Trends in the History of Large Blackouts in the United States

Paul Hines, *Member, IEEE*, Jay Apt, *Senior Member, IEEE* and Sarosh Talukdar, *Fellow, IEEE*

Abstract—Despite efforts to mitigate blackout risk, the data available from the North American Electric Reliability Council (NERC) for 1984-2006 indicate that the frequency of large blackouts in the United States is not decreasing. This paper describes the data and methods used to come to this conclusion, and several other patterns that appear in the data. These patterns have important implications for those who make investment and policy decisions in the electricity industry. Several example calculations show how these patterns can significantly affect the decision-making process.

Index Terms—Power system blackouts, cascading failures

I. INTRODUCTION

The goal of this paper is (1) to determine what trends exist (or do not exist) in the available historical record of large blackouts[1] in the United States and (2) to show how these trends can be important when making decisions that could impact the reliability of bulk power system infrastructures.

There are many trends that one might expect to find in these blackout data. The following are some plausible trend-related hypotheses that will be tested in this paper:

1. Technology improvements and policy changes have resulted in an observable decrease in the frequency of large blackouts, or at least in the frequency per customer.
2. There are no seasonal trends in the data. (Blackout probability does not change with time-of-year.)
3. There are no time-of-day trends in the data. (Blackout probability does not change with time of day.)
4. The fit between the blackout data and a power-law probability distribution is significantly better than the fit to an exponential distribution.
5. Due to complexity of restoration after a large blackout, there is a positive correlation between blackout size and blackout duration.

Paul Hines is with the College of Engineering, University of Vermont, Burlington, VT 05405 (e-mail: phines@cems.uvm.edu).

Jay Apt is with the Tepper School of Business and the Dept. of Engineering & Public Policy, Carnegie Mellon University, Pittsburgh, PA 15213 (e-mail: apt@cmu.edu).

Sarosh Talukdar is with the Electrical and Computer Engineering Dept. and the Dept. of Engineering & Public Policy, Carnegie Mellon University, Pittsburgh, PA 15213 (e-mail: talukdar@ece.cmu.edu).

[1] The term blackout, as used in this paper, refers to any power system event that results in multiple involuntary customer interruptions and lasts for more than 5 minutes.

Robust trends in the blackout data can often be valuable inputs to the investment and policy decision-making process. For example, a power-law relationship between event size and probability, as has been noted by [1-3], indicates that it is important to design energy delivery infrastructure to be robust to large failures rather than focusing narrowly on small failures. Section 4 provides several calculations illustrating how the observed trends may affect the decision-making process.

A. The social costs of large blackouts

When making policy or investment decisions that have the potential to impact blackout frequencies, it is often useful to translate blackout sizes into some measure of social cost, which includes direct economic costs and a monetary measure of societal losses. The Aug. 14, 2003 blackout came with substantial direct economic costs; it is estimated that there were $3 billion in insurance claims [4]. It also resulted in significant non-financial losses, such as subway passengers stranded underground and emergency vehicles stalled in traffic due to failed traffic lights. The social cost of a blackout are a function of many factors including the size of the blackout, the duration of the blackout, its location and the time of day. Clearly blackout costs increase with both the geographic extent of the event and the amount of energy that is left unserved as a result of the grid failure. The impact of a blackout in MW and MWh are useful measures for these two factors. If one had a large database of blackout costs, the following model could be fit and used to predict the costs of future blackouts given its size in MW and MWh.

$$(\text{Cost})_i = \alpha(\text{MW})_i + \beta(\text{MW})_i^2 + \gamma(\text{MWh})_i + \eta(\text{MWh})_i^2 \quad (1)$$

Certainly, other terms could be included in the prediction equation. Terms like event duration, time of day, and location could explain additional variance in event cost. One could also add additional Taylor series terms (cubic, etc). We conjecture, however, that Eq. 1 would capture a majority of the variance. Unfortunately accurate estimates for the social costs of most historical events do not exist, making it difficult to derive good estimates for α, β, γ and η. Some data exist for the commercial costs to individual customers associated with small blackouts, but these data do not provide enough information to build good estimates for the parameters α, β, γ and η. For example, from a study of 24,800 individual customer outages, Larsson et al. [5] found that reported commercial and industry customer costs increased, but not linearly, with outage duration. In this study, per kWh blackout costs increased over the first 9 hours and then decreased thereafter. A follow-up study [6] argued from the same data that much of the impact of large blackouts results from the

initial interruption (α) rather than the duration adjusted size (γ). On the other hand, after several hours the non-commercial costs of a blackout may increase substantially as services such as cellular telephone service and water distribution systems begin to fail.

If one were to perform a regression analysis using Eq. 1, one would certainly obtain positive multipliers for α and γ, since blackout costs increase with both the initial size (MW) and the duration adjusted size (MWh). The quadratic terms (β and η), on the other hand, might have opposite terms. One could argue that costs would grow super-linearly with β due to compounding social costs that come from the scale of a blackout. For example, a blackout that disabled all of the traffic lights in an entire city for 1 hour would likely be more costly than 2 blackouts that disabled 1/2 of the city's traffic lights each for 1 hour. The larger blackout might remove all alternate paths for traffic, and cause a much larger traffic problem. On the other hand, the authors of [5] argue that costs scale sub-linearly with duration (i.e. η<0). This may be the result of organizations adapting to the situation as the blackout extends in time. However, this study focused only on commercial and industrial costs, and used some data-filtering methods that may have distorted the outcome. One could also argue that per-MWh non-commercial costs begin to increase after socially valuable services begin to fail after the first day or so. Some commercial customers may recover lost income if the blackout duration is not long. Accurately estimating the costs associated with longer blackouts is challenging.

Regardless, there is broad agreement in the literature that blackout costs increase with size (MW), duration [20] and unserved energy [21, 22]. This paper thus focuses on extracting information from the time-history of blackout sizes in order to eventually provide better models for understanding blackout cost and risk.

B. Related research results

Several recent papers note useful patterns in the North American blackout data available from NERC. Carreras et al. [1,2] show that large blackout sizes follow a power-law probability distribution function (PDF). Talukdar et al. [3] show that the data fit a power-law statistic far better than they do to an exponential (Weibull) PDF. Carreras et al. [2] argue that time-correlations in the blackout data (using the Hurst parameter, which measures auto-correlation over multiple time-scales) is evidence of self-organized criticality, which would provide a plausible explanation for the power-law tail. While some have questioned the self-organized criticality conclusion, arguing that seasonal effects provide a better explanation for the clustering [7], the power-law statistic in the blackout size distribution is not disputed.

In [8], researchers study these data from the perspective of assessing the risk associated with a terrorist attack using a logistic regression model.

The analysis presented in this paper uses a more extensive data set (1984-2006) than the existing studies and filters the data in several ways to remove effects associated with demand growth, supply shortages, extreme natural events and the spotty reporting of smaller events. Without this filtering one can draw potentially misleading conclusions about the risk associated with large blackouts.

Similarly, some authors have used various theoretical blackout models to develop high-level risk measures for cascading failures. For example [9] describes a probabilistic model of cascading failure risk, which is extended in [23] describes a power system failure model that accounts for hidden failures. The trends observed in this paper could be useful input parameters to an effective blackout-risk assessment tool.

C. Overview of this paper

This paper is organized as follows. Section 2 describes the blackout data that are used in this study. Section 3 describes the trends that are found (or not found) in these data. Section 4 provides some example calculations illustrating how these trends affect decision-making. Finally, Section 5 draws some conclusions including a discussion of various explanations for the finding that the frequency of large blackouts has not decreased in time.

II. THE NERC DAWG DATA FOR 1984-2006

Both the US Department of Energy (DOE) and the North American Electric Reliability Council (NERC) require that member organizations submit reports when sufficiently large disturbances occur within their territories. DOE publishes the resulting data as ``Form 417'' reports, and NERC provides the data through its Disturbance Analysis Working Group (DAWG) database. By law, utilities and other load serving entities must report all disturbances that interrupt more than 300 MW or 50,000 customers [10]. Some smaller disturbances are also included in the reports, but on a less predictable basis. Since the NERC DAWG database [11] is the most complete of the two sources, providing data on blackouts from 1984 to 2006, the statistical analysis presented here is based on the NERC data. In total there are 933 event reports in the NERC data. In a some of these reports, several entries from different organizations refer to a single large blackout. For example, the August 14, 2003 event spans across 6 reports. In order to accurately record these blackouts, such reports are combined into a single event record. This reduces the number of events to 861. Also, the majority of the reported events are smaller than the 50,000 customer, 300 MW limit at which reporting is required by law. Since small event reporting is largely optional and spotty, this analysis focuses on the larger events. Table 1 shows some descriptive statistics for these data with and without the smaller events.

TABLE 1. DESCRIPTIVE STATISTICS FOR THE NERC DAWG DATA 1984-2006

	All	≥ 300 MW	≥ 50k cust.
Total # of events	861	277	320
Mean size in MW	584	1,706	1,111
Median size in MW	90	637	274
Std. dev. MW	3,272	5,610	5,163
Mean size in cust.	62,640	288,720	429,180
Median size in cust.	1,000	71,000	149,750
Std. dev. in cust.	87,150	1,020,200	1,076,700

A. Estimating the percentage of all customer interruptions (SAIFI) represented in the NERC data

Given the number of customers interrupted in each blackout, one can calculate the total annual customer interruptions reflected in these data. With this value, one can calculate the apparent System Average Interruption Frequency

Index (SAIFI), a measure of reliability commonly used in the electricity industry. After adjusting for demand growth, such that the data are scaled to year 2000 customers, and dividing by the number of electricity customers in the US in the year 2000, the apparent SAIFI from the NERC data is:

$$\text{SAIFI} = \frac{(219,643,512 \text{ interruptions})}{(23 \text{ years})(127,568,517 \text{ customers})} = 0.075 \quad (2)$$

According to [6] SAIFI in the United States is approximately 1.2 or 1.3. Thus the disturbances in the NERC data represent about 6% of all customer interruptions recorded in the national SAIFI numbers. Since utilities do not report all blackouts in their SAIFI numbers, it is likely that the NERC data represent somewhat less than 5% of all US customer interruptions. It is important to note this fact, because conclusions drawn from these data may be biased by the fact that they only represent a twentieth part of all customer interruptions and, since most blackouts are very small, a much smaller portion of all blackouts.

B. Blackout initiating events

Disturbances recorded in the DAWG data proceed from a wide variety of triggering events including natural disasters, storms, human error and mechanical failure. For this analysis we categorize disturbances into the following initiating-event categories: earthquakes, tornados, hurricanes or tropical storms, ice storms, lightning, wind or rain storms, other cold weather, fire, intentional attack, supply shortage, other external (not-human or equipment) event, equipment failure, operator error, voltage reduction and volunteer power reductions. Some records indicate multiple initiating events. It is important to note that some events were initiated by a natural cause, such as lightning, but grew through a set of cascading outages or operator errors. It is difficult to precisely isolate the events that were exacerbated by cascading failures with certainty, thus making it difficult to calculate the national impact of cascading failures. Still, these data allow us to calculate a rough upper bound on the historical impact of cascading failures, assuming that all of the large cascading failures are included in the NERC records.

Figure 1 shows the relative frequency of blackouts in each cause category. Table 2 provides the same results in tabular form and includes average event sizes.
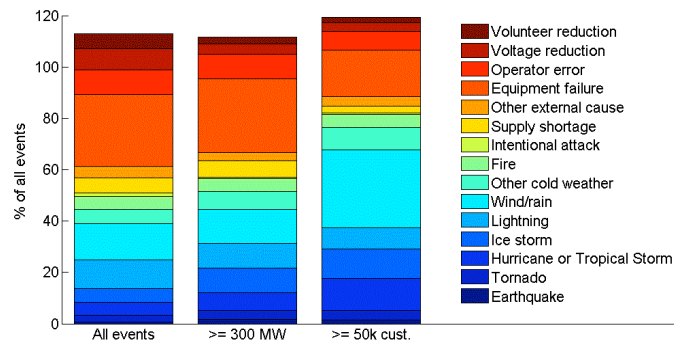

Figure 1. The proportion of blackouts and disturbances in the NERC DAWG data for 1984-2006 with various initiating events. The totals are greater than 100% because some events fall into multiple initiating-event categories.

C. Data filtering

In order to ensure the reliability of the trend-analysis described in Section 3, the data are filtered in several ways. Firstly, we remove events in the "volunteer reduction" and "voltage reduction" categories since these do not generally disrupt electricity service and reporting may be unreliable. Secondly, some of the records include event size in either MW or customers, but not both. To avoid the loss of potentially useful data, missing MW or customer entries were filled according to average customers per MW from the records where both are provided (about 465 customers per MW). Several other filtering methods are used specific to the individual trend analyses described in Section 3.

TABLE 2. STATISTICS FOR DATA CAUSE CATEGORIES

	% of events	Mean size in MW	Mean size in customers
Earthquake	0.8	1,408	375,900
Tornado	2.8	367	115,439
Hurricane/tropical storm	4.2	1,309	782,695
Ice storm	5.0	1,152	343,448
Lightning	11.3	270	70,944
Wind/rain	14.8	793	185,199
Other cold weather	5.5	542	150,255
Fire	5.2	431	111,244
Intentional attack	1.6	340	24,572
Supply shortage	5.3	341	138,957
Other external cause	4.8	710	246,071
Equipment failure	29.7	379	57,140
Operator error	10.1	489	105,322
Voltage reduction	7.7	153	212,900
Volunteer reduction	5.9	190	134,543

III. TRENDS IN THE BLACKOUT DATA

In the introduction we posit five hypotheses regarding trends in the blackout data. Sections A, B and C below focus on time-trends in the data and Sections D and E focus on the power-law trend and the relationship between blackout size and restoration time.

A. Blackout frequency has decreased with time

In this section we test the hypothesis that the frequency of large blackouts is decreasing with time. Figure 2 shows the number of blackouts per year in various size categories (with size measured in MW) for all 933 event records. At first glance the data indicate a clear increase in the number of blackouts per year. Beginning in 1995 apparent blackout frequency increases almost linearly. But a good portion of this increase occurs within the smaller event categories, indicating that the observed increase could be the result of increased voluntary reporting of smaller blackouts. Also, it is plausible that some of the increase could be a simple result of increases in demand. To control for these two factors and test our hypothesis, we normalize event sizes by either electricity demand[2] or population[3] and eliminate small events (those smaller than 400 MW or 100,000 customers after normalizing). Figure 3 shows the blackout frequency after filtering the data for demand growth and small events. While

---

[2] Electricity demand is measured using net generation as reported to the US-EIA (www.eia.doe.gov).

[3] Population is taken from US census annual population estimates (www.census.gov).

the observed increase is largely eliminated, there is no observable increase.

In order to test our hypothesis more thoroughly, several statistical tests were employed. Table 3 shows the results of several Kolmogorov-Smirnov (K-S) t-tests, testing the hypothesis that two sets of non-Gaussian data (event frequency before and after the mid-point year) come from the same probability distribution. While the results do not show a clear increase in blackout frequency, one can conclude with some certainty that blackout frequency is not decreasing in time, and thus reject our hypothesis.
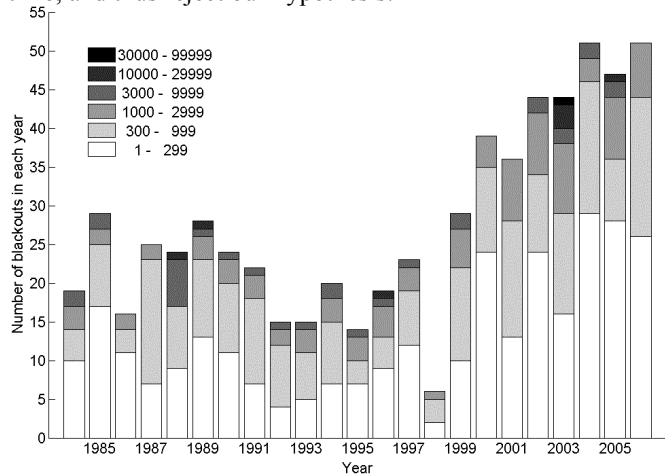

Figure 2. The number of blackouts recorded during each year in the full data set of 933 events. Note that the number of small blackouts reported has increased significantly, but the number of large blackouts is relatively constant.
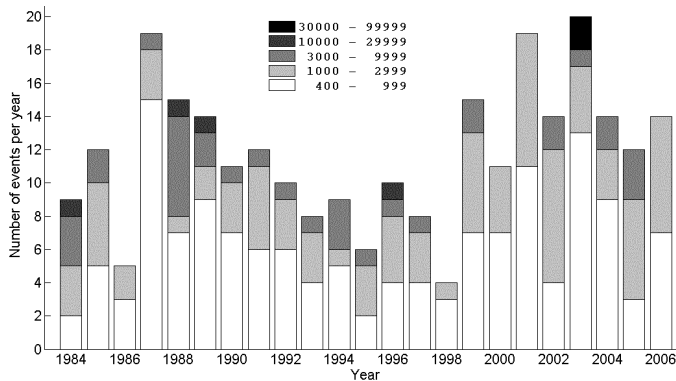

Figure 3. The number of large blackouts per year after removing small events, and controlling for increasing demand. Event size above is shown in year-2000 MW.

TABLE 3. STATISTICAL TESTS FOR THE HYPOTYESIS THAT BLACKOUT FREQUENCY IS DECREASING WITH TIME.

Data[5]	Correlation[4]		'84-'95 Median[6]	'96-'06 Median	P from K-S t-test[7]
	ρ	P			
≥50k cust.	0.59	0.003	10	18	0.047
≥50k y2k cust.	0.46	0.029	10	15	0.147
≥100k cust.	0.53	0.009	7	10	0.147
≥100k y2k cust.	0.34	0.112	8	10	0.985
≥300 MW	0.42	0.046	8	10	0.736
≥300 y2k MW	0.16	0.457	10	10	1.000
≥500 MW	0.40	0.059	5	8	0.147
≥500 y2k MW	0.09	0.690	7	8	0.736

B. Seasonal trends

In this section we test the hypothesis that blackout frequency does not change seasonally. To do so, we use a rolling average to obtain the approximate number of blackouts per month for each month, using three-month windows for averaging. The rolling-average smoothes out some of the noise in the monthly average, showing a clear seasonal trend in the data. Blackout frequency increases substantially during the late summer and mid-winter months. Clearly there is a significant seasonal trend in the blackout data.
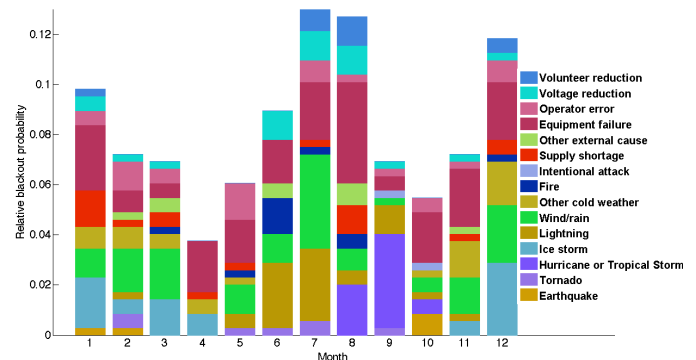

Figure 4. Blackout frequency as a function of time-of-year. Blackouts are more frequent during the late-summer and mid-winter months.

C. Time-of-day trends

In this section we test the hypothesis that blackout probability does not change with the time-of-day. Figure 5 shows a rolling 3-hour hour blackout counts for all 24 hours in the day (after excluding the events that do not have a start time recorded). This figure clearly shows that blackout probability increases substantially during the peak hours. In fact blackouts are about 4 times more likely to start during the late afternoon hours, relative to the early morning hours. Two explanations for this observation are defensible. Storm activity typically increases during the mid-afternoon hours, which may account for some of the increase. In fact, Fig. 5 shows a substantial increase in weather related events during the mid-afternoon hours. Alternatively (or perhaps additionally), this may be the result of power networks being more stressed during mid-afternoon

---

[4] Correlation measures the relationship between the year and the number of blackouts during that year.

[5] "y2k" sizes indicate that the data were scaled to account for demand or population growth using 2000 as a base year.

[6] Medians indicate the median events per year.

[7] The K-S P value measures the probability that the observed differences between the '84-'95 and '96-'06 data are due to chance.

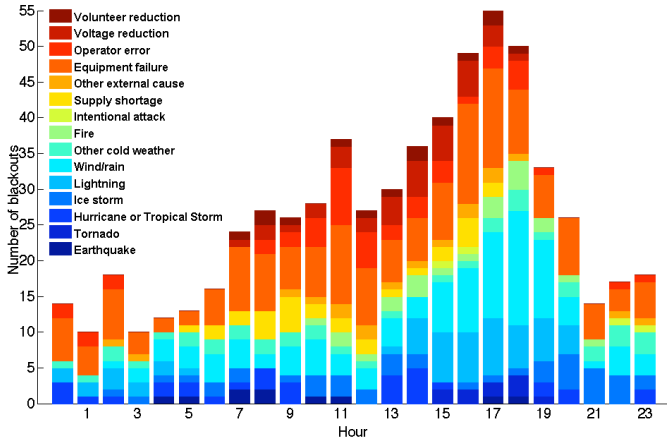hours, indicating proximity to critical points at which blackout probability increases sharply [2, 12, 24].

Figure 5. Blackout frequency as a function of time-of-day. The vertical axis shows the total number of events in 3-hour windows about each hour. This figure indicates that blackouts are 3-4 times more likely during peak hours relative to the early morning hours.

D. Power-laws

It is well known that the sizes of large blackouts in the United States follow a power-law probability distribution [1-3]. International blackout data also show a power-law size-frequency relationship [24], indicating that this relationship is fundamental to the structure of power grids.

There are a number of forms of the power-law probability distribution, but one of the most commonly employed is the Pareto distribution; named after Vilfredo Pareto who found that the distribution of wealth followed a power-law probability distribution. The cumulative distribution function (cdf) for a random variable x with minimum value x_{min}, which follows a Pareto distribution, can be written as follows:

$$P(x \leq X) = 1 - \left(\frac{X}{x_{min}}\right)^k \qquad (2)$$

where k is the scaling exponent. The probability density function (pdf) is:

$$P(x = X) = \frac{k x_{min}^k}{X^{k+1}} \qquad (3)$$

and the expected value (mean) is:

$$E[x] = \begin{cases} \dfrac{k x_{min}}{k-1}, & k > 1 \\ \infty, & k > 1 \end{cases} \qquad (4)$$

Fig. 6 shows the power-law relationship among the sizes of the largest events in the event size data. When fitting the data to a Pareto distribution one obtains exponents (k) in the range of 1.14-1.20. The largest events in this set, particularly with size measured in MW, show an excellent fit to a power-law probability distribution (R^2=0.997 for the scaled MW data).
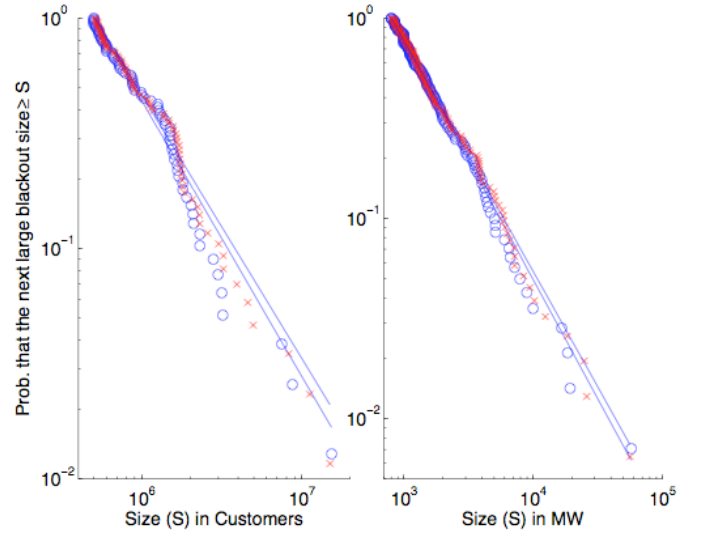

Figure 6. The cumulative probability distribution of blackout sizes in customers (left, for events ≥500k cust.) and MW (right, for events ≥800 MW). X marks indicate blackout sizes adjusted for population/demand growth. O marks indicate un-scaled data. The lines show the power-law fit to the data.

E. Relationship between blackout size and restoration time

Many large blackouts require extended restoration periods. It can take days or event weeks to restore customers after the losses from a natural disaster. Due to difficulties associated with starting large power-plants without off-site power, restoration can take many hours, even if no equipment damage has occurred. It thus seems reasonable to expect to find correlation between the size of a blackout and it's duration.

Table 4 shows the result of correlation tests. In both the MW and customer cases only the events that include both duration and the size measure (MW or customers) were included in the analysis. Surprisingly only a weak, statistically insignificant correlation exists between event size and duration. This is also clear in Fig. 7, which shows no relationship between blackout size and restoration time. In a few cause-categories, there is a significant positive correlation between size and duration (lightning, wind/rain, and "other external cause"), but in some other categories a weak (insignificant) negative correlation exists resulting in no significant correlation for the data set as a whole.
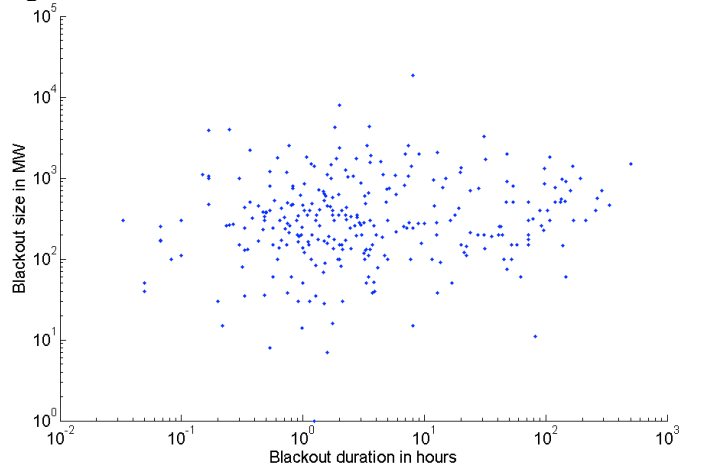

Figure 7. Blackout size in MW plotted against blackout duration. The two variables are almost perfectly uncorrelated. Correlation statistics shown in Table 4 shows the linear correlation. The log-log scale is used here for clarity—the linear scaled figure shows a similar relationship.

TABLE 4. CORRELATION TESTS FOR THE HYPOTHESIS THAT LARGE BLACKOUTS RESULT IN LONG RESTORATION TIMES

	N	Corr coef. (ρ)	P-value
With size in MW	388	0.0135	0.7907
With size in customers	349	0.0971	0.0701

To summarize, we find (A) that blackout frequency has not decreased from 1984 to 2006, (B) that blackouts are substantially more frequent in the summer and winter and (C) during mid-afternoon hours, (D) that large blackouts occur much more frequently than would be expected from a exponential statistics, and (E) that there is no apparent correlation between blackout size and restoration time.

IV. RELEVANCE OF TRENDS TO POLICY PROBLEMS

The following section describes how the trends described in this paper can significantly affect the decision making process. In all of the calculations below we assume that blackout cost scales linearly with blackout size as measured in MW, essentially assuming that $\beta=0$, $\gamma=0$, and $\eta=0$ in Eq. 1. While an imperfect assumption, the lack of correlation between blackout size and duration make this assumption somewhat more reasonable. If there is no correlation between size and duration and the error is Gaussian, the expected size of a randomly chosen blackout in MWh will be a constant function of its size in MW.

A. Relevance of the fact that blackout frequency is not decreasing in time

This is important because it allows us to assume that without dramatic changes within the industry blackout risk will remain roughly constant. This assumption is employed in the expected value calculations that follow in Sections B and C.

B. Relevance of the time-of-day and time-of-year trends

Because blackout risk changes with time-of-day and time-of-year, it is valuable to focus the bulk of our blackout risk reduction efforts on peak periods. Consider a utility considering two policy changes. Under option 1 it deems that by doubling the number of on duty operators during all hours it can reduce the blackout probability at all hours by 50%, without changing the blackout size distribution. The cost of this option is C_1. Under option 2 it deems that it can reduce blackout period during the highest risk hours by doubling is operator staff during only these hours. Let us assume that it can increase its staff during 50% of all hours for 50% of C_1 ($C_2=C_1/2$). The cost of option 2 is C_2. Let $P(B|h)$ and $P'(B|h)$ represent the probability of a blackout at hour h before and after the policy change, c_b be the average per MW blackout cost and E[S] be the expected value for the utility's blackout size distribution. The expected value of either decision can be calculated as follows:

$$E[V] = c_b \sum_{h=1}^{8760} \left(P(B|h) - P'(B|h) \right) E[S]^8 \qquad (5)$$

Because c_b and E[S] remain unchanged, option 1 will reduce blackout costs in the utilities service area by 50%. Option 2 however will not reduce risk as much, but does so during peak

---

[8] This equation assumes that blackout size and blackout probability are independent.

---

periods. If option 2 reduces risk by 50% during the 12 highest risk hours (7am-7pm) given the time-of-day blackout frequency trend shown in Fig. 5, option 2 would reduce overall blackout costs by 34%. The utility can obtain most of the desired risk reduction with half the cost by doubling staff only during peak periods.

C. Relevance of the power-law relationship between blackout size and frequency

The existence of a power-law probability distribution is important because it indicates that large events are substantially more common than one would predict from exponential distributions such as a Gaussian or Weibull, which are commonly used in engineering reliability analysis. The end result is that a blackout of any size (up to the extent of the entire network) has a non-zero probability. More practically, this result indicates that blackout mitigation efforts should focus on the largest events in nearly equal proportion to the smaller events. As evidence of this effect, figure shows the relative contributions from blackouts of various sizes to the overall impact of large transmission system blackouts. Another effect of the power-law distribution is apparent when calculating the size of a 100-year blackout, using methods commonly used for storm impact assessment. Given that the sizes of large blackouts (size greater than 800 MW) follow a Pareto distribution with k=1.15, and given that an event equal in size to Aug. 14, 2003 occurs once in every 23 years (the extent of the available data) the following gives the size of the 100 year blackout:

$$S_{23} = 56,465\,MW$$

$$P_{23} = P(S \geq 56,465 \mid S \geq 800) = 0.006$$

$$P_{100} = P_{23}\left(\frac{23\,years}{100\,years}\right) = 0.00138$$

$$S_{100} = 800(10^{-\log P_{100}/1.15}) = 246,000\,MW$$

where S_X indicates the size of the X-year event, with P_X giving the conditional probability of an event with this size. By comparison, according to DOE/EIA data, the peak demand (EIA: ``Net Internal Demand'') for the continental US in 2000 (the base year for the size measures) was 681,000 MW. Thus, if the observed statistical pattern holds for very large blackouts, and if the US were to see a 100-year blackout next year, it would interrupt about one third of all electricity service in the continental US.

In terms of Eq. 5, the power-law relationship between size and probability can significantly effect the calculation of E[S]. Let us assume that one would like to evaluate a technology that will reduce the probability of all blackouts larger than 300MW by a fixed portion. If we calculate E[S] over this range by just taking the average blackout size from the observed data we get E[S] = 1744 MW. However, if we use the power-law statistic to get the expected size, we get E[S] = 2400 MW. By ignoring the power-law relationship one would undervalue this technology by about 30%.

V. CONCLUSIONS

In this paper we show that the frequency of large blackouts in the United States has not decreased during the years 1984-2006, for which data is available. In fact by some measures, a

slight increase in blackout frequency is observable. This is true, despite the fact that the electricity industry has invested substantially to improve system-wide reliability.

We also show that other trends do exist in the historical blackout data. Blackout frequency changes substantially with the time of day and the time of year, and in confirmation of existing results, there is a clear power-law relationship between blackout size and probability. Several example calculations show that these trends can have a significant effect on the value of various decisions within the industry.

A. Explanations for the lack of reduction in blackout frequency

Given the lack of a significant decrease in the frequency of cascading failures and given the significant investment in technologies and policies to control the risk of such failures, it is natural to look for an explanation for the existing trend (or lack thereof). Unfortunately the data alone do not provide an explanation, as the granularity is not sufficient to empirically evaluate the effects of any particular policy or technical change. Nevertheless, some discussion of the commonly espoused explanations may be useful.

1) Market restructuring

The restructuring of the electricity industry, beginning with FERC Order 888 which required open access to transmission capacity, has been blamed for numerous problems in the US electricity industry. While it is likely that open access has resulted in additional use of transmission paths for long distance transfers, it is difficult to say from these data that restructuring has had a direct effect on blackout risk. Even if the increased use of transmission infrastructure has increased the risk somewhat, this explanation does not help much in providing a solution as it would be very difficult for the industry to return to a market structure with substantially less open access to transmission services. Any such analysis would require state-level or LSE-level CAIDI, SAIDI, and SAIFI statistics, which are generally not easily accessible (and in some states require a Freedom of Information Act request).

2) Inadequate transmission investment

Industry members often assert that a lack of transmission system investment has led to unsatisfactory performance of the transmission system. The national transmission grid study notes that the frequency of transmission loading relief (TLR) events (a rough measure of system stress) has increased simultaneously while transmission system investment has decreased. Hirst [13] shows that the quantity of available transmission has, over the period (1999 - 2002), steadily decreased when normalized by summer peak demand. Vajjhala and Fischbeck [14] show that in many US states where new transmission is most needed, it is particularly difficult to build new transmission.

On the other hand, perhaps due to the attention that this issue has received, transmission investment has increased fairly steadily since 1999 [13]. And there are many ways to increase the capability of transmission systems without actually building new lines. Composite conductors can increase the thermal ratings, and phase-shifting transformers or FACTS devices can relieve bottleneck constraints by changing the apparent impedance of transmission lines. Finally, Blumsack [15] shows that some transmission

construction can have a negative impact on reliability. While transmission investment can, but is not guaranteed to [25], have a positive impact on cascading failure risk, and reliability, transmission constriction alone is a costly, and potentially ineffective, solution to reliability problems.

3) A lack of enforceable reliability rules and system-wide management

After the August 14, 2003 blackout, many in industry and government argued that the voluntary reliability rules, as established and operated by NERC, were an insufficient instrument for managing reliability in a competitive electricity industry. As a result of this discussion, the Energy Policy Act of 2005 gave FERC the authority to appoint an Electricity Reliability Organization (the role that NERC now fills), with the authority to design and enforce mandatory reliability rules nationwide.

Relatedly, Apt et al. [16, 17] argue that insufficient system-wide management of the electricity network (similar to FAA's management of the air-traffic control system in the US) contributes to the overall blackout risk. Apt et al. argue that a systems approach to risk mitigation has significantly reduced the accident frequency in commercial air travel, and that similar actions within the US electricity system could result in similar risk reductions.

The events of Aug. 2003 do provide some evidence that unenforceable reliability rules contributed to the cascading failure. But, since 2003 FERC and NERC have been fairly diligent in their oversight of transmission assets, and the annual number of large blackouts remains relatively constant. While systems policy changes are necessary to solving the blackout problem, they are not sufficient. The cascading failure problem is the result of both policy and technology failures.

4) System protection and problem formulation

Another explanation, argued in [Hines et al, CMU T&D], is that the design of the protection system in electrical power networks is poorly aligned with the objectives of the system as a whole. Protective relays remove components from the network when they experience significant stress. While this approach effectively ensures that problems in the transmission system will not damage equipment, the strategy is frequently sub-optimal with respect to the goal of the system as a whole---delivering energy to customers. A better strategy would protect the equipment, while also seeking to deliver energy to customers. This paper is an attempt to correct this problem through the design of a decentralized control system that aligns the goals of the components with the goals of the system as a whole.

B. Future work

The observed blackout frequencies, when sorted according to time-of-day, time-of-year, location, cause, etc., can be used to create a set of conditional blackout probabilities. When combined one could use these probabilities to create a Baysian risk model. The risk model could be used to calculate the relative likelihood of a blackout given a set of input factors. This could be used as a prior probability and adjusted for real-time conditions to give information to operators in real time. The risk model could also be used to project staffing requirements, schedule maintenance, or even for short-term

investment planning. Future work will focus on the design of such a risk model.

VI. Acknowledgment

VII. References

[1] B. A. Carreras, D. E. Newman, I. Dobson, and A. B. Poole, "Initial evidence for self-organized criticality in electric power system blackouts," in *Proceedings of Hawaii International Conference on System Sciences*, 2000.

[2] B. A. Carreras, D. E. Newman, I. Dobson, and A. B. Poole, "Evidence for Self-Organized Criticality in a Time Series of Electric Power System Blackouts," *IEEE Transactions on Circuits and Systems I: Regular Papers*, v. 51, pp. 1733-1740, 2004.

[3] S. Talukdar, J. Apt, M. Ilic, L. Lave, and M. Morgan, "Cascading failures: survival versus prevention," *The Electricity Journal*, 2003.

[4] J. B. Treaster, "The blackout: Business losses; insurers say most policies do not cover power failure," *The New York Times*, p. 6, August 16, 2003.

[5] L. Lawton, J. H. Eto, A. Katz, and M. Sullivan, "Characteristics and Trends in a National Study of Consumer Outage Costs," *Proc. of the Center for Research in Regulated Industries 16th Annual Western Conference,* 2003.

[6] K. H. LaCommare and J. H. Eto, "Cost of Power Interruptions to Electricity Consumers in the United States," *Energy*, v. 31, n. 12, 2005.

[7] R. Weron and I. Simonsen, "Blackouts, risk, and fat-tailed distributions," in *Proceedings of the 3rd Nikkei Econophysics Symposium*, 2005.

[8] J. S. Simonoff, C.E. Restrepo, and R. Zimmerman, "Risk-Management and Risk-Analysis-Based Decision Tools for Attacks on Electric Power," *Risk Analysis*, v. 27, n. 3, 2007.

[9] I. Dobson, B. A. Carreras, and D. E. Newman, "A loading-dependent model of probabilistic cascading failure," *Probability in the Engineering and Informational Sciences*, vol. 19, pp. 15-32, 2005.

[10] "Form OE-417, Electric emergency incident and disturbance report," U.S. Dept. of Energy, Office of Electricity Delivery and Energy Reliability, online: http://www.eia.doe.gov/oss/forms.html, 2005.

[11] "DAWG Database: Disturbances, Load Reductions, and Unusual Occurrences," North American Electric Reliability Council, Disturbance Analysis Working Group, online: http://www.nerc.com/ dawg/, accessed Nov. 2007.

[12] H. Liao, J. Apt, and S. Talukdar, "Phase Transitions in the Probability of Cascading Failures," in *Proc. of the First Annual Carnegie Mellon Conf. on the Electricity Industry, Electricity Transmission in Deregulated Markets: Challenges, Opportunities, and Necessary R&D Agenda*, 2004.

[13] E. Hirst, "U.S. Transmission Capacity: Present Status and Future Prospects," Edison Electric Institute and Office of Electric Transmission and Distribution, U.S. Department of Energy, Aug. 2004.

[14] Vajjhala, S.P. and P.S. Fischbeck, Quantifying siting difficulty: A case study of US transmission line siting. *Energy Policy*, 2007. 35(1): p. 650-671.

[15] S. Blumsack, L. Lave and M. Ilic, "A Quantitative Analysis of the Relationship Between Congestion and Reliability in Electric Power Networks," *Energy Journal*, Oct, 2007.

[16] J. Apt, L. B. Lave, S. Talukdar, M. G. Morgan, and M. Ilic, "Electrical blackouts: A systemic problem," *Issues in Science and Technology*, vol. 20, no. 4, pp. 55-61, 2004.

[17] Apt, J., L.B. Lave, and M.G. Morgan, "A More Reliable U.S. Electric System." *Issues in Science & Technology*, 2006. 22(4): p. 51-58.

[18] P. Hines, "A Decentralized Approach to Reducing the Social Costs of Cascading Failures," Ph.D. Dissertation, Carnegie Mellon University, 2007.

[19] P. Hines, J. Apt, H. Liao, and S. Talukdar, "The frequency of large blackouts in the United States electrical transmission system: an empirical study," in *Proc. of the Second Carnegie Mellon Conference in Electric Power Systems: Monitoring, Sensing, Software and Its Valuation for the Changing Electric Power Industry*, 2006.

[20] CIGRE Task Force 38.06.01, Methods to consider customer interruption costs in power system analysis, CIGRE report 191, August 2001.

[21] R. Billinton, J. Otengadjei, R. Ghajar, Comparison of 2 alternate methods to establish an interrupted energy assessment rate, IEEE Transactions Power Systems, vol. 2 no 3, pp. 751-757 August 1987.

[22] R. Billinton, R.N. Allan, Reliability evaluation of power systems, second edition, Chapter 13, Plenum Press, New York, 1996.

[23] 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.

[24] 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, 026103, June 2007.

[25] D. Kirschen, G. Strbac, Why investments do not prevent blackouts, Electricity Journal, March 2004, pp. 29-34.

VIII. Biographies

Paul Hines is an Assistant Professor in the School of Engineering at the University of Vermont (UVM). He received a Ph.D. in Engineering and Public Policy from Carnegie Mellon University in 2007 and an M.S. in Electrical Engineering from the U. of Washington in 2001. Formerly he worked as a Research Scientist at the National Energy Technology Laboratory, for the Federal Energy Regulatory Commission, where he studied interactions between nuclear power plants and transmission networks, for Alstom ESCA, where he designed a short-term load-forecasting tool and for Black and Veatch, where he worked on substation design projects. His main research interests are in the areas of complex systems and networks, the control of cascading failures in power systems and energy system reliability.

Jay Apt is Executive Director of the Carnegie Mellon Electricity Industry Center at CMU's Graduate School of Industrial Administration and the Department of Engineering and Public Policy, where he is a Distinguished Service Professor.

Sarosh Talukdar is Professor Emeritus in the departments of Electrical and Computer Engineering and Engineering & Public Policy at Carnegie Mellon University.