Analyzing Cascading Failures and Blackouts Using Utility Outage Data

Ian Dobson
Iowa State University, Ames IA 50011 USA
e-mail: dobson@iastate.edu
November 2023

This is an author preprint to appear as Chapter 2 of "Cascading Failures in Power Grids: Risk Assessment, Modeling, and Simulation", edited by Kai Sun, 2024, Springer Nature Switzerland AG https://doi.org/10.1007/978-3-031-48000-3_2.

Abstract Historical utility data for cascading failure and large blackouts is foundational for understanding and quantifying blackouts. This chapter surveys some of the main ideas in obtaining and exploiting the patterns in this data, beyond the useful lessons that can be learned from each particular blackout. Historical data on blackout size shows a heavy tailed distribution that implies that large blackouts are both rare and will occasionally occur, and that their risk is substantial. Detailed outage data is routinely collected by utilities and can be processed into cascading events or weather-related events. Metrics for these events can then be readily obtained. Almost all of the research on cascading is based on models and simulation, despite the promising and emerging opportunities to also learn from utility outage data. Some of these opportunities are outlined to encourage further work on real data: As well as its obvious key use to ground models and simulations in reality, the historical outage data can enable better contingency lists, be transformed into influence or interaction graphs, replace simulation by sampling, and describe typical blackouts and their restoration with Poisson processes.

Keywords electric power transmission system, utility data, blackout, simulation, sampling, restoration, outages, contingency list, motifs, electric power transmission system, Poisson process, risk, power law, resilience

1 Analysis of particular blackouts

An influential and useful approach to large blackouts considers each specific blackout, analyzes the sequence of events and their interrelated causes in detail, and extracts lessons learned that can be implemented to prevent that particular blackout, or a similar blackout, from happening again. Every large blackout has such analysis and reporting, although the public reporting often is of a summary nature with very limited raw data available for independent analysis. Given the large variety of phenomena involved in cascading blackouts, the details of the analysis and the combination of causes varies considerably. Useful examples of these analyses are [1, 2, 3, 4, 5]; and some of these blackout narratives also describe the mechanisms involved.

At a higher level, the evolution of large blackouts tends to have features in common such as a long complicated chain of unusual events and interactions. Moreover, blackout data shows that there are patterns in blackout statistics that we discuss in the following sections.

2 Probability and risk of large blackouts estimated from data

Many countries record the size of transmission system blackouts [6, 7, 8]. For example, the United States publicly records blackouts above a particular size. The blackout size in MW power interrupted as well as the blackout duration and the number of customers affected are often recorded. The blackout duration is less useful than the other measures since blackout duration suffers from a lack of uniform definition, depends heavily on the last few elements restored, and is inherently more variable than the blackout size [9].

The empirical distribution of blackout size is best presented as a survival function on a log-log scale as shown in Fig. 1. The survival function is the probability of a blackout exceeding a given size as a function of the size. (The survival function is is also called the complementary cumulative distribution function (CCDF) and is one minus the cumulative distribution function (CDF).) The survival function shows the data without the smoothing or binning needed for a probability density function, and the log-log scale shows the large blackouts of small probability and the overall pattern of how the blackout probability decreases as blackout size increases. It is also more useful to consider blackout probability within a given factor (e.g. to within a factor of two or one half) as given by an equal distance above or below the data on the logarithmic vertical axis.

The approximately straight line behavior in the log-log plot of the distribution of blackout size in Fig. 1 shows that the distribution of blackout size has a “heavy tail" or power law region, which is influential on blackout probability and risk. (The power law region is of course limited in its highest extent because every grid has a largest possible blackout in which the entire grid blacks out.) The slope of approximately -1 in the survival function in Fig. 1 implies that doubling the blackout size only halves the probability that a blackout exceeds that size. The slope of –1 in the survival function corresponds to a slope of –2 in a log-log plot of the probability density function, which implies that a blackout of double the size has only one quarter of the probability. This relationship shows the typical heavy tail dependence that as blackout size increases, probability decreases, but it decreases very slowly. In contrast, for a probability distribution with an exponential tail (such as exponential or normal distributions), doubling the blackout size at most squares the probability, so that large blackouts are vanishingly unlikely. A probability distribution with heavy tails indicate that blackouts of all sizes, including large blackouts, can occur. Large blackouts are rare, but they are expected to happen occasionally, and they are not "perfect storms." These qualitative remarks are further supported by a more precise description in subsection 2.1.

The approximate power law dependence of blackout size has been has been confirmed by observed transmission blackout size statistics in many developed countries [6, 7, 8]. Cascading failure is a sequence of dependent events that successively weaken the power system. At a given stage in the cascade, the previous events have weakened the power system so that further events are more likely. It is this dependence that makes the long series of cascading events that cause large blackouts likely enough to pose a substantial risk. (If the events were independent, then the probability of a large number of events would be the product of the small probabilities of individual events and would be vanishingly small.)

The power law region can be explained using ideas from complex systems theory. The main idea is that over the long term, the power grid reliability is shaped by the engineering responses to blackouts and the slow load growth and tends to evolve towards the power law distribution of blackout size [6, 10]. It is notable that the power law region appears both in blackouts attributed to weather and blackouts not associated with weather [11].

Blackout risk can be defined as the product of blackout probability and blackout cost. One simple assumption is that blackout cost is roughly proportional to blackout

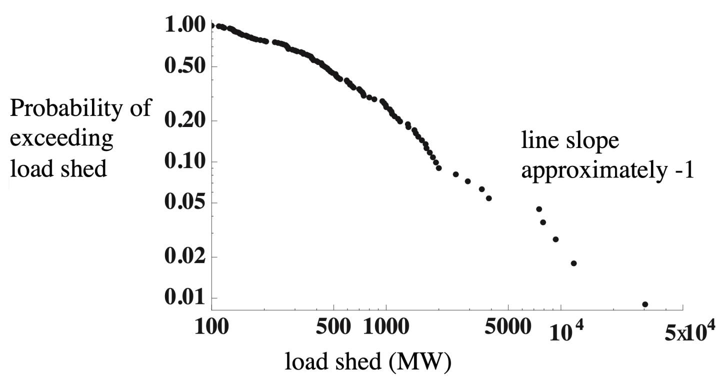

Fig. 1 Empirical survival function of blackout size from Western interconnection in the USA from 1984 to 2006.

size, although larger blackouts are expected to have costs (especially indirect costs) that increase faster than linearly. In the case of the power law dependence, the larger blackouts can become rarer at a similar rate as costs increase, and then the risk of large blackouts is comparable to or even exceeding the risk of small blackouts [8]. Mitigation of blackout risk should consider both small and large blackouts, because mitigating the small blackouts that are easiest to analyze may inadvertently increase the risk of large blackouts [12].

2.1 How the blackout probability decreases as size increases

Since the way that blackout probability decreases as blackout size increases is so important to large blackout risk [8] and justifying the study of large cascades, it is useful to describe an idealized form of this relationship more precisely. Suppose that the blackout size above some minimum blackout size $x_{\min}$ MW is the Pareto random variable X with the probability density function

$$f_X(x) = \frac{x_{\min}}{x^2}, \quad x \geq x_{\min} \tag{1}$$

Then

$$\log f_X(x) = -2\log x + \log x_{\min}, \quad x \geq x_{\min} \tag{2}$$

so that the plot of $\log f_X(x)$ versus $\log x$ is linear with slope –2 as shown in Fig. 2.

Recall, for example, that probability density $f_X(500)$ per MW at $X = 500$ MW means that the probability of a blackout between 500 MW and 501 MW is $f_X(500)$ and the probability of a blackout between 500 MW and 500+dy MW is $f_X(500)dy$. More generally, the probability of a blackout between a MW and b MW is $\int_a^b f_X(y)dy$. The survival function $\overline{F}_X(x)$ is the probability that the blackout size exceeds x so that

$$\overline{F}_X(x) = \int_x^\infty f_X(y)dy = \int_x^\infty \frac{x_{\min}}{y^2}dy = \frac{x_{\min}}{x}, \quad x \geq x_{\min} \tag{3}$$

And

$$\log \overline{F}_X(x)) = -\log x + \log x_{\min}, \quad x \geq x_{\min} \tag{4}$$

so that the plot of $\log \overline{F}_X(x)$ versus $\log x$ is linear with slope –1.

Then doubling the blackout size from x to $2x$ gives one quarter of the probability density and one half of the survival function:

$$f_X(2x) = \frac{x_{\min}}{(2x)^2} = \frac{1}{4} f_X(x) \tag{5}$$

$$\overline{F}_X(2x) = \frac{x_{\min}}{2x} = \frac{1}{2} \overline{F}_X(x) \tag{6}$$

It is straightforward to repeat this calculation more generally with a power law x^{-s} of the tail of the probability density of blackout size. Then the survival function has a power law tail x^{1-s}. And on log-log plots the slope of the probability density is $-s$ and the slope of the survival function is $1 - s$. Doubling the blackout size gives 2^{-s} times the probability density and 2^{1-s} times the survival function.

In contrast, one can consider a probability distribution with an exponential tail such as

$$f_X(x) = \alpha e^{-\alpha x}, \quad x \geq 0 \tag{7}$$

$$\overline{F}_X(x) = e^{-\alpha x}, \quad x \geq 0 \tag{8}$$

Then doubling the size from x to $2x$ has an effect proportional to squaring both the probability density and the survival function:

$$f_X(2x) = \alpha e^{-2\alpha x} = \frac{1}{\alpha} (f_X(x))^2 \tag{9}$$

$$\overline{F}_X(2x) = e^{-2\alpha x} = (\overline{F}_X(x))^2 \tag{10}$$

It follows that if the tail of blackout size were exponential (which it is not), this squaring effect would have made large blackouts vanishingly unlikely as illustrated in Fig. 2.

3 Processing utility outage data

This section describes detailed utility outage data and how to automatically process a series of outages into events or cascades, and further into successive generations of outages within the cascades. The event processing is fundamental for studying cascading and resilience with real data. This differs from the traditional processing for reliability, which tends to study the outages and repairs in a steady state that is averaged over the year.

It is important to note that utility outage data includes events of all sizes and all causes. In particular, large events include those mainly caused by cascading outages within the transmission grid as well as those mainly caused by extreme weather. The outages caused by extreme weather can be identified by their location and timing matching weather records or by cause codes in utility data. Most of the large blackouts are associated with extreme weather [13] (cascading effects could also contribute to the weather blackouts, but the extent of this contribution is not known). The blackouts caused by extreme weather are studied in resilience whereas the routine outages occurring singly or in small groups are dominant in the study of reliability, partly

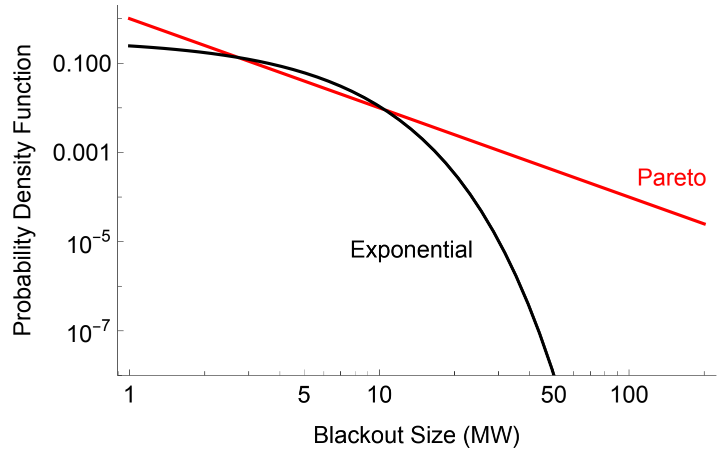

Fig. 2 Comparing Pareto and exponential probability density functions on a log-log plot showing that large size events have vanishingly small probability for exponential whereas the more slowly decreasing Pareto can have rare large events. Slope of Pareto probability density function on this log-log plot is –2. Pareto (Eq. 1) has $x_{\rm min} = 1$ and exponential (Eq. 8) has rate $\alpha = \ln\sqrt{2}$ so that both distributions have median 2.

because small outages are much more numerous. The study of cascading includes large blackouts that can be considered part of resilience and can also include short cascades with a few outages that quickly stop. There are significant overlaps and no clear boundaries between cascading, resilience, and reliability. The important point for data analysis is that if one considers all the automatic outages observed in some real utility data, they encompass and exemplify aspects of cascading, resilience and reliability, and the differences between these categories mostly depend on the processing and interpretation of the data. Detailed utility outage data are foundational for the study of cascading, resilience, and reliability.

While the focus here is on transmission system outages, it should be noted that Murphy [14, 15] describes correlated generator outages with detailed utility generator outage data from the Generator Availability Data System (GADS) of the North American Electric Reliability Corporation (NERC) in the USA. The increased generator outage rates during extremes of cold or heat and loading have a significant impact on generation adequacy, as for example in the February 2021 blackout in Texas. The increased generator failures are correlated by common external weather rather than by cascading via interactions within the power system.

3.1 Transmission utility outage data

3.1.1 Utility outage data

Transmission systems operators and regulators often systematically collect detailed time-stamped outage data for lines, transformers, and other equipment. For example, the North American Electric Reliability Corporation (NERC) has been collecting North American automatic (momentary and sustained) outage data for transmission elements operating at 200 kV and above in its Transmission Availability Data System (TADS) since January 2008 [16]. While the collection of detailed outage data is standard practice in North America and many other countries, there are only a few publicly available sources for this data [17, 18]. The data typically includes the outage start time, end time, bus name(s), rating, cause codes, and whether the outage is automatic (forced) or planned. The processing of the data typically extracts the automatic outages and neglects the planned outages. Some data such as TADS includes only the automatic outages. There may also be an annual inventory of components in the data.

3.1.2 Extracting events from detailed outage data

A key step in processing real data for cascading and resilience is automatically extracting events. The events are often called cascades or resilience events. These events range in size from isolated single outages to events with hundreds of outages.

One approach to defining and extracting events looks at the outage starting times to detect when they bunch up. This approach was developed to study cascading due to all causes in utility data [20]. Note that outages bunching up due to bad weather is well known in reliability [21, 22]. The event definition is simple: if successive outages start within a time threshold such as one hour, they are in the same event. The time threshold of one hour should be adjusted to the data being processed. Although unrelated outages can start within one hour, it is much more likely that the outages are dependent. The processing simply moves through the outage start times in the order of their occurrence and makes a new event when there is a gap between successive outages of more than one hour.

The outages in a cascade event identified by the bunching can be further divided into generations of outages. The initial outages of the cascade are the first generation; these are followed by the second generation of outages, and so on until the cascade ends. The generations of outages within each cascade arise from the fast time scale of less than a minute of the protection system that implements the outages. Multiple automatic outages occurring simultaneously or in very quick succession are grouped together in the same generation, and in particular, successive outages that start within one minute are in the same generation [20].

The described processing of outages into a series of events, each with one or more generations of outages, structures the outage data for cascading analysis. The majority of the events have only one outage and one generation, but there is considerable interest in examining the larger events, which are rarer but more consequential.

Another approach to defining and extracting events considers how outages overlap in time. Two outages overlap in time if the first outage is not restored when the second outage starts. This approach was developed primarily to study weather-related events in utility data [23, 19, 13]. Motivations for considering overlap include the weakening of the grid while an outage persists and considering outages of some positive duration as more significant or impactful than momentary outages. Consider the set union of the durations of the outages as illustrated in Fig. 3. This set union of the duration times of the outages is composed of disjoint intervals, and each disjoint interval defines the duration of one event. The outages can be processed into these events using the number of unrestored outages [23], which varies with time as outages occur and are restored. The number of unrestored outages or, more usually, the negative of the number of unrestored outages is familiar in resilience [25, 26, 9] and is sometimes called the resilience performance curve. Under normal conditions the number of unrestored outages stays near zero because outages are generally infrequent and are restored quickly. But under stressed conditions, outages are more frequent and accumulate before they can be restored, and the number of unrestored outages has excursions away from zero. These accumulations of outages are the events. The events can be extracted from the data by detecting when the number of unrestored outages passes and returns to zero. In practice, some limitation of the outage duration before processing the data may be necessary to avoid an outage that has an exceptionally long repair time creating an unrealistically long event.

Another approach to find weather events is to identify the time period of bad weather in weather data and then find the outages that occur in that time period [27].

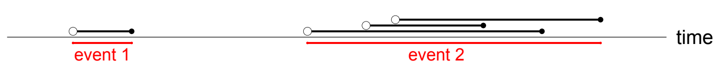

Fig. 3 Two events formed by the union of overlapping outages. For each outage duration shown above the axis, the open circle is the outage start time and the dot is the outage restore time.

Combinations of bunching and overlap and elaborations restricting the geographic region and handling momentary outages may be applied in practice [19, 13, 28, 9]. For example, in [28, 9], for each interconnection, the automatic outages are grouped together into events based on the bunching and overlaps of their starting times and durations. We quote from [28] the algorithm used: “Every outage in an event has to either start within five minutes of a previous outage in the event or overlap in duration with at least one previous outage in the event that has a difference in starting time not exceeding one hour. In applying this algorithm, repeated momentary outages of the same element are neglected if they occur within 5 minutes of each other." Then events that contain at least one outage with a weather-related initiating or sustained cause code are defined as weather-related.

3.2 Statistical patterns in number of outages and generations

The processing of outages into events enables the empirical statistics of the events to be determined, and their statistics augments the heavy tailed behavior of blackout size measured in load shed considered in section 2. The probability distribution of the number of outages in events has a heavy tail and a power law character. For example, Fig. 4 shows the distribution of the number of outages in events with more than 20 outages in 5 years of North American transmission events [28]. Moreover, as shown in Fig. 5, the number of generations in these cascades approximates a Zipf distribution [29, 13], which also has a heavy tail and is a power law distribution. The slope of the dashed line in Fig. 5 has been proposed as the SEPSI index of cascading [29, 13], where SEPSI stands for System Event Propagation Slope Index.

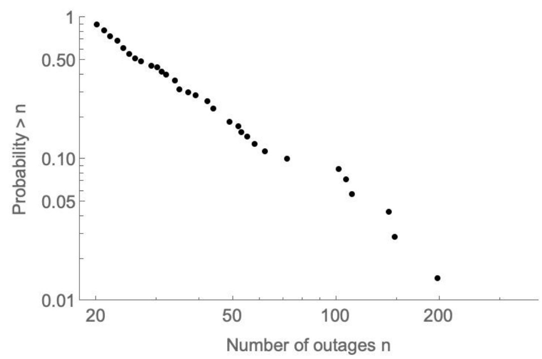

Fig. 4 Log-log plot of survival function of the number of outages in North American transmission events with at least 20 outages. Data are from [28].

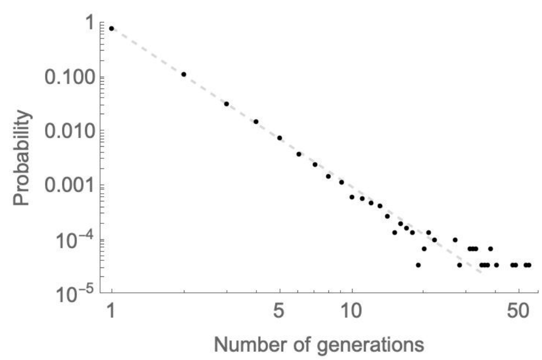

Fig. 5 Log-log plot of probability distribution of number of generations for North American transmission events (dots) with dashed gray line showing the slope of the fitted Zipf distribution. Data are from [13].

4 Probabilistic models directly driven by utility data

Except for the useful analysis of individual blackouts, almost all of the research on cascading is based on models and simulation. Since there are many and varied mechanisms of cascading, and blackouts typically involve complicated combinations of these mechanisms, cascading is particularly hard to model realistically. There is substantial progress modeling a subset of the mechanisms of cascading to produce cascades that could plausibly occur, but many of the most studied mechanisms of cascading are only those with tractable models, and there is much less work giving definitive validation of the cascading models and simulations [57, 58]. And the restoration from cascading blackouts is not often modeled.

Historical blackout data has an obvious key role to ground cascading models and simulations in reality with calibration and validation. But there are further promising and emerging opportunities to directly exploit historical blackout data. To encourage this direction of research, this section shows how historical outage data can enable better contingency lists, be transformed into influence or interaction graphs, replace simulation by sampling, and describe typical blackouts and their restoration with Poisson processes.

4.1 Contingency motifs for multiple outages initiating cascading

When assessing the risk of cascading outages with a simulation, it is usual to sample multiple initial line outages with some sort of equal probability assumption, such as independent, equal probabilities for the individual line outages that make up each multiple outage [30, 31], or equal probabilities for all N-2 outages [24]. One reason is that networks are designed and operated with the N-1 criterion so that outage of only one element in an idealized simulation does not generally initiate a cascade. Analysis of real utility data shows that these equal probability assumptions are pragmatic but unrealistic, making the cascading simulations start from a set of contingencies that are unlikely in practice and making the resulting cascading risk estimates less credible.

A contingency motif [32] is a spatial pattern or small subgraph of the network of multiple outages that occur much more frequently than corresponding multiple outages chosen randomly with equal probability from the utility network. For example, in real data, the N-2 contingency motif  of two lines with a common bus initially outaging occurs much more frequently in initial outages than a random selection of any two lines in the network. Contingency motifs are quite similar to the motifs conventionally considered in network theory [33], except that the frequency of contingency motifs is high compared to a corresponding random selection from the specific power grid being analyzed, rather than high compared to a corresponding random selection from a random graph with similar characteristics. Ren et al. propose conventional network motifs as an indicator of cascading outage risk [34]. They show that phases of cascading outages as the load level increases correspond to the decrease of the frequency of network motifs. The frequency of motifs reflects the connectivity of the power grid; hence, it can be a warning sign of the cascading outage risk. Other researchers have studied conventional network motifs as an indicator of power grid robustness and reliability using techniques from network science [35, 36].

The initial outages of a cascade can be obtained from real outage data by grouping the outages into cascades and generations within each cascade as in subsection 3.1.2. Then the initiating outages are those in the first generation of each cascade. A utility network on which the outages can be located can be deduced directly from the outage data [37], so that the patterns of the initial outages on the utility network can be detected, classified, and counted. This allows the N-k contingency motifs to be obtained, as those that are, for example, ten times more likely to occur in the initial outages than by random selection of k lines from the network. It turns out for two North American publicly available historical outage data sets that the three contingency motifs , ,  account for most of the probability of multiple line outages [32]. Preferentially including these contingency motifs in a contingency list according to their observed frequencies gives a much more realistic sampling of the initiating contingencies and contributes to a better cascading risk analysis via simulation. Of course, for a complete risk analysis, it is necessary to augment the multiple outages with the more tractable single outages, whose probability can be more accurately determined (at least for lines of a given voltage rating or other class), and whose sampling is straightforward.

4.2 Influence/interaction graphs driven by utility data

The way in which line outages influence or interact with each other in cascades of outages can be described by graphs (networks) in which the transmission lines are the nodes and the influences or interactions between outages are the links between the nodes. These graphs are called influence graphs or interaction graphs and are obviously different networks than the physical utility network. Influence/interaction graphs were originally developed in [38, 39, 24] and are reviewed in [40, 41]. Influence graphs can be formed from simulated or utility data, and this section only describes some particular issues of influence graphs driven by utility data.

The first issue with processing utility data into an influence graph is that one is limited to the amount of recorded data. Methods to mitigate this include combining together the data for several generations to estimate the transitions for multiple generations at once while exploiting the overall form of cascade propagation and Bayesian methods [41], and exploiting memory of the interactions between generations of the cascades [42].

The second issue with processing utility data into an influence graph is that real data (usually quantized in time to one minute) has multiple outages occurring in the same minute, which is mainly due to the fast time scale of the protection system. (Simulations, at least so far, avoid this by idealizing single outages occurring in sequence.) This can be addressed either by regarding the multiple outages as additional nodes in the influence graph [41] or by running an expectation maximizing algorithm to allocate the influences among the multiple outages [42].

4.3 Sampling from utility data to replace simulation

Cascading failure analysis is dominated by simulations of models of a subset of the mechanisms of cascading. However, there is another way to generate realistic cascades by computer, which is to sample from the observed outage statistics that characterize how outages propagate and spread on the network. This has been done as part of models of transmission system resilience [43, 44] and to estimate how much cascading follows the initial damage from an earthquake [45]. One can either sample directly for the empirical distributions, or fit a distribution to the empirical distribution and sample from the fitted distribution. The advantage of using a fitted distribution is that it smooths the rarer data and allows some extrapolation. The sampling based on empirical data only captures the overall statistical form of the cascading, but it requires no modeling assumptions, and is very fast compared with model-based approaches. The sampling approach, while necessarily an approximate description, has some sound basis in reality, especially if the empirical data of the network under study is used.

There are two questions to address with sampling cascade outages: How many outages are in the cascade and where are these outages located on the network? One way to characterize the number of line outages in a cascade uses a branching process statistical model that is estimated from the observed data [20]. The propagation parameter from generation k to generation $k+1$ is the expected number of outages in generation $k+1$ that are produced by each outage in generation k. The propagation parameters for all the generations can be directly calculated from observed outage data grouped into cascades and generations. Then the distribution of the number of line outages can be calculated from the number of initial outages and the propagation parameters [20, 45]. Another way is to empirically estimate and sample from the distribution of the total number of outages in cascades [43, 44]. The form of the distribution of the total number of line outages can have the form of a Zipf distribution for two or more line outages [44]. The sampling approach can extend to other aspects. For example, it is also feasible [44] to sample the number of generator outages from observed generator outage statistics [14, 15] and to sample transmission line restoration times from observed line restoration statistics [46].

The spread of the outages on the network can be statistically described by the empirical distribution of network distance between line outages in cascades [43, 45, 44]. To obtain these statistics it is necessary to locate the historical outages on the utility network, and this network can be formed directly from the outage data [37]. After an initial outage or outages are selected on the network, the outages remaining in the total number of outages in the cascade can be located on the network by successively sampling from the network distance distribution. Resampling is used if there are no intact lines available at the sampled network distance.

The sampling approach to cascading is particularly effective when cascading is only one part of the entire calculation and fast samples of cascades are needed. For example, a network model can be used to evaluate the impact of the sampled cascade outages. There is a problem with straightforward sampling of the cascades that the large cascades of the most interest are rare, so that straightforward sampling of cascades is very inefficient in calculating the high impact low frequency cascades. This problem can be solved by stratified sampling with strata of different cascade sizes [47, 44]. Stratified sampling allows sufficient samples at each range of cascade sizes, the results of which are then weighted by the probability of that range of cascade sizes, which in our case is known from the distribution of cascade sizes.

4.4 Statistical models of outage and restore processes

There are several useful ways to describe events by tracking their outages. For each event, the outage process $O(t)$ is the cumulative number of outages at time t and the restore process $R(t)$ is the cumulative number of restores at time t. Both processes start at zero at the beginning of the event and increase to the total number of outages n, as shown in the example with n=12 outages in Fig. 6. Resilience studies

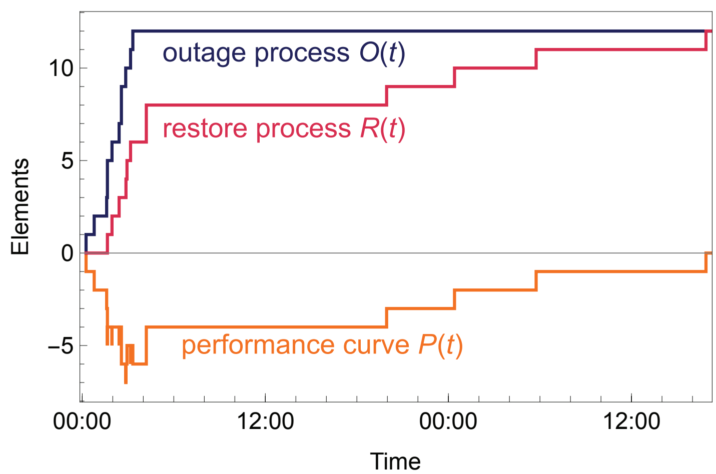

Fig. 6 Processes for a transmission system event with 12 outages. Image from [9] is licensed under CC BY 4.0.

[26, 25, 48, 49] often define for each event a performance (or resilience) curve $P(t)$, which is the negative of the number of *unrestored* outages at time t. The performance curve decrements for each outage and increments for each restore as shown in Fig. 6.

Indeed, the performance curve is related to the outage and restore processes by

$$P(t) = R(t) - O(t).$$

The performance curve can be uniquely decomposed into its outage and restore processes, and it contains the same information as the outage and restore processes [50].

The outage and restore processes and the performance curve, while straightforward, are fundamental to analyzing events in real outage data. Note that the event analysis is at a systems level and is not focused on tracking individual components: it only counts the numbers of outages and restores and it does not track which outaged component restored when or the order in which components restore. Also, the forms of the outage and restore processes and performance curve readily lead to resilience metrics that describe each process; in particular, it is useful to have separate metrics describing the outage process and the restore process and the performance curve [28, 9]. Examples of the metrics are:

Outage process metrics: number of outages, outage duration, outage rate (can be time dependent),

Restore process metrics: restore duration (e.g. time to 95% restore, geometric mean of positive restore times [9]), time to first restore, restore rate (can be time dependent)

Performance curve metrics: area, nadir (maximum components out)

The outage, restore, and performance processes can track quantities other than the number of outages by simply changing the quantity on the vertical axis. For example, the total MVA ratings of the outaged components can be tracked [28, 51]. The list of metrics above assumed that the processes tracked the number of components out. Changing the quantity tracked gives a new set of metrics, except for those metrics quantifying duration. Calculating metrics for utility outage data gives basic and useful information about the various types and magnitudes of outages and blackouts [13, 28, 51, 52]

The outage and restore processes can be stochastically modeled as Poisson processes of time varying rate [53, 54, 9, 55]. Consider an outage process of rate $\lambda_O(t)$ that depends on time t. Then, to first order in a small time interval $[t, t+h]$, the probability of a single outage in $[t, t+h] = \lambda_O(t)h$. The restore process can similarly be modeled as a Poisson process of time varying rate $\lambda_R(t)$. The rates of the outage and restore processes can be obtained from utility data. For example, analysis of seven years of North American transmission data shows that the most typical event with n outages has an outage process of constant rate λ_O over a short interval $[0, o_b]$ and a restore process that starts at time r_a after the start of the outages with a rate proportional to a lognormal distribution:

$$\lambda_O(t) = \lambda_O \tag{11}$$

$$\lambda_R(t) = \frac{n}{(t - r_a)\sigma\sqrt{2\pi}} \exp[-(\ln(t - r_a) - \mu)^2/(2\sigma^2)], \quad t \geq r_a \tag{12}$$

The parameters of these models such as λ_O, r_a, μ, σ can be estimated from utility data [9]. The lognormal distribution of restore rate implies that the rate of restoration typically quickly becomes high and then slows over a long period of time, with the last few restores very delayed. Note that the most typical North American events are dominated by weather events, so it cannot be assumed that the typical cascading event not related to weather is the same as the most typical event. For example, the constant rate outage process of the typical weather event in Eq. 11 is probably not appropriate to cascading events since cascading events show an accelerating rate of outages [56].

One can also consider the mean of the outage and restore processes and the resulting performance curve to obtain a typical event for the data that is processed. We denote these mean processes by $\overline{O}(t)$, $\overline{R}(t)$, and $\overline{P}(t)$. Of course, the Poisson processes vary about their mean value, but the mean value is a representative case. For example, Fig. 7 shows a typical North American transmission event when all the events over seven years are processed.

The mean processes give a nice formula [55] for the area $\overline{A}$ of the mean performance curve $\overline{P}(t)$ tracking the number of components outages. $\overline{A}$ is also the mean of the area A of the performance curve $P(t)$. Moreover, $\overline{A}$ is also the area between the mean outage and mean restore processes (see Fig. 7). Then, writing $\overline{o}$ for the mean outage time and $\overline{r}$ of the mean restore time,

$$\overline{A} = n(\overline{r} - \overline{o}) \tag{13}$$

so that (13) can be understood as the height n of this area times its average width.

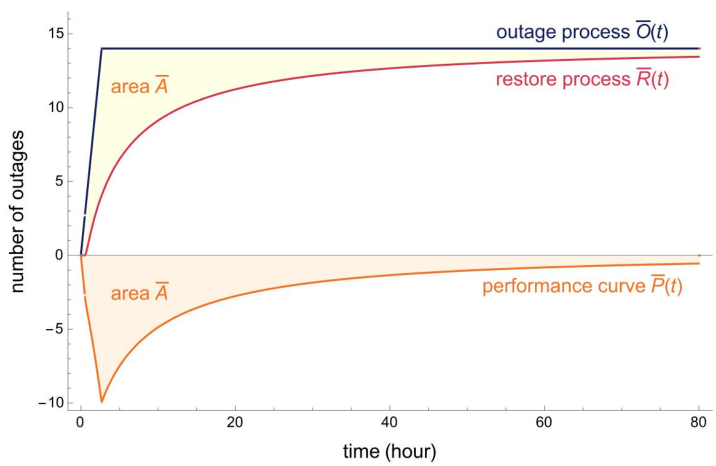

Fig. 7 Typical North American transmission system resilience event tracked by number of outages on the vertical scale. The shaded area $\overline{A}$ of $\overline{P}(t)$ is a useful metric and is the same as the shaded area between $\overline{O}(t)$ and $\overline{R}(t)$. Image from [55] is licensed under CC BY 4.0.

Another way to structure the modeling of events arose in distribution systems [53, 54]. Instead of directly modeling the restore process from data, this approach represents a Poisson outage process arriving at a queue that repairs the outages to produce a restore process. In [54], the response to hurricane Ike is modeled this way using processes that vary in both time and space.

4.5 Validating and calibrating simulations with statistical data

The statistical patterns and metrics that can be extracted from utility data are of fundamental importance in calibrating and validating models and simulations to ensure that they can approximate reality well enough. Overall there has been some progress towards calibration and validation for cascading failure [57, 58, 30, 59, 60], but much remains to be done.

Since cascading outages encounter many thresholds for discrete actions such as tripping a line or not tripping a line, similar simulations with similar data (or even the real power system on successive days) may behave differently under very similar conditions. One simulation or model may trip the line and while another may not and this can have a large effect on the way that a particular cascade evolves. Therefore it is too stringent to require an exact match of the simulation to real blackout data. However, one can compare the statistical patterns or the overall metrics produced by the simulation to the observed statistical patterns and metrics. For example, the distributions of blackout sizes or how much the cascades propagate can be compared.

4.6 Researcher access to utility data and the path forward

There are significant opportunities and much further work to be done processing and analyzing transmission system data. Given the few sources of public detailed outage data [17, 18], one challenge is: How can utilities and regulators be encouraged to share detailed and comprehensive outage data with researchers? One encouraging development is methods to anonymize data. For example, with only some loss of usefulness, bus names can be encrypted and geographic location can be coarse grained. Potential legal liabilities for sharing the data could likely be mitigated by using older data, such as data more than seven years old in the United States. Further collaborative efforts between researchers, regulators, and industry are needed to realize the benefits from the data that has already been collected.

References

1. D. Kosterev, C. Taylor, W. Mittelstadt, Model validation for the August 10, 1996 WSCC system outage, IEEE Trans. Power Systems, vol. 14, 1999, pp. 967-979.
2. V. Venkatasubramanian, Y. Li, Analysis of 1996 Western American electric blackouts, Bulk Power System Dynamics and Control - VI, Cortina d'Ampezzo, Italy, Aug 2004.
3. US-Canada Power System Outage Task Force (2004) Final Report on the August 14, 2003 Blackout in the United States and Canada.
4. Federal Energy Regulatory Commission and the North American Electric Reliability Corporation (2012) Arizona-Southern California Outages on September 8, 2011: Causes and Recommendations.
5. IEEE PES PSDP Task Force on Blackout experience, mitigation, and role of new technologies, blackout experiences and lessons, Best practices for system dynamic performance, and the role of new technologies, IEEE Special Publication 07TP190, July 2007.
6. 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.
7. 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.
8. B.A. Carreras, D.E. Newman, I. Dobson, North American blackout time series statistics and implications for blackout risk, IEEE Trans. Power Systems, vol. 31, no. 6, November 2016, pp. 4406–4414.
9. I. Dobson, S. Ekisheva, How long is a resilience event in a transmission system?: Metrics and models driven by utility data, early access in IEEE Trans. Power Systems, doi: 10.1109/TPWRS.2023.3292328.
10. 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.
11. B.A. Carreras, D.E. Newman, I. Dobson, A.B. Poole, Evidence for self-organized criticality in a time series of electric power system blackouts, IEEE Trans. Circuits and Systems, Part 1, vol. 51, no. 9, September 2004, pp. 1733-1740.
12. D.E. Newman, B.A. Carreras, V.E. Lynch, I. Dobson, Exploring complex systems aspects of blackout risk and mitigation, IEEE Trans. Reliability, vol. 60, no. 1, 2011, pp. 134-143.
13. S. Ekisheva, R. Rieder, J. Norris, M. Lauby, I. Dobson, Impact of extreme weather on North American transmission system outages, IEEE PES General Meeting, Washington DC USA, July 2021.
14. S. Murphy, J. Apt, J. Moura, F. Sowell, Resource adequacy risks to the bulk power system in North America, Applied Energy, vol. 212, p. 1360-1376, 2018 (Also see supplementary information).
15. S. Murphy, F. Sowell, J. Apt, A time-dependent model of generator failures and recoveries captures correlated events and quantifies temperature dependence, Applied Energy, vol. 253, 113513, 2019 (Also see supplementary information).
16. NERC webpage www.nerc.com/pa/RAPA/tads
17. Bonneville Power Administration transmission services operations & reliability, [Online]. Available: https://transmission.bpa.gov/Business/Operations/Outages/.
18. N.K. Carrington, I. Dobson, Z. Wang, Transmission grid outage statistics extracted from a web page logging outages in Northeast America, North American Power Symposium, College Station TX USA, November 2021.
19. M. Papic, S. Ekisheva, E. Cotilla-Sanchez, A risk-based approach to assess the operational resilience of transmission grids, Applied Sciences, vol. 10, no. 14, 2020, 4761.
20. 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, pp. 2146-2155, 2012.
21. R. Billinton, G. Singh, Application of adverse and extreme adverse weather: modeling in transmission and distribution system reliability evaluation, IEE Proc.-Gener. Transm. Distrib., vol.153, no.1, Jan. 2006.

22. R. Billinton, G. Singh, J. Acharya, Failure bunching phenomena in electric power transmission systems, Proc. Inst. Mechanical Engineers Part O: Journal of Risk and Reliability, vol. 220, no. 1, June 2006.
23. N.K. Carrington, S. Ma, I. Dobson, Z. Wang, Extracting resilience statistics from utility data in distribution grids, IEEE PES General Meeting, Montreal, Canada, Aug. 2020.
24. P. D.H. Hines, I. Dobson, P. Rezaei, Cascading power outages propagate locally in an influence graph that is not the actual grid topology, IEEE Trans. Power Syst., vol. 32, no. 2, pp. 958-967, Mar. 2017.
25. M. Panteli, D.N. Trakas, P. Mancarella, N.D. Hatziargyriou, Power systems resilience assessment: hardening and smart operational enhancement, Proc. IEEE, vol. 105, no. 7, July 2017, pp. 1202-1213.
26. A. Stanković et al., Methods for analysis and quantification of power system resilience, IEEE Trans. Power Systems, vol. 38, no. 5, pp. 4774-4787, Sept. 2023.
27. E.A. Morris, K.R. Bell, I.M. Elders, Spatial and temporal clustering of fault events on the GB transmission network. Probabilistic Methods Applied to Power Systems conference, Beijing China, October 2016.
28. S. Ekisheva, I. Dobson, J. Norris, R. Rieder, Assessing transmission resilience during extreme weather with outage and restore processes, Probabilistic Methods Applied to Power Systems, Manchester UK, June 2022.
29. I. Dobson, Finding a Zipf distribution and cascading propagation metric in utility line outage data, preprint arXiv:1808.08434 [physics.soc-ph], 2018.
30. B.A. Carreras, D.E. Newman, I. Dobson, N.S. Degala, Validating OPA with WECC data, in Proc. 46th Hawaii Int. Conf. System Sciences (HICSS), Maui, HI, USA, Jan. 2013, pp. 2197-2204.
31. 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.
32. K. Zhou, I. Dobson, Z. Wang, The most frequent N-k line outages occur in motifs that can improve contingency selection, early access in IEEE Trans. Power Systems, 2023. doi: 10.1109/TPWRS.2023.3249825
33. R. Milo et al., Network motifs: simple building blocks of complex networks, Science, vol. 298, no. 5594, pp. 824-827, Oct. 2002.
34. Q. Chen, H. Ren, C. Sun, Z. Mi, D. Watts, Network motif as an indicator for cascading outages due to the decrease of connectivity, IEEE PES General Meeting, Chicago, Illinois, USA, Jul. 2017.
35. A.K. Dey, Y.R. Gel, H.V. Poor, What network motifs tell us about resilience and reliability of complex networks, Proceedings of the National Academy of Sciences, vol. 116, no. 39, pp. 19368-19373, Sep. 2019.
36. A. Tajer, S. M. Perlaza, H. V. Poor, Advanced data analytics for power systems, Cambridge University Press, 2021.
37. I. Dobson, B.A. Carreras, D.E. Newman, and J.M. Reynolds-Barredo, Obtaining statistics of cascading line outages spreading in an electric transmission network from standard utility data, IEEE Trans. Power Systems, vol. 31, no. 6, November 2016, pp. 4831-4841.
38. P.D.H. Hines, I. Dobson, E. Cotilla-Sanchez, M. Eppstein, “Dual graph" and “random chemistry" methods for cascading failure analysis, Proc. 46th Hawaii Intl. Conf. System Sciences, Maui, HI, USA, Jan. 2013, pp. 2141-2150.
39. J. Qi, K. Sun, and S. Mei, An interaction model for simulation and mitigation of cascading failures, IEEE Trans. Power Systems, vol. 30, no. 2, pp. 804-819, Mar. 2015.
40. U. Nakarmi, M. Rahnamay-Naeini, M. J. Hossain, M. A. Hasnat, Interaction graphs for reliability analysis of power grids: A survey, Energies, vol. 13, no. 9, p. 2219, May 2020, doi: 10.3390/en13092219.
41. K. Zhou, I. Dobson, Z. Wang, A. Roitershtein, A.P. Ghosh, A Markovian influence graph formed from utility line outage data to mitigate large cascades, IEEE Trans. Power Systems, vol. 35, no. 4, July 2020, pp. 3224-3235.
42. J. Qi, Utility outage data driven interaction networks for cascading failure analysis and mitigation, IEEE Trans. Power Systems, vol. 36, no. 2, pp. 1409-1418, March 2021.

43. M.R. Kelly-Gorham, P.D.H. Hines, K. Zhou, I. Dobson, Using utility outage statistics to quantify improvements in bulk power system resilience, Power Systems Computation Conference, Porto, Portugal, June 2020 and Electric Power Systems Research, vol 189, 106676, December 2020.
44. M.R. Kelly-Gorham, P.D.H. Hines, I. Dobson, Ranking the impact of interdependencies on power system resilience using stratified sampling of utility data, early access in IEEE Trans. Power Systems. doi: 10.1109/TPWRS.2023.3260119
45. B. Cheng, L. Nozick, I. Dobson, Investment planning for earthquake-resilient electric power systems considering cascading outages, Earthquake Spectra, vol. 38, no. 3, pp.1734-1760, 2022.
46. S. Kancherla, I. Dobson, Heavy-tailed transmission line restoration times observed in utility data, IEEE Trans. Power Systems, vol. 33, no. 1, January 2018, pp. 1145-1147.
47. F. Faghihi, P. Henneaux, P.E. Labeau, M. Panteli, An efficient probabilistic approach to dynamic resilience assessment of power systems, Congrès Lambda Mu 22 "Les risques au coeur des transitions, (e-congrès)-22e Congrès Maîtrise des Risques Sûreté Fonctionnement, 2020.
48. C. Nan, G. Sansavini, A quantitative method for assessing resilience of interdependent infrastructures, Reliability Engineering System Safety, vol. 157, Jan. 2017, pp. 35-53.
49. S. Poudel, A. Dubey, A. Bose, Risk-based probabilistic quantification of power distribution system operational resilience, IEEE Systems Journal, vol. 14, no. 3, Sept. 2020, pp. 3506-3517.
50. N.K. Carrington, I. Dobson, Z. Wang, Extracting resilience metrics from distribution utility data using outage and restore process statistics, IEEE Trans. Power Systems, vol. 36, no. 2, Nov. 2021, pp. 5814-5823.
51. NERC, 2022 State of reliability, An assessment of 2021 bulk power system performance, July 2022. Available: www.nerc.com.
52. M. Barkakati, A. Pal, A comprehensive data driven outage analysis for assessing reliability of the bulk power system, IEEE PES General Meeting, Atlanta, GA, USA, 2019.
53. C.J. Zapata, S.C. Silva, H.I. Gonzalez, O.L. Burbano, J.A. Hernandez, Modeling the repair process of a power distribution system, IEEE/PES T&D Conf. & Exp.: Latin America, Bogota, Columbia, 2008.
54. Y. Wei, C. Ji, F. Galvan, S. Couvillon, G. Orellana, J. Momoh, Non-stationary random process for large-scale failure and recovery of power distribution, Applied Mathematics, vol. 7, no. 3, 2016, pp. 233-249.
55. I. Dobson, Models, metrics, and their formulas for typical electric power system resilience events, IEEE Trans. Power Systems, vol. 38, no. 6, pp. 5949-5952, November 2023. doi: 10.1109/TPWRS.2023.3300125
56. M. Noebels, I. Dobson, M. Panteli, Observed acceleration of cascading outages, IEEE Transactions on Power Systems, vol. 36, no. 4, July 2021, pp. 3821-3823.
57. IEEE Working Group on understanding, prediction, mitigation and restoration of cascading failures, Benchmarking and validation of cascading failure analysis tools, IEEE Trans. Power Systems, vol. 31, no. 6, November 2016, pp. 4887-4900.
58. P. Henneaux et al., Benchmarking quasi-steady state cascading outage analysis methodologies, Probability Methods Applied to Power Systems, Boise, Idaho, USA June 2018.
59. B.A. Carreras, J. M. Reynolds Barredo, I. Dobson, D.E. Newman, Validating the OPA cascading blackout model on a 19402 bus transmission network with both mesh and tree structures, Fifty-second Hawaii International Conference on System Sciences, Maui, HI January 2019.
60. M. Noebels, R Preece, M. Panteli, AC cascading failure model for resilience analysis in power networks, IEEE Systems Journal, vol. 16, no. 1, pp. 374-385, March 2022.