Cascading Failures: Extreme Properties of Large Blackouts in the Electric Grid

By Paul D.H. Hines, Benjamin O'Hara, Eduardo Cotilla-Sanchez, and Christopher M. Danforth

Introduction

On August 14, 2003, at 1:42pm a power grid operator at the US Midwest Independent System Operator remarked, casually, to an operator at Louisville Gas and Electric, "Hey, what is going on man?" [1] Two and a half hours later, 50 million people across the Northeastern US and Southeastern Canada had lost access to power from the electric grid (see Figure 1). Three years later, on November 4, 2006 at about 9:30 in the evening, German grid operators disconnected a pair of transmission lines across the Ems River, to allow the safe passage of the "Norwegian Pearl" cruise boat. Within a half hour 15 million Europeans were sitting in the dark.

How is it that on relatively ordinary days a few small problems in the US Midwest, or the passing of a cruise boat in Germany, can trigger nearly complete failures of the electricity infrastructures that support some of the world's most technologically advanced societies? Because electric energy is fundamental to almost every aspect of modern societies, understanding cascading failures is vitally important.

Power grids are almost universally agreed to be complex systems, which means that it is not possible to fully understand "the grid" by just looking at its parts. Power grids, which we define here to include all of the physical infrastructure and human individuals and organizations that jointly work to produce, distribute and consume electricity, have many properties that are common to other complex systems. Like the international financial system, power grids are operated by many millions of physical (hardware/software) and human agents. Like the Internet, power systems are frequently subjected to both random failure and malicious attack. Like the weather systems interacting to form hurricanes, there are strong, non-linear connections among the components, and between the components and society at large. And power systems occasionally exhibit spectacularly large, and costly, failures.

This essay attempts to help us to understand these failures by highlighting key mathematical properties of cascading failures in complex systems in general, and in power grids in particular. We focus particularly on the mathematical challenges of measuring cascading failure risk in large power grids, and discuss some techniques that may provide better information to power grid operators regarding cascading failure risk.

Cascading Failures, Risk and Power Laws

According to [2], risk is an exposure to the "chance of injury or loss." Suppose we have some event X (an event where something bad happens). We'll denote the risk associated with X as $R(X)$. We define $R(X)$ to be the probability that this bad event occurs, $\Pr(X)$, times how much the event will cost, $c(X)$. Thus risk of X is

$$R(X) = \Pr(X) \cdot c(X). \tag{1}$$

For example, suppose Jim buys stock in a fictitious widget company for $50. Let X be the event where the stock tanks and Jim loses all his money, and suppose the stock tanks one out of every ten times he buys it. The risk, then, associated with buying the stock is

$$R(X) = \left(\frac{1}{10}\right) \$50 = \$5.$$

Oftentimes, events like X can have varying degrees of severity. For example, assume that one out of ten times Jim loses all his money, one out of ten times he loses half, one out of ten times he loses a quarter, and one out of ten times he loses nothing. In the case of the power grid, severity can be measured by counting the number of people affected by a blackout, or the amount of electric power demand that was disconnected by the blackout. Because the probability of a blackout that affects exactly 52,189 people or interrupts 1000 MW is vanishingly small (and not very useful), we are typically most interested in the risk associated with events falling within a particular *range* of sizes. For example, suppose Jim wants to find the risk of losing somewhere between half and all his money. To find the total risk, Jim adds up the individual risks: the risk of losing half his money and the risk of losing it all. If R_1, P_1, c_1 and R_2, P_2, c_2 are the respective risks, probabilities (1/10 for each event in this case), and costs ($25 and $50), the total risk R_{total} is
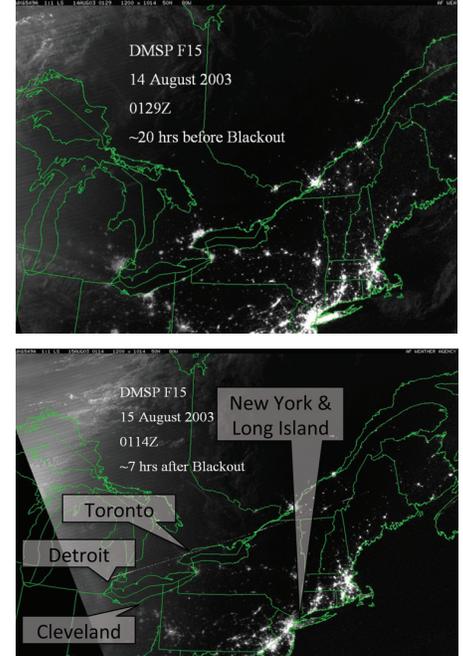
Figure 1. *Night time satellite photographs before and after the blackout of August 14, 2003. Cities that lost most or all of their access to the grid are highlighted at the bottom. 50 million people lost access to electric power as a result of this blackout.*

$$R_{\text{total}} = R_1 + R_2 = P_1 c_1 + P_2 c_2$$
$$= \frac{1}{10}(\$25) + \frac{1}{10}(\$50) = \$7.50.$$

In general, total risk is the sum of the risks associated with a collection of N events,

$$R_{\text{total}} = \sum_{n=1}^{N} \Pr(X_n) \cdot c(X_n).$$

When event sizes can fall anywhere along a continuous space, we replace the probability mass function $\Pr(X)$ with a probability density function $p(x)$, and replace the sum with an integral. Probability density functions tell us the rate at which probability accumulates as x increases. The well-known Gaussian bell curve is probably the most well-known probability density function. For a continuous range of bad things that could happen, with event sizes between S_1 and S_2 the expected risk is:

$$R(S_1, S_2) = \int_{S_1}^{S_2} p(x) \cdot c(x) \cdot dx. \qquad (2)$$

Assuming that we know the probability density function $p(x)$ and the cost function $c(x)$, equation (2) will allow me to evaluate my risk.

In order to understand risk a bit more, let us consider a manufacturing company. Suppose the company owns a factory that produces one million widgets per day. Typically, the factory produces about 50 defective widgets per day, although the number of defective widgets may vary from day to day. When there are no statistical connections between the production of defective widgets, it is common to use a statistical model called the *Poisson* distribution to describe the probability of different numbers of defective widgets. The Poisson distribution tells us the probability of producing x defective widgets in a given day:

$$\Pr(x) = \frac{50^x e^{-50}}{x!}. \qquad (3)$$

Figure 2 shows the shape of this distribution, which is a skewed bell curve with a peak at 50. Suppose the widget company loses \$100 for every defective widget. The cost of producing x defective widgets is

$$c(x) = \$100 \cdot x. \qquad (4)$$

If we want to know the risk of a bad day where 100 or more defective widgets are produced, we can put equations (3) and (4) into the risk formula and add up the risks of 100 to 1,000,000 bad widgets,

$$R(\text{bad day}) = \sum_{x=100}^{1,000,000} \left(\frac{50^x e^{-50}}{x!} \right) (\$100 \cdot x). \qquad (5)$$

Evaluating this sum gives the risk of a bad day as \$0 for all practical purposes. Even a moderately bad day with 70 to 90 defects gives a total expected risk of only about \$10. For comparison, the risk associated with an ordinary day with between 40 to 60 defects is roughly \$4,300. This means that the factory should spend its risk reduction effort to reduce ordinary losses, and not really worry about the very bad days.

When we look into the statistical properties of large blackouts in power grids, we see a very different story. Consider the number of US and Canadian customers who lose electricity service due to large blackouts, which we will define as a blackout that interrupts 50,000 or more customers. Over the period 1984–2006, according to [3,4], blackouts of this sort resulted in about 160 million customer interruptions, or an average of 19,000 customer interruptions per day. If we make the same statistical assumptions as we did in equation (5), namely that the events are well described by a Poisson distribution with $\lambda = 19,000$, the risk associated with large blackouts of 50,000 or larger is roughly zero. Clearly we have made a poor assumption, because large blackouts do happen on a regular basis and they are indeed costly. The problem is that we cannot use standard statistical assumptions when it comes to large blackouts.

To illustrate why, Figure 3 shows the sizes of the largest blackouts in North America from 1984–2006. The horizontal axis represents a blackout of size S, and the vertical axis shows the probability that a randomly chosen blackout will be size S or bigger. Essentially this figure shows us what would happen if we put all of the big blackouts into a hat and then draw out a random one: the horizontal
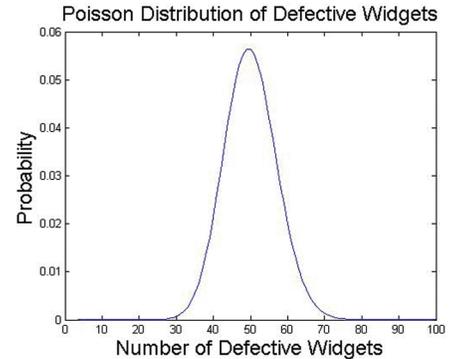
Figure 2. Poisson distribution showing the probability of different numbers of defective widgets. The curve resembles a bell curve with a peak at 50 widgets.
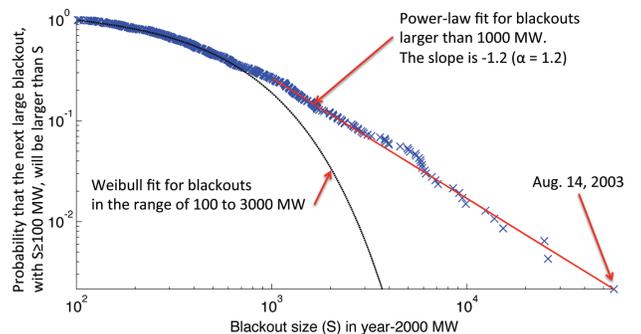
Figure 3. The probability that a large blackout will interrupt service to S or more megawatts. The numbers are scaled to adjust for population growth and growth in electricity demand. Figure adapted from [4]. Power-law first published in [5].

axis shows us the chance of drawing out a blackout of size S or larger. We use a logarithmic scale to highlight the large events and small probabilities. If we try to use exponential statistics like the common bell curve (the Gaussian or Poisson distributions) or the Weibull distribution (as shown), we will grossly underestimate the probability of very large blackouts. In fact, after about 1000 MW, the data are fit far better by a power-law probability distribution, which has a cumulative probability function of the form:

$$\Pr(X \geq x) = \left(\frac{x}{x_{min}}\right)^{-\alpha}, \tag{6}$$

where X is a randomly chosen blackout and x_{min} is a minimum value. The probability density function for (6) is:

$$p(x) = \frac{\alpha x_{min}^{\alpha}}{x^{\alpha+1}}. \tag{7}$$

Power-law probability functions have strange properties, particularly when we are considering risk. To understand why this is important, let's go back to our factory. Assume for the moment that rather than following the Poisson distribution, the number of defective widgets per day follows a power-law distribution, and that the cost is still $100 per defective widget (equation (4)). Let us also assume that the factory continues to produce, on average, 50 defects per day (which gives a power-law exponent of about $\alpha = 0.82$), that defective widgets can be modeled as a continuous variable, and that the factory produces at least one bad widget per day ($x_{min} = 1$). We then can calculate the risk of 100 or more defective widgets from equation (2), using $S_1 = 10^2$ and $S_2 = 10^6$.

$$R(\text{bad day 2}) = \$100 \int_{S_1}^{S_2} x \cdot p(x) \cdot dx \tag{8}$$

$$= \$100 \int_{S_1}^{S_2} x \cdot \left(\frac{\alpha}{x^{\alpha+1}}\right) \cdot dx \tag{9}$$

$$= \$100 \int_{S_1}^{S_2} x^{-\alpha} \cdot dx \tag{10}$$

$$= \$100 \frac{\alpha}{-\alpha+1} [\, x^{-\alpha+1} \,]_{100}^{1,000,000} \tag{11}$$

$$= \$100 \frac{\alpha}{1-\alpha} \left(S_2^{1-\alpha} \right) - S_1^{1-\alpha} \tag{12}$$

$$= \$100 \frac{0.82}{0.18} \left(10^{6 \cdot 0.18} - 100^{0.18} \right) \tag{13}$$

$$\cong \$4,400.$$

The risk of a bad day has gone from $0.00 to more than $4,000. The risk of a very bad day, in which the factory produces 1000 or more defective units is only slightly smaller at about $3,900. The power-law probability function causes the bulk of the risk to be associated with the infrequent, but large losses, in contrast to the Poisson distribution, where the bulk of the risk was caused by small losses. If the factory is a power-law factory, it should focus its risk management efforts not on the ordinary days in which it produces 1–100 defective items, but on the seemingly unlikely days in which the factory produces thousands of defective items.

To further illustrate the strange properties of the power-law factory, let us assume that the factory has no upper bound on the number of defective units that it can produce in a given day.* If the average number of defective items per day remains unchanged at 50 units per day, the power-law exponent for the infinite power-law factory is $\alpha = 1.02$. In this case the risk of producing 1 to 1000 defective widgets is less than $700, even though the factory will only produce more than 1000 defects only once in every 1,151 days. However, the risk of producing 1,000,000 or more defective widgets is nearly $3,800, even though equation (6) tells us that a failure of this size will only occur once every 3,000 years! When power laws are present in risk distributions it becomes extraordinarily important that we pay careful attention to very unlikely events that could have enormous consequences. Given these considerations, the fact that blackout sizes in power grids follow a power law is tremendously important.

When we fit the blackout size data in Figure 3 to equation (6), we get the exponent $\alpha = 1.2$. This is quite close to the critical exponent of $\alpha = 1$ at which the system risk becomes infinite if the size

---

*Note that there are very few cases in which it is reasonable to assume that there are no upper bounds to the size of the risky disaster. However there are some cases where the largest imaginable disaster is very, very large, such as with worst case climate change, or massive collapses in the international financial system.

of the worst failure is unlimited (like the power-law factory that could produce an unlimited number of defective widgets). To see why the risk becomes infinite, consider equation (12), which gives us the risk of a bad day. As long as the power-law exponent $\alpha > 1$, the term $S_2^{1-\alpha}$ will approach zero as S_2 increases toward infinity. If $\alpha < 1$, the risk associated with 2–3 million defects exceeds the risk associated with 1–2 million defects; and the risk of 3–4 million exceeds that of 2–3 million. This pattern leads to infinite risk when all failure sizes are considered. In reality no real system can truly have infinite risk, because real systems are finite in size. The worst-case blackout cannot be larger than the size of an interconnected power grid. For the US, this worst case is about 550,000 MW (the peak load in the Eastern North American grid), which is about 10 times as large as the August 2003 blackout. In many cases blackout sizes follow a power-law for a range of sizes, and then have an exponential tail (like the bell curve) as the event sizes approach the worst case.

Explaining the Power Law

The fundamental question that you may be asking at this point is, "why is it that blackouts have a power-law probability distribution?" While this remains the topic of some debate among those who study these things, it is clear that the fundamental reason is that if a few things fail on a given bad day, there is an increased chance that a lot of things will also fail. The Poisson distribution will provide a good estimate of the chance of failures of various sizes so long as individual failures are statistically independent (i.e., there are no systematic relationships between failures). When relationships among the failures exist the chance of large failures increases. There are two potential sources of these relationships. The first is common cause failures. Large hurricanes, ice storms and earthquakes cause very large blackouts because they damage large numbers of components. The damage caused by hurricanes also follows a power law [6], and, at least in some places, the statistics of blackouts seem to correspond to the statistics of weather and climate [7].

However, weather is not the whole story. When natural disasters are removed from the blackout data, the power-law remains [8]. The second reason for relationships between small events and large events is cascading failure. When a component in the grid gets overstressed, it will tend to disconnect itself from the grid through an automated process. For example, if a generator's automatic control system causes it to begin to rotate too fast, the generator will shut down. If a transmission line has too much current flowing on it, it will heat up, causing the metal conductors to expand and sag. If the underlying vegetation (e.g., trees) is too close, electric current from the line will arc to the tree and a circuit breaker will disconnect the transmission line. When this happens, the stress that was borne by the original component redistributes itself nearly instantly throughout the remainder of the network. The redistribution of stress may cause another component to fail. This process can iterate rapidly and result in very large blackouts. The events of August 14, 2003 in North America and November 4, 2006 in Europe are good examples of cascading failure.

It is clear that cascading failures can cause large blackouts, but cascading failures do not necessarily produce power-law probability distributions in blackout sizes. The leading hypothesis (see [9]) is that competing pressures to both use the electricity infrastructure efficiently to reduce costs, and upgrade the infrastructure when it fails, give rise to a self-organizing process that makes large cascading failures more frequent than they might be otherwise. The real complexity in the system comes from not merely the physics of electricity systems, but from self organization that results from interactions between the physics of cascading failure and the decisions of humans that operate the grid. Even without human involvement, self organization produces power-law failure distributions in other complex systems, like earthquakes, forest fires and landslides [10].

Cascading Failures in Other Complex Systems

Cascading failures exemplify a property common to many complex systems, namely the nonlinear system-wide spreading of a local disturbance. Due to the consequences of nonlinearity, small fluctuations in the pressure of the Earth's atmosphere off the west coast of Africa can result in large changes in the path and strength of North Atlantic hurricanes [11,12]. The introduction of a new species to an ecosystem might trigger ecological changes that ultimately result in massive environmental damage. And internet memes, viral media, and other social phenomena spread rapidly around the world through social networks like Facebook and Twitter.

Of course, there can be considerable difficulties identifying the cause or source of any spreading phenomenon. A match dropped in a dry forest can kindle nearby trees and eventually ignite a massive forest fire. One might correctly conclude that the match was responsible, but overlook the fact that the forest was ready to burn no matter where the spark initiated. This is a characteristic common among many network phenomena, namely that the topology of connections is the dominant factor governing behavior [13].

Traditionally, scientists have used mathematical models to make probabilistic predictions of the future behavior of these complex systems. Numerical weather and climate forecasts rely on sophisticated computer simulations integrating the physics of billions of degrees of freedom, while ecologists

and epidemiologists use differential equations and agent-based models to evaluate potential mitigation scenarios. These modeling efforts can be particularly difficult when the most relevant nodes are incentivized to actively disguise their influence, for example in a terrorist network [14].

In the near future, simulations of large networked systems (like the weather or power grids) will be informed by incredible amounts of real-time observational data. Scientists will couple the data to computational simulations, improve their mathematical models, and potentially improve our understanding of complex phenomena [15]. For example, social scientists studying culture, influence and contagion (e.g., of happiness) can rely less on surveys and small experiments, and more on observing and describing our collective interactions, allowing the data to suggest theories for how we behave [16–19].

Improving our Understanding of Blackout Risk

In the power grid, a similar transition from a data poor to a data rich environment is underway. In the past, a few sensors placed at only the most important locations in the power grid provided just enough data to estimate the operating state of the network about once every minute. Occasionally this state estimation process failed when sensors provided erroneous data, or computer systems malfunctioned. State estimator failure was one of the contributors to the August 14, 2003 blackout [20]. A dramatic transition is now in progress. At the high voltage level, hundreds of sensors (called synchronized phasor measurement units) are being installed to provide massive amounts of detailed measurements about 30 times per second. At the lower voltage level (between your neighborhood substation and your home) power grid operators have until recently had almost no information. If a storm damaged your local power lines, the utility had almost no way to know, other than to wait for you to call them. Now, many electric utilities are installing meters at every home that will report back detailed measurements every 5–15 minutes. This new "Smart Grid" technology might make blackouts more infrequent and/or enable the grid to support more renewable, intermittent power sources like wind and solar. It is, however, possible that the smart grid will have unintended consequences due to complex interactions between the humans and the physical infrastructure of the grid. Smart grid could turn out to be tremendously costly, overwhelm operators and computer systems with raw data, or create new privacy, security and reliability risks.

Therefore, turning this massive quantity of data into information that is useful to both electricity consumers and system operators is an important challenge. The statistical nature of blackouts indicates that methods that seek to reduce blackout risk need to carefully consider not only the everyday small blackouts, but also the very large blackouts that are relatively rare, but can be tremendously costly. Because of this complexity, new mathematical methods are needed to make the electric energy infrastructure of the future not just data rich, but truly intelligent.

References

[1] USHEC, "Blackout 2003: How did it happen and why?: House energy committee meeting transcript." in *US House of Representatives, Energy Committee*, Sept. 3, 2003.

[2] M. Morgan and M. Henrion, *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge: Cambridge University Press, 1990.

[3] Disturbance Analysis Working Group, "DAWG Database: Disturbances, Load Reductions, and Unusual Occurrences," North American Electric Reliability Council, Tech. Rep., 2006.

[4] P. Hines, J. Apt, and S. Talukdar, "Large blackouts in north america: Historical trends and policy implications," *Energy Policy*, vol. 37, pp. 5249–5259, 2009.

[5] 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 the 33rd Hawaii International Conference on System Sciences*, 2000.

[6] W. D. Nordhaus, "The economics of hurricanes and implications of global warming," *Climate Change Economics*, 2010.

[7] D. Xianzhong and S. Sheng, "Self-organized criticality in time series of power systems fault, its mechanism, and potential application," *IEEE Transactions on Power Systems*, vol. 25, no. 4, pp. 1857–1864, Nov. 2010.

[8] 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*, vol. 51, no. 9, pp. 1733–1740, 2004.

[9] I. Dobson, B. Carreras, V. Lynch, and D. Newman, "Complex systems analysis of series of blackouts: cascading failure, critical points, and self-organization," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 17, p. 026103, 2007.

[10] D. L. Turcotte, B. D. Malamud, F. Guzzetti, and P. Reichenbach, "Self-organization, the cascade model, and natural hazards," *Proceedings of the National Academy of Sciences*, vol. 99, pp. 2530–2537, Feb. 19 2002.

[11] D. S. Nolan, E. D. Rappin, and K. A. Emanuel, "Tropical cyclogenesis sensitivity to environmental parameters in radiative–convective equilibrium," *Quart. J. Roy. Meteor. Soc.*, vol. 133, pp. 2085–2107, 2007.

[12] C. Landsea, "How do tropical cyclones form?" [Online]. Available: http://www.aoml.noaa.gov/hrd/tcfaq/A15.html

[13] D. Watts and P. S. Dodd, "Influentials, networks, and public opinion formation," *Journal of Consumer Research*, vol. 34, Dec. 2007.

[14] J. C. Bohorquez, S. Gourley, A. R. Dixon, M. Spagat, and N. F. Johnson, "Common ecology quantifies human insurgency," *Nature*, vol. 462, pp. 911–914, 17 Dec. 2009.

[15] M. Schmidt and H. Lipson, "Distilling free-form natural laws from experimental data," *Science*, vol. 324, no. 5923, pp. 81–85, 3 Apr. 2009.

[16] P. S. Dodds and C. M. Danforth, "Measuring the happiness of large-scale written expression: Songs, blogs, and presidents," *Journal of Happiness Studies*, vol. 11, no. 4, 2009.

[17] D. Centola, "The spread of behavior in an online social network experiment," *Science*, vol. 329, no. 5996, pp. 1194–1197, 2010.

[18] J. Fowler and N. Christakis, "Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the framingham heart study," *British Medical Journal*, vol. 337, p. a2338, 2008.

[19] J.-B. Michel, Y. K. Shen, A. P. Aiden, A. Veres, M. K. Gray, The Google Books Team, J. P. Pickett, D. Hoiberg, D. Clancy, P. Norvig, J. Orwant, S. Pinker, M. A. Nowak, and E. L. Aiden, "Quantitative analysis of culture using millions of digitized books," *Science*, vol. 331, pp. 176–182, 2011.

[20] USCA, "Final Report on the August 14, 2003 Blackout in the United States and Canada," US-Canada Power System Outage Task Force, Tech. Rep., 2004.

P. Hines, Ph.D., (paul.hines@uvm.edu) is an assistant professor in the School of Engineering at the University of Vermont. B. O'Hara (benjamin.ohara@uvm.edu) is an undergraduate student in the Department of Mathematics and Statistics at the University of Vermont. E. Cotilla-Sanchez (eduardo.cotilla-sanchez@uvm.edu) is a graduate student in the School of Engineering at the University of Vermont. C. Danforth, Ph.D., (chris.danforth@uvm.edu) is an assistant professor in the Department of Mathematics and Statistics at the University of Vermont.