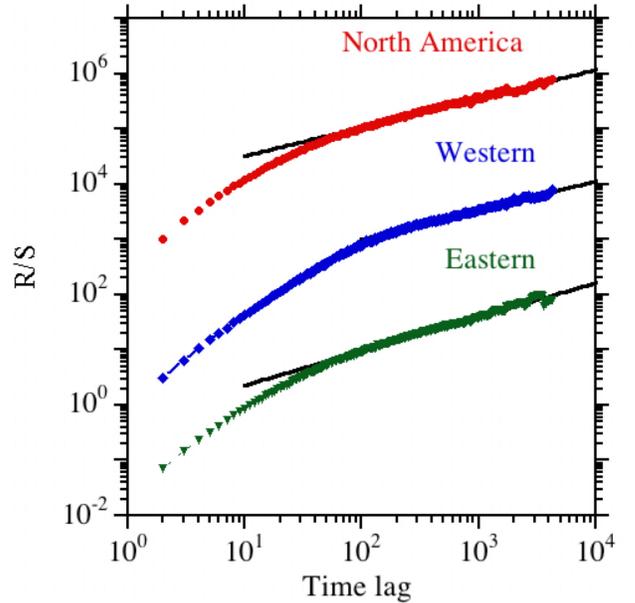


Fig. 4. R/S plots for customers. North American and Western data are multiplied by 10^4 and 10^2 to separate the plots.

in the blackout time series.⁴ The Range over Scale plots shown in Figures 3 and 4 show an algebraic scaling region

⁴Though more points would always be useful, there are enough points to allow for approximately a decade and a half of the asymptotic power law region for the calculation of the Hurst exponent. It is generally agreed that anything over a decade can produce a valid exponent. The R/S method for calculating the Hurst exponent has small uncertainty for the low lag points since they are averaged over many samples, while the uncertainty grows for the higher lag. The last point has only one sample (the entire length) and is therefore not included in the exponent calculation.

TABLE III
HURST EXPONENTS AND WAIT TIME PARAMETERS

	North America	Western	Eastern
Hurst exponent; power shed	0.57	0.46	0.60
Hurst exponent; customers	0.52	0.53	0.62
wait time parameter β_0	0.029	0.015	0.016
wait time parameter β_1	0.18	0.034	0.13

for time lags above 100 days. Table III shows the Hurst exponents for the Eastern and Western interconnections. The Eastern interconnect shows the presence of mild positive long range time correlations for power shed and customers disconnected. The presence of these mild positive correlations is consistent with a long-term dynamical effect by which blackouts, including those far in the past, might influence the present. One mechanism for such a long-term dynamical effect is the complex system self-organization in OPA models of series of blackouts [20], [16], [21]. The long range time correlations for the Western interconnect are weaker and less conclusive.

IV. WAITING TIMES

Waiting times are the times between successive blackouts, and their statistics indicate the nature of the blackout initiating events. In particular, identifying non-random statistical patterns in blackout-initiating events could inform their mitigation. In this section we improve the statistical modeling of blackout-initiating events, resolve questions raised in the literature, and find no evidence of non-randomness. Modeling the statistics of blackout initiating events is also essential for cascading failure simulations.

These statistics were first suggested to be exponential with a constant rate [5], and this approximation was questioned in [7], [10]. In this paper, since the number of blackouts is slowly changing, we assume that the rate of the exponential β is slowly changing over the blackout time interval with a uniform distribution over the interval $[\beta_0, \beta_1]$. Then the probability distribution of wait time W becomes

$$f_W(t) = \int_{\beta_0}^{\beta_1} \beta e^{-\beta t} \frac{1}{\beta_1 - \beta_0} d\beta \quad (1)$$

And the rank function for W is

$$F_W(t) = \int_t^{\infty} f_W(\tau) d\tau = \frac{e^{-\beta_0 t} - e^{-\beta_1 t}}{(\beta_1 - \beta_0)t} \quad (2)$$

The values of β_0 and β_1 shown in Table III are obtained by fitting the rank function for W to the waiting times in the data as shown in Figure 5. We can see that the waiting time data is consistent with an inhomogeneous Poisson process.

The change $\beta_1 - \beta_0$ is smaller for the Western interconnection. This reflects random triggers with a possible small increase in frequency in the last few years. The Eastern interconnect shows a stronger time dependence of the frequency. The present data is consistent with random set of triggers, with frequency changing with time. The assertions in [7] about the non-random origin of the triggers arise from imposing a homogeneous Poisson process model that is incompatible with the data.

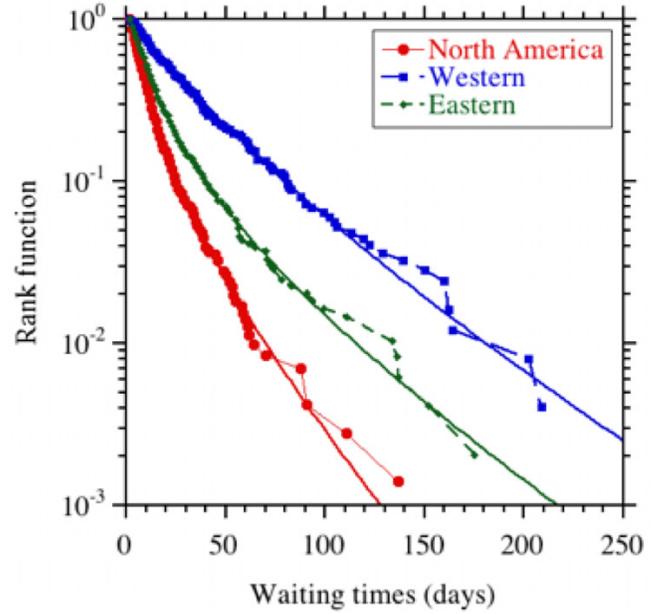


Fig. 5. Distributions of waiting times between blackouts.

Reference [7] claims that self-organized-criticality-type dynamics should exhibit exponential decay in the waiting time distribution. However this claim is not correct: the complex systems properties of blackouts such as self-organization or criticality are independent of the statistics of the waiting times; properties of the complex system are distinct from the properties of the system inputs [25].

V. VARIABILITY OF MEANS

The heavy tail in the probability distribution of blackout size implies high variability in blackout size and that estimates for mean or average blackout size can fluctuate excessively. This section analyzes the NERC data for load shed in the Western interconnection to quantify the excessive fluctuation. This leads to our recommendation that mean values for blackout size should be avoided.

The processing of the data is now described. In order to establish a uniform minimum blackout size of 100 MW, blackouts smaller than 100 MW were omitted from the data. This left 111 blackouts observed over 22 years. There are on average five blackouts larger than 100 MW per year.

The estimated mean is the sum of the samples divided by the number of samples. Under suitable conditions, the estimated mean becomes close to the actual mean as the number of samples becomes large enough. In the case of blackouts, there is always a blackout of maximum size (the entire interconnection blacks out). Therefore, in theory, estimates of the mean eventually converge to the mean as the number of samples becomes infinite. However, in practice, this convergence may require a very large number of samples, and there can be substantial variation in the estimated mean for practical sample sizes. For real blackout data, it is impractical to wait too long for enough blackout samples to accumulate. (For simulated

blackouts, there are similar difficulties with excessively long run times.)

We estimate the mean and standard deviation of load shed from the Western interconnection NERC data. The raw data of 111 load shed amounts was not directly used to estimate the mean and standard deviation because of the high variance in these quantities for the relatively small numbers of observed blackouts ... indeed this is exactly the problem under investigation! Instead, we use a nonparametric bootstrap [26] to generate 200 000 samples of blackout load shed.⁵ (The 200 000 samples correspond to 40 000 years of operation of the WECC.) The mean of the 200 000 samples is approximately 1500 MW and the standard deviation is approximately 4500 MW. This procedure may well underestimate the mean and standard deviation since the fit of the CCDF assigns a very small probability to blackouts in excess of the largest observed WECC blackout of 30 390 MW. Note that it is assumed that the distribution of load shed does not change over the 22-year period of observation.

Having estimated the mean and standard deviation of the distribution of load shed, it is now feasible to estimate the standard deviation of the estimated mean of a given number of independent load shed samples. Suppose that the blackouts are observed for y years. Since there are 5 reportable blackouts per year, there are $5y$ samples of load shed in y years. The standard deviation of the load shed is 4500 MW. It follows that the standard deviation of the estimated mean of $5y$ samples is $4500/\sqrt{5y}$ MW. For example, by setting $y = 1$ we obtain that the annual mean blackout load shed has standard deviation of about 2000 MW. It takes 50 years of observation, or 250 blackout samples, to reduce the standard deviation of the estimated mean to 280 MW. The point is that it takes many years to reduce the standard deviation of the estimated mean load shed enough to make the estimated mean load shed useful.

To further illustrate this point, we assume one or ten years of observation (that is, 5 or 50 samples) and use the generated load shed samples to generate an empirical distribution of mean load shed.⁶ The estimated annual mean load shed is distributed so that:

- with probability 0.66, $0 < \text{annual mean} < 1000$.
- with probability 0.15, $1000 < \text{annual mean} < 2000$.
- with probability 0.09, $2000 < \text{annual mean} < 3000$.

Recalling that the actual mean is approximately 1500 MW, the annual mean is estimated to be within 500 MW of the actual mean only in 15 percent of cases.

The estimated 10-year mean load shed is distributed so that:

- with probability 0.28, $0 < 10\text{-year mean} < 1000$.
- with probability 0.53, $1000 < 10\text{-year mean} < 2000$.
- with probability 0.16, $2000 < 10\text{-year mean} < 3000$.

⁵The details of the bootstrap are that we form the empirical CCDF by ranking the blackouts in decreasing size and dividing the rank by the number of blackouts, fit a function F to the empirical CCDF using linear interpolation, and then generate the samples of blackout load shed by applying F^{-1} to samples from a uniform distribution between 0 and 1.

⁶We avoid using the central limit approximation of normality for this calculation since this may fail for heavy tailed distributions.

The 10-year mean is estimated to be within 500 MW of the actual mean in only slightly more than half the cases.

To summarize, to reliably and accurately estimate the mean load shed in Western interconnection blackouts from observations would require several decades of observation. This is impractically long. Moreover, even if the annual mean were estimated reliably and accurately, it would not be very useful since there is such large variability of blackouts about the mean that the mean value is not representative.

According to Table I, the power law exponent α for the Eastern interconnection blackout load shed is less than the power exponent for Western interconnection blackout load shed so that the distribution of load shed for the Eastern interconnection has a slightly heavier tail than the Western interconnection. Since a heavier tailed distribution has higher variability of means, the result in this section that Western interconnection mean load shed is highly variable and unrepresentative also applies to the Eastern interconnection.

CCDF or equivalent distributions are good ways to present extreme events statistics. These distributions are not a single index and the raw data is preserved. There is sometimes a push to summarize highly variable data with one number, even when this does not make sense and modern information technology does not require the reduction to one number.

We comment on the exclusion of large events from highly variable data. It is a common practice to exclude large events when gathering data or computing reliability indices. This practice has important implications when assessing extreme event risk. The criteria for which large events are excluded vary significantly between organizations, and include whether the outage was planned, blackout size or duration, and the cause of the blackout such as storms, or the history of recent blackouts. Moreover, the extreme events are inconsistent with the most common probability models that lack heavy tails, and are sometimes wrongly excluded because they do not fit the assumed but wrong models.

Some commonly used indices are annual mean values. One of the reasons for the exclusion is that if the extreme events are not excluded, the annual means are highly variable. The high variability of an index, unless explained as an expected consequence of the heavy tailed phenomenon being assessed, can produce problems of perception and interpretation. We recommend that all data be carefully collected without exceptions, and that processing and display of the data should describe any filtering or censoring of the data.

Cascading outages result from initial outages that then propagate across the grid. One way to avoid the problems of metrics directly characterizing heavy tailed distributions of blackout sizes is to instead characterize both the initial outages and the average propagation of the outages. For transmission line outages, branching process models can predict the blackout size distribution from the initial outages and the average propagation. Moreover, the average propagation does not have heavy tails and can be estimated from one year of outage data from a large utility [27]. Extensions to similarly quantify load shed statistics are being pursued [28], [29], [30].

The high variability in blackout size and the consequent high variability in estimated mean blackout size is an intrinsic

property of blackouts that arises from the nature of cascading events. We suggest that power transmission grid reliability can best be maintained by metrics and policies that account for this observation. In particular, despite their widespread use in the blackout literature, mean blackout size and other statistics that compute an average such as conditional value at risk are too highly variable to be useful or representative statistics. If there is any doubt about this conclusion for blackouts in a particular country or region, it is straightforward to apply the bootstrap sampling described in this section to the appropriate blackout size data to quantify the variability of the statistic.

VI. RISK OF LARGE BLACKOUTS

This section roughly estimates the risk of medium size and large Western interconnection blackouts based on NERC historical data. For this calculation, risk is defined as probability times cost, medium size blackouts have load shed between 100 MW and 1000 MW, and large blackouts have load shed greater than 1000 MW. We make the simple and pragmatic assumption that direct blackout costs are proportional to energy unserved. Energy unserved is load shed times blackout duration. Analysis of the NERC North American data suggests the rough and conservative approximation that large blackout duration is proportional to the square root of load shed, as stated in [31] and summarized in the Appendix. This implies that blackout cost is proportional to load shed to the power 1.5. Determination of blackout costs poses difficulties that impact the estimation of risk and further progress in determining blackout costs would repay systematic data collection and further investigation [12], [13].

Suppose that the blackout load shed is s MW and that blackout direct cost C is proportional to s^γ . Let $f(s)$ be the probability density function of blackout size and $F(s)$ be the cumulative density function of blackout size. These probability densities are conditioned on a blackout happening. Then, given that a blackout happens, the probability that a blackout has load shed between s_1 MW and s_2 MW is

$$P[s_1, s_2] = \int_{s_1}^{s_2} f(s) ds = F(s_2) - F(s_1). \quad (3)$$

If blackouts occur with frequency b per year, then the expected number of blackouts per year between s_1 and s_2 MW is $bP[s_1, s_2]$.

Let $R[s_1, s_2]$ be the risk of a blackout having load shed between s_1 MW and s_2 MW. Risk is probability times cost. Therefore, given that a blackout happens, the risk of a blackout with load shed between s_1 MW and s_2 MW is

$$R[s_1, s_2] = \int_{s_1}^{s_2} f(s) s^\gamma ds. \quad (4)$$

This can be rewritten in terms of the cumulative density function as

$$R[s_1, s_2] = F(s_2) s_2^\gamma - F(s_1) s_1^\gamma - \int_{s_1}^{s_2} F(s) \gamma s^{\gamma-1} ds. \quad (5)$$

The cumulative density function has better properties than the probability density function, so (5) may have some advantages over (4) when computed from data.

TABLE IV
APPROXIMATE RISK OF MEDIUM AND LARGE BLACKOUTS

cost exponent γ	largest blackout	Risk (arbitrary units)	
		medium blackout	large blackout
1.5	30 390 MW	60	700
1.0	30 390 MW	3	8
1.5	150 000 MW	60	3300
1.0	150 000 MW	3	16

Note: Medium blackouts range between 100 MW and 1000 MW.
Large blackouts exceed 1000 MW.

The first line of Table IV shows the result of roughly estimating the risk of medium and large blackouts from the NERC WECC data with (5). As explained above, the best currently available, but uncertain approximation $\gamma = 1.5$ is used for the cost exponent. Since the cost has arbitrary units, so does the risk. The largest blackout in the NERC WECC data is 30 390 MW, so the first line of Table IV reflects the actual historical data. It can be seen in the first line of Table IV that the risk of a large blackout is an order of magnitude greater than the risk of a small blackout. Note that the calculation is conservative in only accounting for direct costs; the indirect costs of large blackouts can sometimes be very large.

The first line of Table IV is the best rough estimate available of the relative risks of medium and large blackouts, but the risk estimates are very sensitive to the assumptions. Lines 2, 3, and 4 of Table IV examine some changes in these assumptions.

Line 2 of Table IV changes the cost assumption in a conservative way that reduces the risk. The cost assumption for line 2 is that blackout cost is proportional to the load power shed so that $\gamma = 1.0$. The effect of this change in the cost assumption is that the risk reduces by approximately an order of magnitude, and that the risk of large blackouts is approximately double the risk of the medium blackouts.

The risk estimates also depend strongly on the largest blackout that occurs in the data. The largest possible blackout of WECC blacks out all of WECC, and has size approximately 150 000 MW. A hypothetical (and hopefully never to be experienced) blackout of 150 000 MW was added to the NERC historical data set to illustrate the sensitivity of the risk estimates to also observing the largest possible blackout. The results are shown in Lines 3 and 4 of Table IV. Including a hypothetical largest possible blackout in the data increases the estimated risk of the largest blackouts by a factor of two to five.

The risk estimates of medium and large blackouts are rough approximations, with significant uncertainties that depend on assumptions about cost and which of the largest blackouts occur in the data. The estimates could be improved in the future, but it is already plausible to conclude that the historical data shows that the risk of large blackouts exceeds the risk of the medium size blackouts, and possibly by a large amount. It is clear that further work to better determine the cost of large blackouts would improve the risk estimation. Our analysis is not inconsistent with the analysis of [32] that finds momentary interruptions to be important, because [32] does not address the largest blackouts. In [32], the surveys used to assess blackout cost do not ask about extended blackouts, and the

larger blackout data is eliminated together with outlier data when trimmed means are computed.

According to Table I, the power law exponent α for the Eastern interconnection blackout load shed is less than the power exponent for Western interconnection blackout load shed so that the distribution of load shed for the Eastern interconnection has a slightly heavier tail than the Western interconnection. Therefore the result in this section that the risk of large blackouts exceeds the risk of medium size blackouts for the Western interconnection also applies to the Eastern interconnection.

VII. COMMUNICATING BLACKOUT RISK

Since transmission system blackouts greatly impact the public, business, policy makers, regulators, and the entire power industry, and the understanding and perception of blackout risk influences everyone's response to blackouts, it seems appropriate not only to highlight the importance of the communication of blackout risk, but also to make a brief start in this section towards improving its communication to both non-technical and technical audiences. The effective communication of blackout risk can minimize misunderstanding and undue exaggeration or undue minimization of the risk.

For an initial approach, we suggest that blackouts be divided into small, medium, and large blackouts. For example, in the Western interconnection, one could consider small blackouts as less than 100 MW load shed, medium blackouts as between 100 and 1000 MW load shed, and large blackouts as more than 1000 MW load shed. For the 22 years of blackouts recorded by NERC, the probability of a Western interconnection blackout being greater than 100 MW is 0.8. Consider the CCDF of load shed shown in Fig. 6. This distribution is conditioned on the load shed being greater than 100 MW. According to Fig. 6, given that a blackout of more than 100 MW occurs (that is, the blackout is not small), the probability of medium blackouts is 0.74 and the probability of large blackouts is 0.26.

The main guideline for communicating probabilities from Gigerenzer's book [33] is to use natural frequencies (counts of events) rather than probabilities or conditional probabilities. A natural frequency approach to explain the CCDF in Fig. 6 is as follows: Over a period of 200 years there would be 1000 blackouts in the data. Then, on average, 190 of these blackouts will be less than 100 MW and 810 of these blackouts will be greater than 100 MW. On average, there would be 441 blackouts greater than 500 MW, 243 blackouts greater than 1000 MW, 45 blackouts greater than 5000 MW, and 18 blackouts greater than 10000 MW.

Another simple way of indicating blackout frequency specifies the average return time between blackouts. For example, on average, there is a blackout bigger than 10000 MW every 11 years.⁷

⁷One caveat with average return times is that they should not be confused with more likely outcomes. Suppose that the times between blackouts are exponentially distributed. Then, since the exponential distribution is always decreasing, intervals between blackouts shorter than 11 years are more likely. For example, there is a probability greater than 0.5 that the interval between blackouts bigger than 10000 MW is less than 8 years.

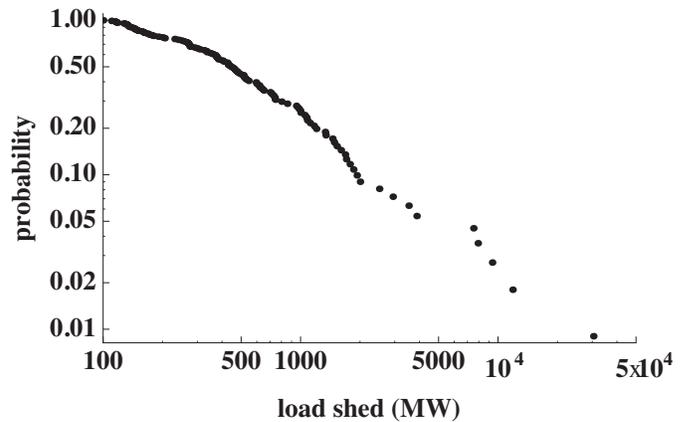


Fig. 6. CCDF of load shed in WECC blackouts.

Although we do not discuss blackout duration in the paper due to deficiencies in the quality of the duration data, it should not be forgotten that blackout duration also contributes strongly to blackout impact.

VIII. CONCLUSIONS

We use historical blackout data from NERC to give an improved analysis of reported blackouts of the Eastern and Western electric power transmission interconnections in North America, including estimates of the blackout size distributions in terms of load shed and customers disconnected, the statistics of waiting times, and long-term correlations between blackouts.

The observed blackout size distributions have heavy tails with a power law characteristic. Therefore, confirming previous studies, large blackouts are rare but expected to occur occasionally. Also accounting for the costs of large blackouts, the risk of large blackouts exceeds the risk of medium size blackouts in North America. This conclusion appears to be robust despite uncertainties in the blackout costs.

The heavy tails in the probability distribution of blackout sizes imply that blackout size is inherently highly variable. In particular, we make calculations based on observed data to show that, when large blackouts are not censored, annual mean values of blackout size are also inherently highly variable and are not representative or useful blackout statistics. We briefly discuss good ways to communicate the blackout statistics.

The statistics of the times between blackouts are well described by a Poisson distribution with a mean rate that slowly changes in time. There are mild long-term time correlations present between blackouts.

We reaffirm the importance of continuing to collect and analyze comprehensive blackout data for power transmission systems.

REFERENCES

- [1] B. A. Carreras, D. E. Newman, I. Dobson, A. B. Poole, Initial evidence for self-organized criticality in electric power system blackouts, 33rd Hawaii Intl. Conf. System Sciences, Maui, HI, January 2000.
- [2] M. Amin, National infrastructures as complex interactive networks, chapter 14 in Automation, Control, and Complexity: An Integrated Approach, T. Samad and J.R. Weyrauch, Eds., New York: Wiley, June 2000, pp. 263-286.

- [3] B. A. Carreras, D. E. Newman, I. Dobson, A. B. Poole, Evidence for self-organized criticality in electric power system blackouts, 34th Hawaii Intl. Conf. System Sciences, Maui, HI, Jan. 2001.
- [4] J. Chen, J.S. Thorp, M. Parashar, Analysis of electric power system disturbance data, 34th Hawaii Intl. Conf. System Sciences, Maui, HI, Jan. 2001.
- [5] B.A. Carreras, D.E. Newman, I. Dobson, A.B. Poole, Evidence for self-organized criticality in electric power system blackouts, IEEE Trans. Circuits & Systems, part I, vol. 51, no. 9, Sept. 2004, pp. 1733-1740.
- [6] M. Amin, Scanning the technology: Energy infrastructure defense systems, IEEE Proceedings, vol. 93, no. 5, May 2005, pp. 861-875.
- [7] R. Weron, I. Simonsen, Blackouts, risk, and fat-tailed distributions, Proc. 3rd Nikkei Econophysics Symposium, Springer Tokyo, 2005/6.
- [8] J. Simonoff, C. Restrepo, R. Zimmerman, Risk management and risk analysis-based decision tools for attacks on electric power, Risk Analysis, vol. 27, no. 3, 2007, pp 547-570.
- [9] P. Hines, J. Apt, S. Talukdar, Large blackouts in North America: Historical trends and policy implications, Energy Policy, vol. 37, no. 12, 2009, pp. 5249-5259.
- [10] D. Cornforth, Long tails from the distribution of 23 years of electrical disturbance data, Power Systems Conf. & Exp., Seattle WA, Mar. 2009.
- [11] D. Cornforth, Applications of data mining to time series of electrical disturbance data, IEEE PES General Meeting, 2009.
- [12] E. Fisher, J.H. Eto, K. H. LaCommare, Understanding bulk power reliability: the importance of good data and a critical review of existing sources, 44th Hawaii International Conference on System Sciences, Kauai, Hawaii, January 2011.
- [13] D.E. Newman, B.A. Carreras, V.E. Lynch, I. Dobson, Exploring complex systems aspects of blackout risk and mitigation, IEEE Transactions on Reliability, vol. 60, no. 1, March 2011, pp. 134-143.
- [14] B.A. Carreras, D.E. Newman, I. Dobson, N.S. Degala, Validating OPA with WECC data, Forty-sixth Hawaii International Conference on System Sciences, Maui, Hawaii, January 2013.
- [15] A. Clauset, C.R. Shalizi, M.E.J. Newman, Power-law distributions in empirical data, SIAM Review, vol. 51, no. 4, Nov. 2009, pp.661-703.
- [16] 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, 026103 June 2007.
- [17] NERC (North American Electric Reliability Council). 1994 system disturbances, October 1995.
- [18] NERC (North American Electric Reliability Council). 1996 system disturbances, August 2002.
- [19] D.N. Kosterev, C.W. Taylor, W.A. Mittelstadt, Model validation for the August 10, 1996 WSCC system outage, IEEE Trans. Power Systems, vol 14, no 3, August 1999 pp. 967-979.
- [20] 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.
- [21] H. Ren, I. Dobson, B.A. Carreras, Long-term effect of the n-1 criterion on cascading line outages in an evolving power transmission grid, IEEE Trans. Power Systems, vol. 23, no. 3, Aug. 2008, pp. 1217-1225.
- [22] S. Mei, X. Zhang, M. Cao, *Power Grid Complexity*, Tsinghua University Press, Beijing and Springer, Berlin 2011.
- [23] H. E. Hurst, Long-term storage capacity of reservoirs, Trans. Amer. Soc. Civil Eng., vol. 116, pp. 770-808, 1951.
- [24] B.B. Mandelbrot, J. R. Wallis, Noah, Joseph, and operational hydrology, Water Resources Research, vol. 4, pp. 909-918, 1969.
- [25] R. Sánchez, D.E. Newman, B.A. Carreras, Waiting-time statistics of self-organized-criticality systems, Physical Review Letters, vol. 88, no. 6, Feb 2002, Article ID 068302.
- [26] A.C. Davison, D. V. Hinkley, Bootstrap methods and their applications, Cambridge, England: Cambridge University Press, 1997.
- [27] I. Dobson, Estimating the propagation and extent of cascading line outages from utility data with a branching process, *IEEE Trans. Power Systems*, vol. 27, no. 4, November 2012, pp. 2146-2155.
- [28] J. Kim, K.R. Wierzbicki, I. Dobson, R.C. Hardiman, Estimating propagation and distribution of load shed in simulations of cascading blackouts, *IEEE Systems Journal*, vol. 6, no. 3, September 2012, pp. 548-557.
- [29] J. Qi, I. Dobson, S. Mei, Towards estimating the statistics of simulated cascades of outages with branching processes, *IEEE Trans. Power Systems*, vol. 28, no. 3, August 2013, pp. 3410-3419.
- [30] J. Qi, K Sun, Estimating the propagation of several cascading outages with multi-type branching processes, arXiv preprint arXiv:1405.6431, 2014.
- [31] M. Morgan, I. Dobson, et al., *Extreme Events Phase 2*, California Energy Commission Publication number CEC-MR-08-03, 2011.
- [32] K.H. LaCommare, J. Eto, Cost of power interruptions to electricity consumers in the United States, *Energy*, vol. 31, no.12, Sept 2006, pp. 1845-1855.
- [33] G. Gigerenzer, *Calculated risks: How to know when numbers deceive you*, Simon and Schuster, New York, 2002.

TABLE V
CORRECTIONS TO LOAD SHED DATA FOR 4 LARGE WECC BLACKOUTS

Blackout date	DAWG load shed	Corrected load shed	Source
Jan 17 1994	4235 MW	7500 MW	[17]
Dec 14 1994	5020 MW	9336 MW ^(a)	[17]
Jul 2 1996	2500 MW	11850 MW ^(b)	[18]
Aug 10 1996	0	30390 MW	[19]

(a) comprises 6877 MW firm and 2459 MW interruptible.

(b) sum of load shed in 5 islands.

APPENDIX: CONSERVATIVE APPROXIMATION OF BLACKOUT DURATION DEPENDENCE ON SIZE

The NERC data on blackout data on large blackouts is sparse and has uncertainties due to non-uniformities in blackout duration definitions and reporting. Therefore, we extract from the NERC data a conservative lower bound on the relationship between blackout duration and size rather than the relationship itself. We also improved the NERC duration data by systematically searching for and integrating public data about the NERC blackouts.

There appear to be two patterns of dependence of blackout duration D on load shed s in the improved NERC data. For smaller blackouts, there is a roughly linear dependence and for larger blackouts there is a sublinear dependence below the linear dependence. Fitting a power law curve to the sublinear larger blackout data yields $D = 29.5s^{0.69}$. Fig. 7 shows the improved NERC data for larger blackouts, the upper curve $D = 29.5s^{0.69}$, and a lower curve $D = 29.5s^{0.5}$ with an exponent 0.5 that is smaller than 0.69 and therefore more conservative. Most of the data lies above the lower curve $D = 29.5s^{0.5}$, even while allowing some margin for uncertainty in the data. Therefore we use D proportional to $s^{0.5}$ as a conservative lower bound.

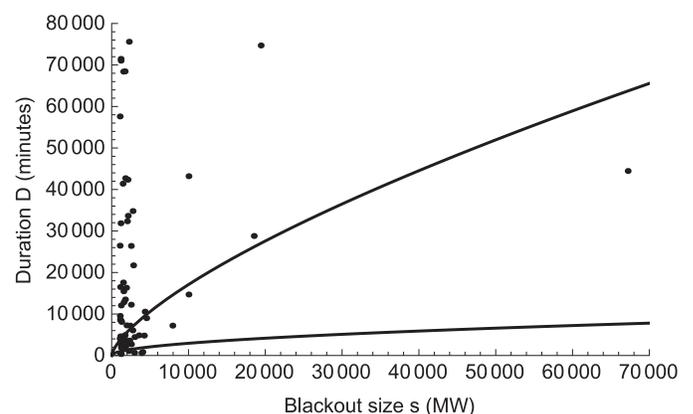


Fig. 7. Dots show improved NERC data for blackout duration versus size for larger blackouts. Upper curve $D = 29.5s^{0.69}$ fits a power law to the sublinear, large blackout data. Since lower curve $D = 29.5s^{0.5}$ lies well below most of data, it is a conservative lower bound.

Benjamin A. Carreras received the Licenciado en Ciencias degree in physics from the University of Barcelona, Spain and the Ph.D. degree in physics from Valencia University, Spain. He has been a Researcher and a Professor at the University of Madrid, Spain, Glasgow University, U.K., Daresbury Nuclear Physics Laboratory, Warrington, U.K., Junta de Energia Nuclear, Madrid, Spain, and the Institute for Advanced Study, Princeton, NJ. He was a Corporate Fellow at Oak Ridge National Laboratory, TN. He is now Principal Scientist at BACV Solutions Inc., Oak Ridge, TN. Dr. Carreras is a Fellow of the American Physical Society.

David E. Newman received the B.S. degree in physics and mathematics from the University of Pittsburgh, Pittsburgh, PA and the Ph.D. degree in physics from the University of Wisconsin, Madison. He was a Wigner Fellow and Research Scientist at Oak Ridge National Laboratory, TN. In 1998, he joined the University of Alaska, Fairbanks, where he is now a Professor of Physics. Dr. Newman is a Fellow of the American Physical Society.

Ian Dobson (F 06) received the BA in Mathematics from Cambridge University and the PhD in Electrical Engineering from Cornell University. He previously worked for British industry and the University of Wisconsin-Madison and is currently Sandbulte professor of engineering at Iowa State University.