The Mathematics of Power Outages

Rachel P. Post

Carthage College

rpost@carthage.edu

April 20, 2016

Abstract

The power grid in the United States suffers from unpredictable blackouts that affect millions of individuals, but what results in a total blackout rather than a simple power outage is not yet known. Through a Mathematica simulation of the power system, we analyze the power grid in the United States in an attempt to discover what factors influence the incidence of blackouts. By utilizing the sandpile theory and the Barabási-Albert model, we aim to uncover what signs can be identified in a simple power station failure that ultimately lead to a blackout. We find that as the failure threshold for a power station varies, there is a phase transition in blackout size, and in addition to this, we also find that the interconnectedness of the network has an unexpected opposing influence on the size and time elapsed between blackouts.

1 Introduction

Power is generated at a power plant, and it then travels to transmission substations and the power substations through transmission lines, and once the voltage is stepped down, it makes its way to the customer whether that be a home, school, or business. Although the general flow of power seems simple, the power grid as a whole is actually a fairly complicated system of stations and connections. The power grid in the United States is made up of over 450,000 miles of high voltage power lines and over 160,000 miles of lower voltage transmission lines which transport the power from power plants to substations and from the substations to customers respectively. All of these power stations and power lines are interconnected throughout the country, which is divided up into three main sections: the Western Interconnection, the Eastern Interconnection, and the Texas Interconnection.

The interconnectedness of the power grid in the United States allows for power sharing to occur, which offers several benefits including reliability, flexibility, and economic competition. The main problem with the system, however, is that energy cannot be stored, so power is generated as it is used. This means that the entire system must respond quickly when there is a shift in demand. As a result, a simple power failure in one portion of the country can very easily spread and affect a huge number of people, showing how the power sharing and interconnection of the power gird can also be disastrous. What results in a total blackout rather than a simple power failure, however, is not yet known, and is the focus of this study.

The initial cause of a power outage can be any number of things in the real world, with weather being the leading cause, but the effect on the components of the power grid is what really matters. When a real world event such as weather causes a power station or transmission line to fail, it will disconnect itself from the grid, causing neighboring stations and transmission lines to have their loads increased to meet the demand, which is a sensible solution unless these neighbors are already operating near their maximum capacity. If this is the case, they will also overload, causing them to fail and disconnect from grid, which in turn causes its neighbors to experience an increase in their load as well. When this happens time after time, a cascading failure occurs in which several power stations and transmission lines disconnect from the grid, resulting in power outages for a large number of people. If this cascading failure is large enough, a significant portion of the country could experience a blackout instead of a small power outage. Therefore, the aim of this study is to determine what factors result in this phenomenon.

2 Definitions and Development

The Sandpile Theory is the basis we will use for simulating how power is shared in the power grid. The basic idea of this theory is best explained through an example. To do this, imagine starting with a pile of sand as the name suggests. At each time step, a grain of sand is added to the pile, and eventually, the pile reaches a threshold or critical point where it can no longer sustain itself, and all of the sand will topple over to its neighbors, which we will call an avalanche.

To illustrate this idea, see Figure 1 below which depicts the initial addition of sand to the grid. In this figure, we start with the blank grid on the left, and suppose sand is randomly added until we reach the grid shown on the right.

0	0	0	0
0	0	0	0
0	0	0	0

→

0	0	0	0
0	4	3	0
0	0	0	0

Figure 1: Sandpile Theory (Initial Addition of Sand)

Now, we will assume that the critical point for this grid is 4, and because one of the piles has reached this threshold, it will spill its sand over to each of its neighbors, resulting in the graphic shown on the right hand side of Figure 2.

0	0	0	0
0	4	3	0
0	0	0	0

→

0	1	0	0
1	0	4	0
0	1	0	0

Figure 2: Sandpile Theory (First Avalanche)

As a result of this process, however, a different pile has reached the critical point of 4, causing it to then spill its sand to each of its neighbors, which we call a cascading avalanche. The result of this cascading avalanche gives us the graphic on the right hand side of Figure 3. At this point, all of the piles in the model are below the critical point, so the grid has become stable. This is the main idea of the Sandpile Theory, and it is what we will use to simulate how power is shared in the power grid.

0	1	0	0
1	0	4	0
0	1	0	0

$\longrightarrow$

0	1	1	0
1	1	0	1
0	1	1	0

Figure 3: Sandpile Theory (Cascading Avalanche)

Definition 1 We define a **network** to be a group or system of interconnected things, specifically a set of nodes or vertices with links or edges connecting them. A synonym for the word network would be a graph. Typically, the term graph is used when we are discussing an abstract object, while the term network is used when we are modeling an application.

For our purposes, we will be using a network to represent the power grid as it allows us to deviate from the lattice structure shown on the left side of Figure 4 below, and represent more complex systems like on the right side of Figure 4.

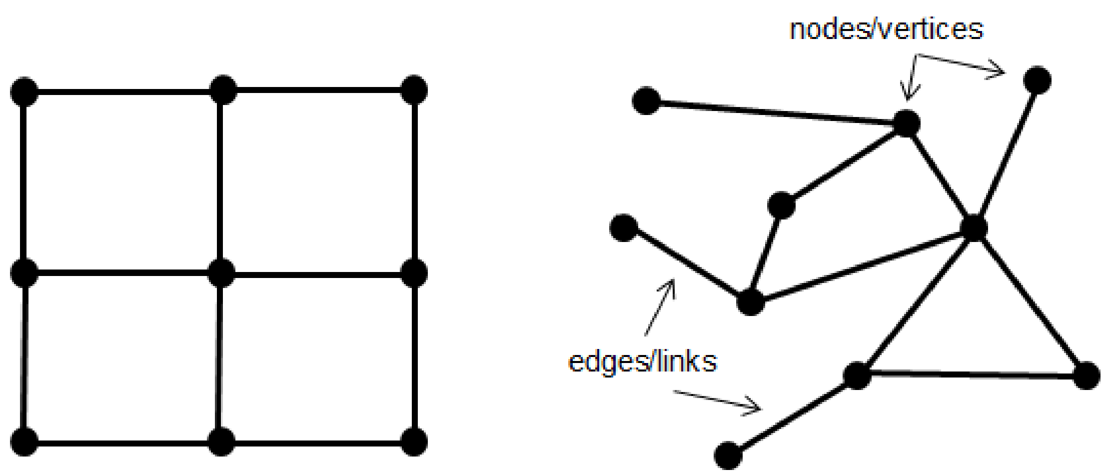

Figure 4: Example of a Network

There are several examples of networks in the real world, and one of these is the World Wide Web in which the vertices in the network represent different web pages, and the edges between them would be the links that get you from one page to the next. Another example is a social network where people are represented by the nodes and there would be connections between nodes that represent people that know one another. In addition, the power grid in the United

States is also nicely represented by a network, so we will be using a network in which the vertices represent power stations, and the edges represent power lines.

Definition 2 The **degree** of a given vertex in a graph or network is the number of edges connected to that vertex.

In order to simulate a real world network, it is necessary to use a model that will accurately represent the interactions between the vertices and edges. Because we want to simulate the power grid, our main concern is that we must use a model that allows for vertices to have different degrees because not all power stations are connected by the same number of power lines. To do this, we will use the Barabási-Albert Model which is named for the two men who proposed the model, Albert-László Barabási and Réka Albert.

Definition 3 The **Barabási-Albert Model** is a model in which connections between nodes are constructed using a specific algorithm. To build a network that follows the Barabási-Albert Model, we start with an initial set of vertices and edges. Then, at each step, a new vertex is added and connected to a given number of vertices that are chosen from all of the vertices in the network with a probability proportional to their degrees.

This means that vertices in the network that already have a lot of edges are more likely to gain another, while vertices with fewer edges are less likely to gain another. This idea is similar to a social network where someone with many friends is more likely to gain more friends than someone with only a few friends. This model produces a network in which each vertex does not have the same degree, as desired.

Description of Simulation Using these ideas, we are able to simulate the power grid in the United States using a program written in Mathematica. Figure 5 below shows an example of what our power grid might look like at the beginning of our simulation. Each vertex in the network represents a power station, and each edge represents a transmission line that gets power from one hub to the next, and eventually to the customer. As you can see in Figure 5, all of the vertices do not have the same number of edges, which is a result of the Barabási-Albert model, as desired.

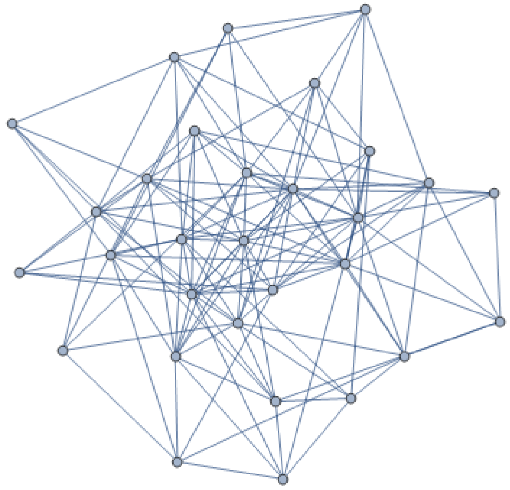

Figure 5: Example of a Random Graph Generated by the Barabási-Albert Model

The graph in Figure 5 was generated with the initial parameters of 30 nodes and a minimum degree for each additional vertex of 2. Throughout the research, these parameters were adjusted to examine how they played a role in the occurrence of blackouts, but this is an example of what one of the graphs would look like with the given parameters.

Before starting the simulation, there are number of inputs that must be entered such as the total number of vertices in the initial network, the number of edges each new vertex gets when it is added to the network, the threshold value for the vertices in the network, and total time for the simulation. Adjusting the total number of vertices simulates adjusting the size of the power grid, changing the number of edges simulates changing how many transmission lines are connecting power stations, and adjusting the threshold changes the point at which a vertex or power station reaches its maximum capacity and must spill its load to its neighbors and disconnect from the grid. Once these variables are defined, the simulation begins.

After the program generates a random network such as the one in Figure 5, time begins to increase, one step at a time. At each time step, a grain of sand is added to a random vertex, simulating the increase in demand for a random power station. This process continues until one of the vertices in the network reaches the threshold that was defined at the beginning of the program. When this occurs, time stops, and the vertex that has reached the threshold point passes its entire load evenly to its neighbors. The neighbors are all of vertices in the network that are connected to the given vertex. Because the load was evenly divided by the number of neighbors, in some cases, this resulted in the neighbors receiving a non-integer amount of sand. This was not an issue, however, since the sand was representing demand. Therefore, for the purposes of this simulation, the non-integer values were used. After the load has been distributed to the neighbors, the given vertex is eliminated from the network.

If in this process, one of the neighbors reaches the threshold point, then it would spill its entire load evenly to its neighbors and then become eliminated from the network as well. On the other hand, if all of the remaining vertices are below the threshold point, then time continues and demand increases again until another vertex reaches the threshold. This process continues until either we reach the full time defined at the beginning of the program, or all of the vertices are eliminated from the network. If all of the vertices are eliminated from the network, we consider this to be a total power blackout.

Variables Measured Once the simulation is complete, there are a number of variables that are analyzed to determine if there is a relationship with blackout occurrence. The first of these is the **percent of the graph remaining**, which tells us how much of our power grid is operating at the conclusion of the program. We are particularly interested in knowing when this variable is zero because if it is, this means that all of the vertices we started with were eliminated from the network, resulting in a total blackout.

Another variable that is checked at the end of the simulation is called **first avalanche**. This variable tells us the time at which the threshold was first reached in the program, so it tells us the time at which the first power station is forced to disconnect from the grid.

The next variable that was measured during each simulation was **time without avalanche** which keeps track of how much time has passed after one avalanche occurs until the next avalanche occurs. Specifically, we want to check when the time between two avalanches is the longest and when the time between two avalanches is the shortest throughout the simulation. This tells us when the model went the longest and the shortest amounts of time without any power stations reaching the threshold and having to disconnect from the grid.

The last variable we looked at after each simulation was called the **biggest avalanche**. This variable measured the largest cascading avalanche that occurred during the simulation, so it keeps track of how many power stations disconnect from the grid consecutively without any additional sand or demand being added to the system.

3 Results

In order to gather all of our data, the simulation was run several times, controlling and testing for each variable so that every combination of variables could be analyzed. For instance, to test how the number of vertices affects the likelihood of a blackout, the simulation was run 100 times with a varying numbers of vertices defined at the beginning of the program while holding all other parameters constant. This method was then repeated for each variable. In general, much of the data collected failed to show interesting patterns between variables, implying that the relationship between those variables is very weak or nonexistent. For example, the data from the simulations in which we analyzed the effect of changing the number of connections on the time of the first avalanche did not show any noticeable trend, leading us to the conclusion that these variables are unrelated in our power grid simulation. A few of the relationships between pairs of variables did give us noteworthy results, however, so we will take a closer look at these relationships and their implications.

The first two noteworthy trends that can be pulled from these simulations come when we look at the amount of graph remaining variable. The two parameters in particular that seemed to have a relationship with the amount of graph remaining at the end of the simulation were the number of vertices and the threshold value. When looking at the number of vertices in the initial network, we found that as the number of vertices at the beginning increases, the percent of the graph remaining when full time was reached also increases (See Figure 6). This means that if the power grid contains more power stations, it is less likely to experience a total blackout. This trend, however, is simply a result of the fact that time was being held constant while the number of vertices was being increased.

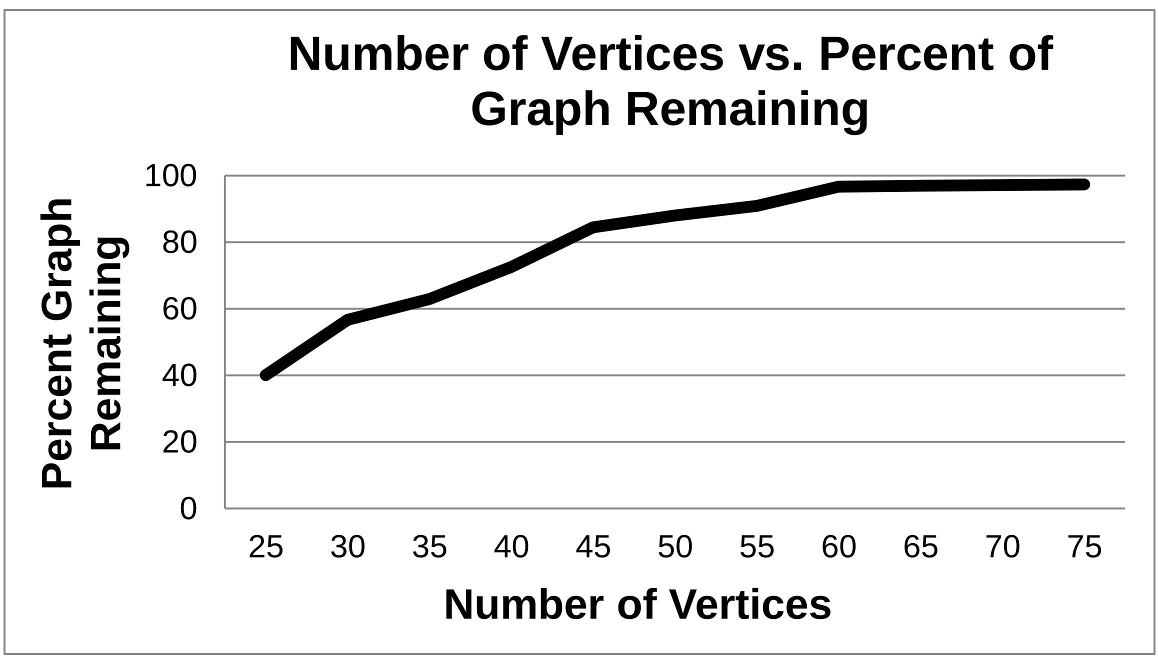

Figure 6: Graph of Number of Vertices vs. Percent of Graph Remaining

A similar trend can be seen when we look at the data involving the threshold values. As the threshold value increases, the percent of the graph remaining when full time was reached also increases (See Figure 7). This means that the higher the threshold is for the power stations, the less likely a blackout is to occur in the grid. What is particularly interesting about this trend, however, is the phase transition that occurs at the threshold value of 20. For any threshold value less than 20, the percent graph remaining at the end of the simulation was 0, and this was true across all variations of all other variables. Regardless of the number of vertices, the number of edges, or the total time, having a threshold value below 20 in our simulation resulted in a blackout, which suggests that the threshold value is of particular importance in the power grid.

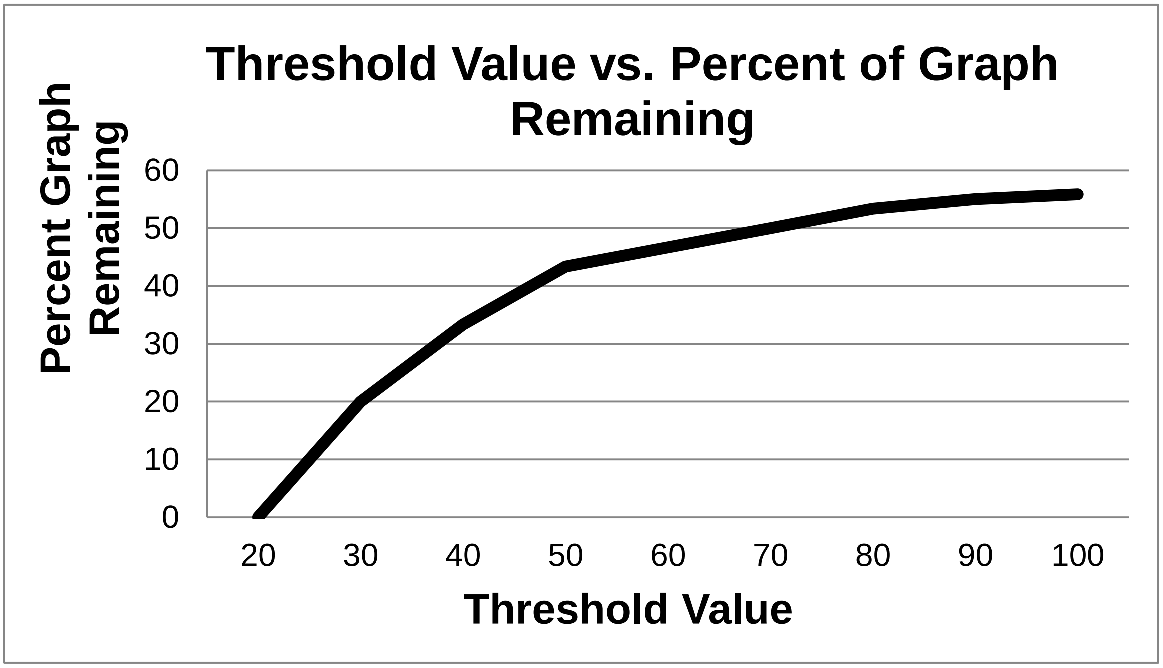

Figure 7: Graph of Threshold Value vs. Percent of Graph Remaining

The next two noteworthy trends come from analyzing the number of edges parameter, which means we are looking at what happens when we change the number of transmission lines between power stations while the actual number of power stations remains constant. Before looking at these results, however, it is important to note the difference in the nature of the biggest avalanche variable and the time between avalanche variable. First, as the size of the biggest avalanche increases, the percent of graph remaining decreases. This indicates that larger avalanches lead to more blackouts. On the other hand, an increase in the time between avalanches results in a greater percent of graph remaining, indicating that more time between avalanches leads to less blackouts. As a result of this, we can see that while an increase in the biggest avalanche variable is undesirable in the power grid, and increase in the time between avalanche variable is actually beneficial for the power grid, which is a distinction that is important when analyzing the next two figures.

First, when considering the relationship between the number of edges and the size of the largest avalanche, we found that increasing the number of edges a vertex gets when it is added to the network leads to an increase in the biggest avalanche size (See Figure 8). Overall, this means that if there are too many connections in the power grid, the size of the largest avalanche will increase, ultimately resulting in a greater likelihood of a blackout occurring.

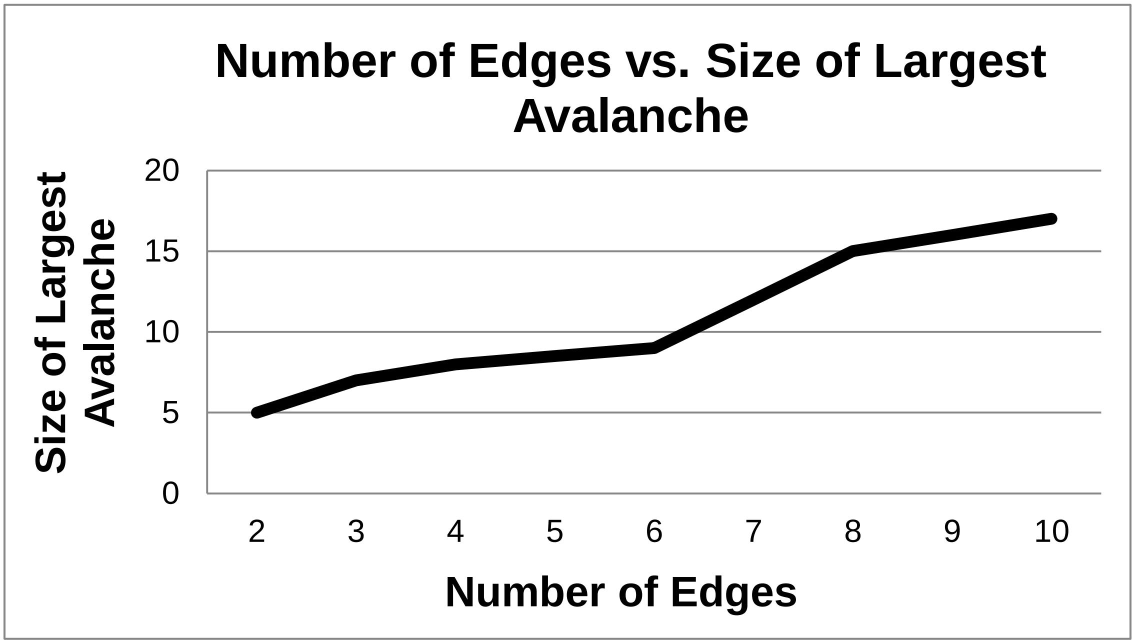

Figure 8: Graph of Number of Edges vs. Size of Largest Avalanche

This implies that reducing the number of edges will reduce the likelihood of a blackout. When we analyze the relationship between the number of edges and the time between avalanches, however, we see a different trend that relies on the distinction made above regarding the nature of our variables. The data from our simulations shows that as the number of edges a vertex gets when it is added to the network decreases, time between avalanches decreases (see Figure 9). This means that if there are too few transmission lines between power stations, cascading failures will occur more often, in turn resulting in the likelihood of a blackout to be greater.

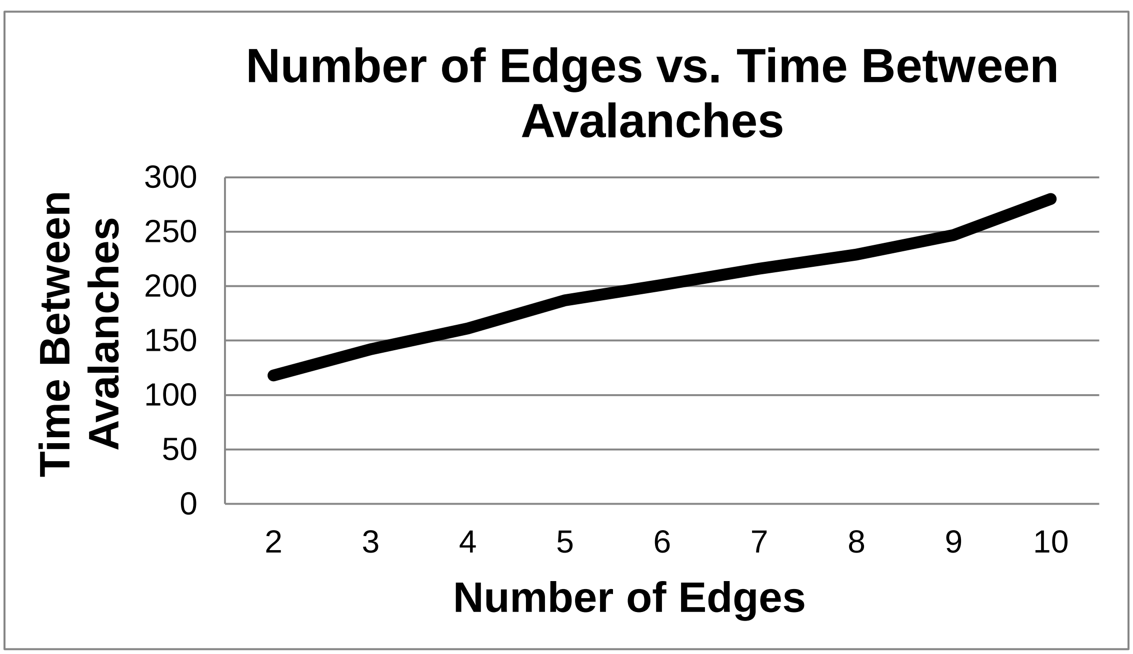

Figure 9: Graph of Number of Edges vs. Time Between Avalanches

Together, these two trends suggest that a balance between too many and too few connections must be achieved in order to prevent large power failures without increasing the frequency of small power failures, which is discussed further in the next section.

4 Conclusion and Directions for Further Research

Ultimately, one conclusion that can be drawn from this study is that increasing the size of the power grid by adding power stations would be an ideal way to reduce the frequency of blackouts, and raising the threshold would also make them less likely. Unfortunately, making these adjustments is not that easy in practice because adding new power stations is not a simple task, and the critical point at which these components operate is not really something that humans have control over. This is because the power grid works based on a concept called self-organized criticality, which is an idea that further research could be focused on. The idea of self-organized criticality says that some phenomena in nature are controlled by a critical point that just occurs naturally and cannot be manipulated. It seems that the critical point at which power stations and transmission lines operate below is one of these critical points. Thus, self-organized criticality, an idea which this study did not look at closely, could be playing an important role in how the United States power grid operates and be worthy of further research.

In addition to this, something interesting seemed to happen when the number of connections was adjusted, which represents changing the number of transmission lines in the power grid. What was particularly noteworthy about this was that having too many connections or too few connections appeared to increase the likelihood of a blackout. This implies that a balance needs to be found in the number of transmission lines in order to avoid the issue of reducing the likelihood of a large blackout but perhaps increasing the frequency of smaller power outages, and of reducing the number of small power outages, but perhaps increasing the likelihood of a large blackout. Once again, this would not be an easy fix to make since it requires restructuring power grid, but it is potentially the most promising option for reducing the occurrence of blackouts in the future.

Additionally, another relationship that could be worth further investigated is the relationship between the number of connections and the size of the largest avalanche. The data gathered in this study did not seem to suggest a pattern between these two variables, but it still seems intuitive to believe that the size of the avalanche is significant and that it is related to the number of connections in the grid. When a large power failure occurs, it makes sense to believe that the number of connections played a role in this phenomenon, so further research into whether or not this is truly the case could be worthwhile.

Lastly, one other direction for future research would be to explore the possibility of addressing the same problem using different models. This study utilized the Barabási-Albert model which used a power law degree distribution when building the network, but there are perhaps other models that use different methods of developing a network that could yield interesting results and lead to other conclusions as well.

References

[1] Albert, R. and Barabási, A-L., Statistical Mechanics of Complex Networks. *Reviews of Modern Physics*, **74** (2002) 47-97.
[2] Barabási, A-L. and Albert, R., Emergence of Scaling in Random Networks. *Science,* **286** (1999) 509-512.

[3] Brain, M. and Roos, D., *How Power Grids Work*. Science – How Stuff Works. http://science.howstuffworks.com/environmental/energy/power.htm
[4] Casrreras, B., Critical Points and Transitions in an Electric Power Transmission Model for Cascading Failure Blackouts. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **12** (2002) 985.
[5] Ellenberg, J., The Amazing, Autotuning Sandpile, *Nautilus*, **023** (2015) Web. 12 Dec 2015.
[6] Hoffmann, H. and Payton, D.W., Suppressing Cascades in a Self-Organized Critical Model with Non-contiguous Spread of Failures. *Chaos, Solitons and Fractals*, **67** (2014) 87-93.
[7] Nedic, D., Criticality in a Cascading Failure Blackout Model. *International Journal of Electrical Power & Energy Systems* **28** (2006) 627-633.
[8] Nykamp, D., *An Introduction to Networks*. Math Insight. http://mathinsight.org/network_introduction.

Appendices

Copy of Mathematica Code:

```
(*Mathematics of Power Outages Simulation*)
(*Rachel Post*)
v=30;
(*set initial number of vertices for the network*)
e=5;
(*set number of edges that each new vertex gets*)
graph=RandomGraph[BarabasiAlbertGraphDistribution[v,e]];
(*create Barabasi-Albert network with above parameters*)
table=Table[0,{i,v}];
(*create variable to check amount of sand/demand at each vertex*)
threshold=50;
time=0;
timestop=2000;
(*set threshold value, initialize time to zero, set end time*)
thresholdcheck=Table[0,{i,timestop}];
(*create variable to keep track of how often the threshold is
reached*)
avalanchesize=0;
(*create variable to track size of largest avalanche*)
While[time<timestop,
 (*run simulation until full time is reached*)
 While[Max[table]<threshold && time<timestop,
(*check if vertices have reached the threshold and if full time
has been reached*)
  x=RandomInteger[{1,v}];
  (*assign a random vertex to the variable x*)
  time=time+1;
  table[[x]]=table[[x]]+1;
thresholdcheck[[time;;timestop]]=thresholdcheck[[time;;timestop]]
+1;
 (*step through time, add sand/demand to the random vertex, and
increase threshold variable*)
While[Max[table]>=threshold,neighbors=AdjacencyList[graph,Flatten
[Position[table,Max[table]]]];
  (*find vertex that reached threshold and all of its neighbors
in the network*)
```

```
  spill=Max[table]/Length[neighbors];
  (*create variable for excess sand*)
  table[[Flatten[Position[table,Max[table]]]]]=0];
 (*set vertex that reached threshold back to zero*)
 VertexDelete[graph,Flatten[Position[table,Max[table]]]]
  (*delete the vertex from the network*)
  For[i=1,
i<Length[neighbors]+1,i++,table[[neighbors[[i]]]]=table[[neighbor
s[[i]]]]+spill];
 (*add spill amount to each neighbor*)
 avalanchesize=avalanchesize+1]
(*increase size of largest avalanche*)
```